Argumentation as a practical foundation for decision theory

Paul-Amaury Matt

Imperial College London
Department of Computing

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Logic and Artificial Intelligence of the University of London.

Abstract

This thesis proposes a methodology for constructing decision making models in decision theory. The method allows decision makers to express their knowledge, beliefs and doubts about decision situations and derive from them a) decision tables, b) utility functions and c) probability distributions, from which optimal decisions can then be computed. The work expounded here is based on two seemingly unrelated theories that recently emerged in the field of artificial intelligence, viz. argumentation theory and the theory of coherent previsions. Decision models can have a dramatic impact on the decisions eventually adopted by decision makers. The methods used to obtain these models must therefore exhibit substantive and normative goodness, i.e. both have practical value and be mathematically coherent. Argumentation has the ability to model human knowledge, expertise and opinions (perhaps to some extent intuition as well) and when carefully employed understandably permits to reach satisfying degrees of substantive goodness. The theory of coherent previsions offers a fully axiomatic and unified view on probability as well as rational decision making: we adopt this theory as a warrant for normative goodness. Various parts of the method are studied in detail and separately applied with success to decision problems pertaining to domains such as finance, trust computing and poker.

Contents

1 Introduction **12**

2 Background **19**

2.1 Decision theory 20

2.1.1 Decision model 20

2.1.2 Decision criteria under strict uncertainty 22

2.1.3 Decision under risk & the expected utility criterion 25

2.1.4 Decision theory: summary 27

2.2 Argumentation for decision making 28

2.2.1 What is argumentation ? 28

2.2.2 General decision making process 30

2.2.3 Intrinsic strength of arguments 32

2.2.4 Conflicts and notions of argument acceptability 34

2.2.5 Decision principles 37

2.2.6 Argumentation for decision making: summary 39

2.3 Summary 40

3 Decision tables & assumption-based argumentation **41**

3.1 Basic influence diagrams 43

3.2 Extended assumption-based argumentation 49

3.2.1 Assumption-based argumentation 49
3.2.2 Extended assumption-based argumentation 51
3.3 Basic influence diagrams as argumentation frameworks 54
3.4 Rationality & the liberal stable semantics 58
3.5 Computing the liberal stable semantics 68
3.5.1 Computing the naive semantics 68
3.5.2 From the naive to the liberal stable semantics 71
3.6 Computing decision tables 75
3.7 Application to the construction of a decision table for poker . . . 77
3.8 Related work 83
3.9 Summary and conclusion 88

4 Argument strength & value functions **89**

4.1 On the construction of value functions 90
4.2 Modeling opinions about goals by abstract argumentation 92
4.3 Strength of arguments in abstract argumentation frameworks . . 97
4.3.1 Rules of the (F, x) game of argumentation strategy 98
4.3.2 Value of the (F, x) game of argumentation strategy 100
4.3.3 Computing the value of the (F, x) game of argumentation strategy and the strength of x 101
4.3.4 Properties of the argument strength measure 105
Elementary properties 106
Properties related to the acceptability of the argument . . 107
Dynamic properties 112
4.4 Social welfare-based aggregation of value arguments 117
4.5 Value functions for poker playing programs 124
4.6 Related work 129
4.7 Summary and conclusion 133

5 Capturing uncertainty with imprecise probabilities **134**

5.1 Background on imprecise probabilities 136

5.1.1 Bayesian probability 136

5.1.2 Confidence intervals for relative frequencies 138

5.1.3 Imprecise probabilities 140

5.1.4 Belief and plausibility functions 144

5.1.5 Coherence of imprecise probabilities 146

5.2 Robust belief and plausibility functions 149

5.3 Modeling expert opinions by abstract argumentation 157

5.4 Combining statistics and expert opinions 160

5.5 Coherence of dialectical probabilities 164

5.6 Properties of dialectical probabilities 170

5.7 Application to trust computing 180

5.7.1 Background on the evidence-based trust model 180

5.7.2 Robust belief functions for trust 181

5.7.3 Arguments for reasoning about trust 182

5.7.4 Dialectical probabilities for trust 183

5.7.5 Performance of dialectical probabilities 185

5.8 Related work 192

5.9 Summary and conclusion 195

6 Statistical inference, risk analysis & decision making with dialectical probabilities **198**

6.1 Background: using imprecise probabilities 199

6.1.1 Statistical inference 200

6.1.2 Decision analysis 200

6.1.3 Financial risk analysis 202

6.2 Statistical inference 203

6.3 Financial risk analysis 207
6.4 Decision analysis 208
6.5 Forcing decisions 212
6.6 Axioms fulfilled by the maximality criterion 216
6.7 Link with decision criteria under strict uncertainty and risk . . . 219
6.8 Application to stock trading 223
6.8.1 Description of the problem 223
6.8.2 Building arguments from technical analysis 227
6.8.3 Decision making 231
6.8.4 Choice of the parameters K and V_A 235
6.8.5 Performance results 236
6.9 Related work 239
6.10 Summary and conclusion 242

7 Conclusion **244**

List of Figures

2.1 General form of a decision table. 21

2.2 Example of a fictitious decision table. 23

2.3 Values associated to each decision under different criteria. 24

2.4 Axioms satisfied by the four classical criteria. 25

3.1 Basic influence diagram for the umbrella problem. 46

3.2 Basic influence diagram for theorem 3.1. 59

3.3 Pathological case of basic influence diagram. 64

3.4 Links between liberal stability and other semantics. 65

3.5 Basic influence diagram where a complete set of assumptions is not liberal stable. 66

3.6 Decision table for the example of Figure 3.2. 76

3.7 Decision table for the umbrella problem. 77

3.8 Poker basic influence diagram (part 1). 79

3.9 Poker basic influence diagram (part 2). 79

3.10 Poker basic influence diagram (part 3). 80

3.11 Poker basic influence diagram (part 4). 81

3.12 Scenarios for the poker decision table. 83

3.13 Poker decision table. 83

4.1 Value and mitigation arguments for the umbrella problem. 96
4.2 Argumentation framework with strength labels. 105
4.3 Equivalent arguments with distinct strength. 110
4.4 Example of attack reducing the strength of an argument. 113
4.5 Adding an attack against a superfluous argument. 114
4.6 Adding an attack against a non-superfluous argument. 114
4.7 Adding a defender to an argument. 115
4.8 Strength of arguments in the umbrella problem. 117
4.9 Framework modeling the uncertainties of a desperate bluffer. . . 127
4.10 Poker decision table. 128
4.11 Value function for the desperate bluffer after a raise. 128
4.12 Best decision for the desperate bluffer (strict uncertainty). 129

5.1 Performance depending on ρ and assuming $\theta = 0$ (no fraud). . . 188
5.2 Influence of θ on the performance of Yu & Singh's belief function. 190
5.3 Influence of θ on the performance of dialectical probabilities. . . 191
5.4 Worst-case performance with respect to θ depending on ρ. 192

6.1 American express stock from August, 11 1998 to August, 8 2008. 225
6.2 Attacks in the stock trading abstract argumentation framework. 230
6.3 Evolution of the gross wealth depending on the probability used. 237
6.4 Risk of holding the stock per unit of money invested. 238

List of Tables

4.1 Examples of argument strength. 105

4.2 Value arguments depending on the psychological and emotional state of the player. 126

4.3 Mitigation arguments depending on the playing scene. 127

6.1 Possible configurations of the abstract argumentation framework. 231

6.2 Strength of arguments depending on the framework configuration. 232

6.3 Performance of the statistical prior, robust beliefs and dialectical probabilities compared to the stock's performance. 237

Acknowledgements

First of all, I would like to express my sincere gratitude to Dr. Francesca Toni, Reader in Computational Logic in the Department of Computing, Imperial College London, who has been my supervisor since the beginning of my study. She has provided me with numerous helpful suggestions, important advice and constant support during the course of this work.

I also wish to express my cordial appreciation to Prof. Anthony Hunter of the Department of Computer Science, University College London, Prof. John Fox of the Department of Engineering Science, University of Oxford and Prof. Phan Minh Dung of the Department of Computer Science, Asian Institute of Technology who gave constructive advice, providing important ideas and useful discussions on the topics of argumentation and subjective probability.

My very keen appreciation goes to the members of the Operations Research and Computational Finance Group of Imperial College London and in particular to Prof. Berç Rustem, Dr. Daniel Kuhn, Wolfram Wiesemann, Steve Zymler and Angelos Tsoukalas for their precious assistance in the fields of decision theory, optimisation, game theory and finance.

I would like to express my sincere gratitude for financial support from ARGUGRID as well as the European Commission which has funded this project under the Framework VI.

Finally, my special appreciation and heartiest thanks go to my parents, Marie-Claude and Dominique Matt, for their unconditional support and encouragement to pursue my research interests.

Chapter 1

Introduction

In 1956, John McCarthy coined the term *artificial intelligence*, which he defined as "the science and engineering of making intelligent machines". According to [Arvey, 1994], the term *intelligence* denotes "the general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience. It is not merely book learning, a narrow academic skill, or test-taking smarts. Rather, it reflects a broader and deeper capability for comprehending our surroundings, catching on, making sense of things, or figuring out what to do". An important branch of artificial intelligence [Russel and Norvig, 2002] is thus concerned with decision making, which can be essentially regarded as the mental process leading to the selection and adoption of a course of action amongst a number of alternatives.

In decision under risk, which constitutes the main branch of decision theory [Fishburn, 1964; Raiffa, 1968; French, 1987], a decision making problem is modeled by a decision table, a utility function and a probability distribution. The decision table represents the decisions that can be taken and all the pos-

sible outcomes of each decision under different scenarios. The utility function quantifies the desirability of each outcome from the decision maker's viewpoint. The probability distribution finally describes the chance of occurrence of the scenarios. Decision theory offers several criteria to select the best decision, such as Wald's maximin return, Hurwicz's optimism-pessimism index, Savage's minimax regret and Laplace's principle of insufficient reason for decision making under strict risk (when no probabilistic information is available) or the expected utility criterion in decision under risk (when the distribution of probability of the scenarios is known precisely).

For very simple decision making problems, constructing a decision model is not an issue. However, in general, modeling a decision making problem may constitute a complex task in itself. The inputs required by decision theory are sufficiently elaborate and far too critical to be improvised. The construction of a substantively good model (having practical value) requires knowledge or even expert knowledge in the domain of application as well as the ability to reason with qualitative, incomplete and often conflicting sources of information. The main objective of this thesis is thus to develop scientific methods for supporting decision makers in modeling decision making problems. The task consists in capturing all pieces of information relevant to the problem and processing it in a rigorous and automated way, so as to eventually obtain models having both substantive and normative goodness (mathematical coherence). Our approach is principally based on two modern theories of artificial intelligence, viz. argumentation theory and the theory of coherent previsions.

We will follow the plan described here-after. Chapter 2 introduces basic elements of decision theory and gives an overview of argumentation-based decision making. Decision theory is good from a normative standpoint and constitutes an elegant computational theory of decision making. It however completely

relies on the assumption that the decision maker is able to build an abstract mathematical model of decision making problems. Argumentation on the other hand appeals to concrete pieces of information: the underlying beliefs of the decision maker and his reasons for believing that a decision is either good or bad. The tools developed in this thesis are based on simple ideas and material originating from the field of argumentation, but their ultimate purpose will only be to ease the process of constructing models for decision theory.

Chapter 3 addresses the problem of constructing decision tables. Decision tables encapsulate a lot of information. They show all the possible decisions the decision maker can choose from. They also describe all the possible outcomes of these decisions, under all possible scenarios. In decision under strict uncertainty, a decision model consists of a decision table plus a utility function. The utility function is itself defined over the set of possible outcomes and measures the desirability of these outcomes. In decision under risk, the decision model additionally requires the specification of a probability distribution, which is defined over the set of scenarios of the decision table. Therefore, decision tables constitute a critical element of the model. In order to help the decision maker constructing a decision table, we introduce a logical and semi-graphical tool called *basic influence diagram* for describing situations of decision making under strict uncertainty. Here, we ignore information concerning either the desirability of outcomes (utility function) or the likelihood of scenarios (probability distribution). We propose an assumption-based argumentation technique [Dung *et al.*, 2006] for reasoning about these diagrams and a semantics called *liberal stable semantics* to identify all the possible decisions, scenarios and outcomes. An existing hypergraph algorithm [Eiter and Gottlob, 1995] is then employed for computing these outcomes and generating decision tables automatically. As a concrete example of application, we construct a decision table for the game of

poker [Mendelson, 2005].

Chapter 4 investigates the possible role of argumentation regarding the construction of utility functions. After decision tables, utility functions are in decision theory the next most important element of decision models, as probability distributions are only required for decision making under risk. The mathematical term of utility function is normally reserved to denote real-valued functions representing the preferences of a decision maker over lotteries with random outcomes [Fishburn, 1964]. We are not concerned with such complex objects in this thesis. Instead, we consider real-valued functions defined over sets of possible outcomes. A function which assigns to each possible outcome a real number is called a *value function.* Only value functions are truly needed in decision theory [Fleming, 1952]. We will present a method for building value functions based on *judgements of desirability* concerning outcomes. These judgements take the form of *abstract arguments* [Dung, 1995] and are interpreted as justified claims [Krause *et al.*, 1995]. Some arguments correspond to claims concerning the degree of desirability of outcomes. The other arguments are used to challenge (undermine) such claims and more finely appreciate their validity. To this effect, we develop a measure of *argument strength* using elements of abstract argumentation [Dung, 1995] and game theory [von Neumann, 1928]. We apply the mathematical formalism of *linear social welfare* functions introduced first by [Fleming, 1952] and later developed by [Harsanyi, 1955] and [Dyer and Sarin, 1978] to aggregate all arguments (according to their strength) into a single value function. We show as an application how to generate value functions for various poker players.

Chapter 5 proposes a method for assessing the probability of occurrence of scenarios and more generally *events* (sets of scenarios), by combination of statistical data and the opinion of experts. In that chapter, opinions are also

modeled using the framework of abstract argumentation and the resulting probabilities are consequently given the name of *dialectical probabilities*. We distinguish between *forecast arguments*, whose role is to support the occurrence of events, and *mitigation arguments*, which are used to mitigate forecast arguments and indirectly, to model conflicts between expert opinions. The measure of argument strength developed in Chapter 4 is used here again to assess the validity of arguments. Dialectical probabilities are proved to have the intuitive form of maximal ϵ-contaminations: one of the most basic types of imprecise probabilities, introduced in Robust Bayesian Analysis [Huber, 1981; Berger and Berliner, 1986] in the 80s. We show that dialectical probabilities also correspond to a special class of belief functions [Shafer, 1976] – a modern uncertainty model that generalises the Bayesian approach and has allowed the recent emergence of many expert systems in artificial intelligence. We show that dialectical probabilities are mathematically coherent in the sense of [Walley, 1991]. We study their properties with respect to argumentation theory as well as generalised uncertainty and information theory [Klir, 2006]. As a first application, we show their efficiency as a computational model of trust in service-oriented computing.

In Chapter 6, we use the property of coherence of dialectical probabilities and the theory of coherent previsions [Walley, 1991; 1996; 2000] to derive a practical framework for statistical inference, risk analysis and decision analysis. Such a framework is needed since decision theory and statistics both assume a Bayesian model of probability, whereas dialectical probabilities are usually non-Bayesian. In order to better understand the decision criteria used for selecting the best decisions under dialectical probabilities, we establish formal links between them and the classical decision criteria for decision making under risk or strict uncertainty. As an application, we show that dialectical probabil-

ities derived from arguments based on technical analysis can lead to significant performance improvements in financial forecasting and stock trading.

We summarise our work and conclude in Chapter 7.

The novelty of this thesis is first to propose *argumentation as a foundation for decision theory.* By this, we mean developing tangible *argumentation-based methods* for obtaining a) *decision tables*, b) *value functions* and c) *probability distributions* as required by decision theory. By doing so, we hope to increase the applicability of decision theory, intend to clarify and strengthen the connection between decision theory and argumentation and also aim at promoting some of the ideas and principles employed in argumentation. Second, we introduce *dialectical probabilities* as a *new type of non-Bayesian probabilities.* We notably provide a concrete recipe for *aggregating statistical data with expert opinions* (seen as arguments) and *quantify experimentally the added value of argumentation* in two important statistical problems, namely trust computing and financial forecasting. Third, we explore and *revisit the classical decision-theoretic criteria* used either in decision making under strict uncertainty or risk in the *general context where probabilities are only imprecisely known.* We thus also propose *new mathematical closed-form formulae for robust statistical inference, decision analysis, trust and financial risk analysis* as well as *algorithms for decision making and robust optimisation.* Finally, we contribute to the field of argumentation, by proposing a new *dialectical measure of argument strength*, a new notion of argument acceptability called *liberal stable semantics*, by showing how to practically *derive arguments in technical analysis* to forecast the prices of stocks and by demonstrating the *value added by arguments in the estimation of subjective probabilities.*

Several of the results obtained in this thesis have already been published in scientific conferences. The results of Chapter 3 concerning the use of assumption-

based argumentation in decision making under strict uncertainty and the properties of the liberal stable semantics have been published in [Matt and Toni, 2008a] at the Second International Conference on Computational Models of Arguments (COMMA 2008). Chapters 4 and 5 rely on the same dialectical measure of argument strength. The foundations, construction and properties of this measure have been published in [Matt and Toni, 2008b] at the Eleventh European Conference on Logics in Artificial Intelligence (JELIA 2008).

Chapter 2

Background

In this chapter, we introduce two very distinct theories of decision making in artificial intelligence, namely decision theory and argumentation for decision making. Both theories are interesting and have their advantages. Decision theory is abstract, general, quantitative and fully axiomatic. It is in other words a coherent mathematical theory of decision making. Unfortunately, its normative goodness is not sufficient to make it practically valuable. It is probably in response to the abuse of abstraction in academic research on decision making that several groups of researchers have started to develop an interest in more qualitative and practical approaches to decision making. Argumentation for decision making – although arguably not fully mature yet – provides an excellent account of what a practical and substantively good theory of decision making should look like. In the first section of this chapter, we provide a brief and formal introduction to decision theory. This part is necessary to understand what this thesis aims at. In the second section, we present the main ideas behind argumentation for decision making and introduce thereby a lot of material that will be re-used throughout the following chapters. Note that some of these chapters

contain additional background on topics such as assumption-based argumentation, hypergraphs, game theory, value theory, Bayesian probability, statistics, imprecise probability or the theory of coherent previsions.

2.1 Decision theory

2.1.1 Decision model

Decision theory [Fishburn, 1964; Raiffa, 1968; French, 1987] is the standard theory of decision making in artificial intelligence [Russel and Norvig, 2002]. The theory aims at identifying the best decision to take. It assumes that the decision maker is able to specify an abstract *decision model* in terms of

- a decision table,
- a utility function, and
- a probability distribution.

A *decision table* is a table of the form shown in Figure 2.1. The lines of a decision table correspond to the *possible decisions* $d_1, \ldots, d_m$, where it is assumed that $m \geq 2$. The decision maker must choose one of them. The columns of the table are *scenarios* $w_1, \ldots, w_n$, where $n \geq 1$. Scenarios occur randomly and are uncontrollable. The set of scenarios $\{w_1, \ldots, w_n\}$ is denoted Ω (omega) and called the *universe*. It is assumed that the universe is exhaustive (one of the scenarios will occur) and that the scenarios are mutually exclusive (only one scenario can happen). Scenarios are said to be random because the decision maker is not certain exactly which one of them will occur. The consequence of a decision is called an *outcome* and depends on which scenario occurs. By convention, the outcome of decision d_i under scenario w_j is denoted by the

symbol $x_{i,j}$ and appears in the cell of the table at the intersection of line i and column j.

.	w_1	w_2	...	w_j	...	w_n
d_1	$x_{1,1}$	$x_{1,2}$	...	$x_{1,j}$	...	$x_{1,n}$
d_2	$x_{2,1}$	$x_{2,2}$	...	$x_{2,j}$	...	$x_{2,n}$
...	...	...	...	...	...	...
d_i	$x_{i,1}$	$x_{i,2}$	...	$x_{i,j}$	...	$x_{i,n}$
...	...	...	...	...	...	...
d_m	$x_{m,1}$	$x_{m,2}$	...	$x_{m,j}$	...	$x_{m,n}$

Figure 2.1: General form of a decision table.

Situations where the number of scenarios is $n = 1$ are referred to as decision making under *certainty*. Decision making under *strict uncertainty* on the opposite corresponds to cases where $n \geq 2$ and the decision maker knows nothing about the likelihood of the scenarios. Finally, decision making *under risk* denotes to situations where $n \geq 2$ and a Bayesian probability distribution defined over the set of scenarios is available / known. A *probability distribution* is a function $P : \Omega \to [0, 1]$ such that $P(w_1) + \ldots + P(w_n) = 1$. For every scenario $w \in \Omega$, the value $P(w)$ is interpreted as the chance of occurrence or likelihood of the scenario w from the decision maker's viewpoint. In the remainder, we will refer to probability distributions as *Bayesian* to distinguish them from *imprecise probabilities* studied in Chapter 5.

In decision making under certainty, strict uncertainty and under risk, it is systematically assumed that the decision maker is able to quantify the desirability of each outcome $x_{i,j}$. Decision theory thus assumes given a function u called a *value function* associating to each outcome $x_{i,j}$ a level of desirability $u(x_{i,j}) \in \mathbb{R}$ (or simply $\mathbb{Q}$ or $\mathbb{Z}$ so as to avoid issues related to the representation of irrational numbers). The level of desirability of an outcome can be either positive (the outcome is seen as desirable) or negative (the outcome is undesirable). This function u represents the preferences of the decision maker over the

set of possible outcomes. We would for instance say that

- $x_{i,j}$ is preferred to $x_{k,l}$ whenever $u(x_{i,j}) \geq u(x_{k,l})$,
- $x_{i,j}$ is strictly preferred to $x_{k,l}$ whenever $u(x_{i,j}) > u(x_{k,l})$, and
- $x_{i,j}$ and $x_{k,l}$ are equally preferred whenever $u(x_{i,j}) = u(x_{k,l})$.

2.1.2 Decision criteria under strict uncertainty

Which decision should be chosen by the decision maker ? In a situation of decision under certainty, the best decision d_i is clearly the one with highest value $u(x_{i,1})$. In decision under strict uncertainty, several criteria for choosing the best decision have been proposed. Usually, a criterion takes the form of a formula that allows to assess numerically the value of a decision, independently of which scenario will occur. The result is a score that is used to rank decisions. The most famous and widely used criteria for decision under strict uncertainty are defined as follows. We will eventually come back to these criteria in Chapter 6 and show that the decision criteria used for making decision under imprecise probabilities, and in particular with the dialectical probabilities defined in Chapter 5 relate to these criteria when the probabilities are maximally imprecise, or in other words, when no probabilistic information is available at all. In the following definitions and the remainder of this thesis, we denote the value $u(x_{i,j})$ of an outcome $x_{i,j}$ simply by the symbol $u_{i,j}$.

Definition 2.1. *Define the security level of decision d_i by $s_i = \min_{j=1}^{n}\{u_{i,j}\}$. **Wald's maximum return** criterion consists in choosing the decision with maximal security level.*

Definition 2.2. *Define the level of optimism of d_i by $o_i = \max_{j=1}^{n}\{u_{i,j}\}$. Given a value $x \in [0,1]$ called the optimism-pessimism index, **Hurwicz's optimism-***

pessimism index *criterion consists in choosing the decision d_i that maximises $xs_i + (1-x)o_i$, where s_i is the security level of d_i.*

Definition 2.3. *Define the regret of an outcome $x_{i,j}$ as $r_{i,j} = \max_{k=1}^{m}\{u_{k,j}\} - u_{i,j}$. Define the regret of a decision d_i as $\rho_i = \max_{j=1}^{n}\{r_{i,j}\}$.* ***Savage's minimax regret*** *criterion consists in choosing the decision with minimal regret.*

Definition 2.4. ***Laplace's principle of insufficient reason*** *criterion consists in choosing the decision d_i that maximises $\sum_{j=1}^{n} \frac{1}{n} u_{i,j}$.*

As an illustration, consider the fictitious decision table displayed in Figure 2.2, where each outcome $x_{i,j}$ has been directly replaced by the value of $u(x_{i,j})$. The calculations of the values of the decisions are given in Figure 2.3. Wald's criterion leads to choose d_2 with maximal maximin return of 1. Hurwicz's criterion leads to choose d_3 for $x = 0$ and $x = 1/2$ with respective optimism-pessimism indexes of 4 and 3/2. Savage's criterion leads to choose d_4 with minimal regret of 1 and finally Laplace's principle of insufficient reason leads to choose d_1 with maximal expected value $\sum_{j=1}^{n} p_j u_{i,j}$ of 5/4 under the uniform probability distribution ($p_j = 1/n$ for every j). We see that these four criteria can disagree and lead to the choice of completely different decisions.

.	w_1	w_2	w_3	w_4
d_1	2	2	0	1
d_2	1	1	1	1
d_3	0	4	0	0
d_4	1	3	0	0

Figure 2.2: Example of a fictitious decision table.

In decision making under strict uncertainty, it is quite important to carefully select the criterion. The selection is usually made by looking at the axioms the criteria fulfil and depends on the practical decision problem considered. Eight important axioms for rational decision making are listed here-below.

.	Wald	Hurwicz ($x = 0$)	Hurwicz ($x = 0.5$)	Savage	Laplace
d_1	0	2	1	2	5/4
d_2	1	1	1	3	1
d_3	0	4	2	2	1
d_4	0	3	3/2	1	1

Figure 2.3: Values associated to each decision under different criteria.

Axiom 2.1 (complete ranking). *The criterion provides a complete ranking (i.e. a total pre-order* [1]*) of all the possible decisions.*

Axiom 2.2 (independence of labeling). *The strict preference of one decision over another should not depend on the order of the decisions and scenarios in the decision table.*

Axiom 2.3 (independence of value scale). *The strict preference of one decision over another should remain the same after application of an affine transformation of the value function.*

Axiom 2.4 (strong domination). *If the value of the outcome of a first decision is strictly greater than the one of a second decision under every scenario, then the first decision should be strictly preferred to the second one.*

Axiom 2.5 (independence of irrelevant alternatives). *The strict preference of a decision over another should not depend on any other possible decision.*

Axiom 2.6 (independence of addition of a constant to a column). *The strict preference of one decision over another should be invariant to the addition of a constant to the value of outcomes under a specific scenario.*

Axiom 2.7 (independence of column duplication). *The ranking of decisions should remain the same after duplication of a column in the decision table.*

[1] A total pre-order is a binary relation that is transitive, total and reflexive.

Note that apart from axioms 2.2, 2.3 and 2.4 which are considered uncontroversial, the other axioms are not universally accepted by the scientific community (the reader may refer to the first chapter of [French, 1987] for a discussion on that topic). Concerning for example axiom 2.5, some economists think it sometimes makes sense to evaluate a decision relatively to one or more other decisions. We also support this idea and will prove in Chapter 6 that the main decision criterion for making decisions under imprecise probabilities violates axiom 2.5 (but satisfies all the other ones). Anyway, as Figure 2.4 shows, none of the classical criteria satisfies all seven axioms. It has been shown in fact that a criterion that satisfies them all could not even exist.

Axiom versus criterion	Wald	Hurwicz	Savage	Laplace
complete ranking (axiom 2.1)	√	√	√	√
ind. of labelling (axiom 2.2)	√	√	√	√
ind. of value scale (axiom 2.3)	√	√	√	√
strong domination (axiom 2.4)	√	√	√	√
ind. of irrelevant alternatives (axiom 2.5)	√	√	×	√
ind. of addition of a constant (axiom 2.6)	×	×	√	√
ind. of column duplication (axiom 2.7)	√	√	√	×

Figure 2.4: Axioms satisfied by the four classical criteria.

2.1.3 Decision under risk & the expected utility criterion

In decision under risk, a probability distribution P is given. Taking a decision d_i is then equivalent to entering some kind of non-deterministic game where the possible prices are the outcomes $x_{i,1}, \ldots, x_{i,n}$ and the probability of winning them are respectively $P(w_1), \ldots, P(w_n)$. This kind of game in which prices are random constitutes what is called a *lottery*. In general, a lottery with prices $x_1, \ldots, x_r$ and respective probabilities $p_1, \ldots, p_r$ is denoted

$\langle p_1, x_1; p_2, x_2; \ldots; p_r, x_r \rangle$. The lottery associated with decision d_i is thus

$$\langle P(w_1), x_{i,1}\,;\, P(w_2), x_{i,2}\,;\, \ldots\,;\, P(w_n), x_{i,n} \rangle$$

The utility theory developed by [von Neumann and Morgenstern, 1944] uses three axioms to define the rationality of a preference relation $\succeq$ over lotteries

- completeness: for any two lotteries L and L', either $L \succeq L'$, $L' \succeq L$ or both ($L \approx L'$)
- transitivity: if $L \succeq M$ and $M \succeq N$, then $L \succeq N$
- Archimedean: if $L \succeq M$ and $M \succeq N$, then there exists $p \in [0, 1]$ such that $pL + (1 - p)N \approx M$

Given these axioms (referred to as von Neumann-Morgenstern postulates) and a rational preference $\succeq$ specified over the set of lotteries with prices in X, utility theory shows that there always exists a unique value function u (up to a scaling change) that represents $\succeq$. This means that for every pair of possible outcomes x_i and x_j, $x_i \succeq x_j \Leftrightarrow u(x_i) \geq u(x_j)$ and that for every pair of lotteries $\langle p_1, x_1\,;\, p_2, x_2\,;\, \ldots\,;\, p_r, x_r \rangle$ and $\langle p'_1, x_1\,;\, p'_2, x_2\,;\, \ldots\,;\, p'_r, x_r \rangle$:

$$\langle p_1, x_1\,;\, p_2, x_2\,;\, \ldots\,;\, p_r, x_r \rangle \succeq \langle p'_1, x_1\,;\, p'_2, x_2\,;\, \ldots\,;\, p'_r, x_r \rangle \quad \text{if and only if}$$

$$\sum_{i=1}^{r} p_i u(x_i) \geq \sum_{i=1}^{r} p'_i u(x_i)$$

In summary, assuming given

- a decision table $T : D \times \Omega \rightarrow X$, where $D = \{d_1, \ldots, d_m\}$ is a set of possible decisions, $\Omega = \{w_1, \ldots, w_n\}$ is a universe and X is the set of possible outcomes,
- a value function $u : X \rightarrow \mathbb{R}$, and

- a probability distribution $P : \Omega \to [0, 1]$

the decisions of D are ranked and chosen as follows.

Definition 2.5. *Define for every decision $d_i \in D$ the expected utility of d_i as the value $E(d_i) = \sum_{j=1}^{n} P(w_j)u(x_{i,j})$. The **expected utility** criterion consists in choosing the decision with maximal expected utility.*

Remark that the expected utility criterion fulfils all the axioms from 2.1 to 2.6. Axiom 2.7 can only be adopted to study criteria which do no make use of probabilistic information, but this is not the case here. In Chapter 6, we will see that the criteria used for decision making with imprecise probabilities also collapse with the expected utility criterion when probabilities become precise, thus offering a unified approach to rational decision making.

2.1.4 Decision theory: summary

In summary, we have seen that decisions are modeled by a) a decision table describing the outcomes of a set of possible decisions under different scenarios, b) a value function measuring the desirability of the outcomes and optionally c) a probability distribution describing the chance of occurrence of the scenarios. The decision model must be provided by the decision maker and its construction is not the object of the theory. If no probability distribution is available (strict uncertainty) then the best decision can be chosen by applying criteria such as Wald's maximin return, Hurwicz's optimism-pessimism index, Savage's minimax regret or Laplace's principle of insufficient reason. Wald's criterion, by only looking at the worst-case scenario of a decision, is safe but very pessimistic. Hurwicz's criterion combines both the worst and best case scenario and thus allows a decision maker to adapt its level of optimism to the situation. Savage's criterion focuses on avoiding regrets that may result from making a

wrong decision. Finally, Laplace's criterion gives equal importance to all possible scenarios. Despite their rationality (see list of axioms fulfilled), the criteria may lead to different decisions. Finally, if a probability distribution is available, then the decision maker can choose the best decision by applying the expected utility criterion which weights the utility of every possible outcome of a decision against its likelihood.

2.2 Argumentation for decision making

2.2.1 What is argumentation ?

Uses of argumentation form a recent branch of logic and artificial intelligence. From a historical point of view [Chesñevar *et al.*, 2000], modeling arguments appears to be at the foundation of artificial intelligence's understanding of rule-based systems where rules can come into conflict. When a rule supporting a conclusion may be defeated by new information, it is said that such reasoning is defeasible. When we chain defeasible reasons to reach a conclusion, we have *arguments*, instead of proofs. Broadly speaking, arguments can therefore be interpreted as justified claims or objects that have the form of logical proof, but that do not have the force of logical proof [Krause *et al.*, 1995].

Nonetheless, the concepts of mathematical proof and argument can be defined in the same manner [Bench-Capon and Dunne, 2007], by defining means for expressing assertions, accepted bases on which to build theorems from axioms, procedures by which further theorems may be derived from existing theorems and axioms, and precise concepts of termination for derivation. The concept of argument generalises the one of proof, insofar as arguments can be derived from uncertain information, taking the form of defeasible rules or defeasible beliefs (referred to as assumptions in this chapter). This characteristic is what gives

them their defeasible character. Since assumptions are adopted subjectively, arguments can also be said to be subjective. Therefore, arguments are the basis of opinions and it is necessary to keep in mind that their primary role is to persuade, not to find universal truth.

Research in argumentation covers a rather wide collection of topics. The earliest studies of argumentation are of philosophical nature and concerned with the nature, structure and representation of arguments [Wigmore, 1940; Toulmin, 1958]. In the early 90s, argumentation started to be used as a way of representing uncertain knowledge [Krause and Clark, 1993], to model common sense [McCarthy, 1990] and default reasoning [Pollock, 1992] as well as reasoning under uncertainty [Fox *et al.*, 1993]. Abstract argumentation [Dung, 1995] then offered a unified theory of non-monotonic reasoning and logic programming.

Abstract argumentation has since then favoured the development of logical approaches to epistemic reasoning [Bondarenko *et al.*, 1997; Prakken and Sartor, 1997; Besnard and Hunter, 2000; Dung *et al.*, 2002; 2006; Toni, 2007; Besnard and Hunter, 2008; Toni, 2008]. A parallel evolution has been the development of logics of argumentation for reasoning under uncertainty [Krause *et al.*, 1995; Bonet and Geffner, 1996; Parsons, 1997; Fox, 2003] which together with abstract argumentation have strongly stimulated the emergence of argumentation-based methods for practical reasoning [Amgoud and Prade, 2004; Amgoud, 2005; Amgoud and Prade, 2006; Prakken, 2006].

Recent results in the field of argumentation are refinements concerning the notion of argument acceptability [Amgoud and Cayrol, 1998; 2002; Caminada, 2006; Dung *et al.*, 2007] and their computation [Vreeswijk and Prakken, 2000; Dung *et al.*, 2007] or complexity [Dimopoulos *et al.*, 2002; Dunne, 2007; 2008; 2009], studies on the persuasiveness and explanatory power of arguments [Bench-Capon and Prakken, 2006; Morge and Mancarella, 2007; Amgoud and de Saint-

Cyr, 2008], the design of argument strength measures [Jakobovits and Vermeir, 1999; Besnard and Hunter, 2000; Fox, 2003; Cayrol and Lagasquie-Schiex, 2005], the analysis of strategies of argumentation [Dunne and Bench-Capon, 2003; Rahwan and Larson, 2008] , the design and study of expert systems and decision support systems [Fox, 1993; Fox *et al.*, 1998; Jefferys *et al.*, 2006; Mozina *et al.*, 2007] and the automated interpretation and generation of human-like (enthymematic) arguments [Zuckerman, 2009].

Amongst these topics, some are of particular relevance for analysing problems of decision making under uncertainty. These are epistemic reasoning, reasoning under uncertainty, practical reasoning, the acceptability of arguments, the strength of arguments and the design and study of expert systems and decision support systems. Our objective in the next subsections is to give an overview of argumentation for decision making. This will allow us to illustrate the differences that exist between argumentative decision making and decision theory. We shall recall the reader that our purpose in this thesis is not to provide an argumentation-based alternative to decision theory, but rather to develop argumentation-based tools complementing decision theory. Thus, the next subsections will in fact allow us to introduce basic ideas and material for developing new decision-theoretic tools using argumentation.

2.2.2 General decision making process

A general approach for argumentative decision making has been recently outlined in [Amgoud, 2005] and [Amgoud, 2009]. Assuming given a set of possible decisions $D = \{d_1, \ldots, d_m\}$, where $m \geq 1$ and a set or language $\mathcal{L}$ of propositions relevant to the decision making situation and notably including D, let us call *belief* any proposition in $\mathcal{L}$ which is not a possible decision. The decision making process is decomposed into five successive steps:

1. constructing arguments in favour or against decisions or beliefs
2. evaluating the intrinsic strength of each argument
3. determining the conflicts among arguments
4. evaluating the acceptability of arguments
5. comparing decisions on the basis of the accepted arguments

Arguments in favour or against a decision are called *practical arguments.* Practical arguments essentially serve to highlight the advantages and disadvantages of decisions under various circumstances. Such arguments are therefore built both from the preferences of the decision maker (what is good or bad) and his beliefs. Arguments in favour or against beliefs are instead called *epistemic arguments.* Their role is to justify or invalidate beliefs, or, in other words, to mitigate uncertainty.

Epistemic arguments are uncertain claims according to which some beliefs are or are not true. Thus, it is clear that the set of epistemic arguments implicitly refers to a universe. As we will see in the next subsection, the intrinsic strength of epistemic arguments plays essentially the same role as the probability distribution P in a decision-theoretic model. Practical arguments contain information concerning the (desirability of the) possible outcomes of decisions. It is as if these practical arguments were built from an implicit value function u. The fourth step, which consists in selecting a subset of "accepted" arguments has no equivalent in decision theory. The selection of acceptable arguments allows to simplify the decision making, by ignoring part of the available information and basing the decision comparison on the information that is the most credible. In decision theory, this would be tantamount to selecting the most likely scenario and ignoring the others, reducing then the problem to a very simple situation of decision making under certainty. Decisions can then

be compared on the basis of the accepted practical arguments only via *decision principles.* Decision principles play the same role as decision criteria in decision theory, viz., the one of establishing a ranking of the decisions. We ignore the construction of arguments at the first step but describe in detail the other steps in the next subsections.

2.2.3 Intrinsic strength of arguments

The second step of the argumentative decision making process consists in establishing and specifying preference orders between the arguments. Three pre-order relations [2] are required. These are noted $\geq_e$ to compare epistemic arguments, $\geq_p$ to compare practical arguments and $\geq_m$ to compare epistemic arguments with practical arguments. For any arguments x and y, the relation $x \geq y$ (for any of these orders) can be read as "x is stronger than y". The simplest of these pre-orders is $\geq_m$, as one considers in argumentation for decision making that epistemic arguments are stronger than practical arguments. We now explain in turn the pre-orders used for comparing epistemic arguments and practical arguments respectively.

One writes $x \geq_e y$ when x and y are both epistemic arguments and x is "stronger" than y. One may compare the strength of arguments either subjectively, as in the view of [Perelman and Olbrechts-Tyteca, 1969] and the work of [Bench-Capon, 2003], where the strength of an argument depends on the social or moral values that it advances – or objectively, by treating the strength of an epistemic argument as its probability of holding for true. Some authors such as [Zuckerman, 2009] use nowadays Bayesian networks to compute these probabilities. Other probabilistic models are of course possible. A simple approach employed by[Poole, 1993; Krause *et al.*, 1995] to evaluate the probability of an epistemic argument consists in decomposing the underlying assumptions of the

[2] A pre-order is a binary relation that is reflexive and transitive.

argument into statistically independent propositions and multiply their individual probabilities. In probability theory, to say that two events are independent intuitively means that the occurrence of one event makes it neither more nor less probable that the other occurs. For instance, the existence of cats, dogs or other quadrupled life-forms on Mars are clearly dependent. Thus, the existence of life on Mars cannot simply be refuted by enumerating a sufficiently large number of known quadrupled life-forms. It is also possible to use possibility and necessity measures [Amgoud and Prade, 2004] instead of probability distributions to assess the strength of an epistemic argument. A very general presentation of the mathematical properties and examples of intrinsic strength measures defined over numerical or symbolic scales of confidence can be found in [Krause *et al.*, 1995] and [Ambler, 1996]. Concrete examples of measures over the interval $[0, 1]$ are: the weakest link measure, whereby the confidence of an argument is taken as the minimum of the confidences of its assumptions; the product rule whereby the confidence of an argument is the product of the confidence of its assumptions; or a probabilistic valuation (sum-product rule), where the confidence in a proposition is the sum of the confidence of the arguments supporting it and the confidence of each argument is the product of the probability of its assumptions.

Informally speaking, *practical arguments* denote justified claims according to which a decision is good or bad. The justifications can be that the decision allows or fails to achieve a desirable or undesirable outcome. Justifications are build from facts and assumptions. Practical arguments in favour of a decision are called arguments *pro* and those against a decision arguments *con*. One writes $x \geq_p y$ when x and y are both practical arguments and x is "stronger" than y. The strength of a practical argument is a function of two variables:

- the level of certainty C of the argument, which depends on the overall

confidence the decision maker has in the assumptions of the argument,

- the importance I of achieving or avoiding that outcome.

The certainty level can be assessed in the same way as the intrinsic strength of epistemic arguments. The importance of the outcome is represented by a number. The strength of the practical argument is taken by [Amgoud, 2009] as equal to the minimum of the two values C and I. It goes without saying that for the comparisons to be meaningful, the importance and certainty levels must all be expressed on a common scale and in the same unit. If probability is chosen to express the certainty level of practical arguments, then levels of importance must also be specified by values ranging between 0 and 1. The idea of using the importance of achieving an outcome to appreciate its strength is originally due to [Perelman and Olbrechts-Tyteca, 1969], who rather than using the term importance speaks of *value* advanced or promoted by an argument.

2.2.4 Conflicts and notions of argument acceptability

When an epistemic argument x is expressed against a belief p and that p constitutes one of the underlying assumptions of the epistemic or practical argument y, we say that "x attacks y". In a similar fashion, if x is an argument in favour of p and the negation $\neg p$ of p is an assumption of y, then we say that "x attacks y". The set Arg of all epistemic and practical arguments together with the set att of all pairs of arguments x and y such that x attacks y can be seen as a directed graph (Arg, att) in which 1) nodes are arguments and 2) arcs represent attacks between the arguments. Such a graph displaying arguments and their conflicts is called in the field of argumentation an *abstract argumentation framework* [Dung, 1995].

Definition 2.6. *An **abstract argumentation framework** is a pair $F = (Arg, att)$, where Arg is a set of arguments and att is a binary relation on Arg.*

Abstract argumentation overlooks the internal structure of arguments, hence the use of the term *abstract*. In decision making, the graph of attacks among arguments is transformed into a defeat graph which merges attacks and information concerning the intrinsic strength of the arguments. Note that abstract argumentation itself prevents the derivation of an intrinsic measure of strength for arguments, as having access to the internal structure of an argument is essential to assess the probability of an argument holding for true. The attack binary relation and the preference relation $\geq$ determined by the strength given to practical and epistemic arguments are merged into a new binary relation between arguments called the *defeat* relation.

Definition 2.7. *For every pair of arguments x and y, we say that x **defeats** y if and only if x attacks y and it is not the case that $y > x$, i.e. y is not strictly preferred to x* [3].

This notion of defeat extends the notion of defeat first proposed in value-based argumentation [Bench-Capon and Dunne, 2002; Bench-Capon, 2003], whereby the preference of a practical argument over another only depends on the values respectively promoted by these arguments and the preferences existing over these values (but does not depend on their intrinsic strengths).

Various acceptability semantics can be given, allowing to identify which arguments to select to make a decision. These semantics were originally defined based on the attack relationship, as follows. For every arguments $a, b \in Arg$ and sets of arguments $A, B \subseteq Arg$, we write

- a attacks b iff $(a, b) \in att$
- A attacks b iff there exists $a \in A$ such that a attacks b
- a attacks B iff there exists $b \in B$ such that a attacks b

[3] By definition, $y > x$ when $y \geq x$ and $x \not\geq y$.

- A attacks B iff there exist $a \in A$ and $b \in B$ such that a attacks b

Definition 2.8 (acceptability semantics). *For any $A \subseteq Arg$, A is*

- ***conflict-free*** *iff it is not the case that A attacks itself*
- ***naive*** *iff A is maximally conflict-free with respect to set inclusion*
- ***admissible*** *iff A is conflict-free and attacks all arguments that attack it*
- ***preferred*** *iff A is maximally admissible with respect to set inclusion*
- ***complete*** *iff A is admissible and contains every argument x such that A attacks all arguments attacking x*
- ***semi-stable*** *iff A is complete where $A \cup \{x | A \text{ attacks } x\}$ is maximal with respect to set inclusion*
- ***stable*** *iff A is conflict-free and attacks every argument it does not contain*
- ***grounded*** *iff A is minimally complete with respect to set inclusion*
- ***ideal*** *iff A is a maximally admissible set of arguments that is contained in every preferred set of arguments*

The conflict-free, admissible, preferred, stable, complete and grounded semantics have been introduced by [Dung, 1995], the naive semantics by [Bondarenko *et al.*, 1997], the semi-stable semantics by [Caminada, 2006] and the ideal semantics by [Dung *et al.*, 2007]. These semantics have the following properties and relationships [Dung, 1995; Bondarenko *et al.*, 1997; Caminada, 2006; Dung *et al.*, 2007]. Here, *extensions* simply correspond to sets of arguments.

Property 2.1. *There always exists conflict-free, naive, admissible, preferred, semi-stable, complete, grounded and ideal extensions.*

Property 2.2. *The grounded and ideal extensions are unique.*

Property 2.3. *Every conflict-free extension is contained in a naive extension.*

Property 2.4. *Every admissible extension is contained in a preferred extension.*

Property 2.5. *Every stable extension is a preferred extension.*

Property 2.6. *Every preferred extension is a complete extension.*

Property 2.7. *Every stable extension is a semi-stable extension.*

Property 2.8. *The grounded extension is contained in the ideal extension.*

In argumentation for decision making, the attack relation is replaced by the defeat relation in the definition of these semantics. For the decision making process to take place, a notion of acceptability is chosen. We refer to the set of arguments deemed acceptable by the chosen semantics as *accepted* arguments.

2.2.5 Decision principles

After the selection of acceptable arguments comes the final step of comparing the possible decisions. Decision are compared according to decision principles which have the same role as decision criteria in decision theory. Principles can be either *unipolar*, *bipolar* or *non-polar* [Amgoud and Prade, 2008].

Unipolar principles are the simplest ones. They are called *unipolar* because they only take into account one type of practical argument, which can be either the arguments in favour of decisions (pro) or against decisions (con). A typical example of unipolar principle is the *counting principle.*

Definition 2.9. *d_1 is **preferred** to d_2 ($d_1 \succeq d_2$) if and only if*

- *the number of accepted arguments pro d_1 is greater than the number of accepted arguments pro d_2, or alternatively*
- *the number of accepted arguments con d_1 is smaller than the number of accepted arguments con d_2.*

This preference relation naturally defines a total pre-order on the set of decisions. Another way of comparing decisions is to compare the intrinsic strength of their arguments pro (or con). So, we may use a different unipolar principle, referred to as *focus principle.*

Definition 2.10. *d_1 is **preferred** to d_2 ($d_1 \succeq d_2$) if and only if*

- *there exists an accepted argument pro d_1 whose intrinsic strength is greater than the intrinsic strength of any accepted argument pro d_2, or alternatively*
- *the intrinsic strength of any accepted argument con d_1 is smaller than the intrinsic strength of one of the accepted arguments con d_2.*

Bipolar principles take into account both accepted arguments pro and con. These consequently generalise the unipolar principles. A possible generalisation of the counting principle is the *bipolar counting principle.*

Definition 2.11. *d_1 is **preferred** to d_2 ($d_1 \succeq d_2$) if and only if the number of accepted arguments pro d_1 is greater than the number of accepted arguments pro d_2 and the number of accepted arguments con d_1 is smaller than the number of accepted arguments con d_2.*

The focus principle is generalised by the *bipolar focus principle.*

Definition 2.12. *d_1 is **preferred** to d_2 ($d_1 \succeq d_2$) if and only if there exists an accepted argument pro d_1 whose intrinsic strength is greater than the intrinsic strength of any accepted argument pro d_2 and the intrinsic strength of every accepted argument con d_1 is smaller than the intrinsic strength of one of the accepted arguments con d_2.*

Note however that neither of these two generalisations leads to total preorders. Sometimes, decisions become incomparable under such principles. This

problem also exists in decision theory (axiom 2.1 concerning the complete ranking of decisions implied by any "rational" decision criterion is not universally accepted).

In order to obtain a complete pre-order on the set of possible decisions, one needs to aggregate all arguments pro and con a decision and compute a single value. Decision principles that rely upon such an aggregation criterion are called non-polar. For instance, one may count the total number of accepted arguments pro a decision and subtract from it the total number of accepted arguments con. The weakness of this simplistic approach is that it does not take into account the strength of the practical arguments. Thus, a more satisfying approach would consist in adding the strength of all accepted arguments pro a decision and subtracting from it the strength of all accepted arguments con. Such a non-polar principle obviously leads to a total pre-order on the entire set of possible decisions.

2.2.6 Argumentation for decision making: summary

We have defined argumentation as the study of defeasible proofs and have discussed the origins and objectives of this branch of artificial intelligence. We have then outlined a methodology for decision making based on the construction and use of arguments in favour or against decisions (practical arguments) or beliefs (epistemic arguments). In this approach, each argument is first examined individually and its intrinsic strength is assessed, by examining either its probability or the values it promotes. The existing approaches for estimating the intrinsic strength of epistemic arguments have been briefly introduced. The strength of a practical argument in favour or against a decision obviously depends on the level of confidence (intrinsic strength) the decision makers has in its truth. It nevertheless also depends on the importance of the specific benefits or disad-

vantages shown by the argument. Practical arguments are based on beliefs, some of which may be uncertain, because the decision maker has found reasons (epistemic arguments) to question them. At the dialectical level, practical and epistemic arguments conflict. Such conflicts can be formally modeled by a directed graph of attacks. This graph of attacks is transformed into another graph of defeats between the arguments. The conflicts can be resolved by employing notions of argument acceptability on the defeat graph. Decision criteria such as the bipolar counting or focus principle finally easily allow to rank decisions by comparing the accepted practical arguments in favour or against them.

2.3 Summary

We have defined the concepts of decision table, value function and Bayesian probability distribution and have introduced the main criteria for making decisions in decision theory. We have introduced argumentation in the context of artificial intelligence and have presented argumentation for decision making in detail. In particular, we have defined the notion of abstract argumentation. Abstract argumentation will be used in Chapters 4 and 5 to construct value functions and estimate probability distributions. In Chapter 3, a special instance of abstract argumentation called *assumption-based argumentation* will be used to construct decision tables. In this thesis, argumentation will be used as a practical foundation for decision theory, unlike argumentation for decision making, which constitutes an alternative to decision theory.

Chapter 3

Decision tables & assumption-based argumentation

As discussed in Chapter 2, decision theory is concerned with decision making under strict uncertainty and under risk. Decision tables are needed to support both types of decision making. This chapter is concerned with the practical construction of decision tables. Some researchers in the field of artificial intelligence [Vanthienen and Dries, 1993] have observed that very little had actually been done by 1993 within decision table theory [Schmidt and Kavanagh, 1964; King, 1968; Milner, 1977; Reilly *et al.*, 1987] to support and automate the construction of decision tables. This is still true today.

The construction of a decision table is not straightforward [Verhelst, 1980], for the following reasons. First, the decision maker must find a way of representing the decisions, scenarios and outcomes. Second, the decision maker must use

his knowledge and understanding of the domain to link decisions and scenarios (causes) to outcomes (effects). Finally, he shall process all relevant data and arrive at a complete, consistent and comprehensible identification of the possible decisions and scenarios and precisely determine the outcome of each decision under every scenario. Good decisions cannot be made without good decision tables: the decision maker should not overlook any possible scenario or decision and must carefully examine each outcome before comparing their likelihood or desirability and committing to any decision.

We propose *basic influence diagrams* as a tool for describing situations of decision making under strict uncertainty. The objective of this chapter is – given a basic influence diagram as input – to propose an automated and rational way to process the diagram and construct a decision table. We propose to use a particular model of argumentation called assumption-based argumentation [Bondarenko *et al.*, 1997; Dung *et al.*, 2006; Toni, 2007] and the new *liberal stable* semantics to obtain decision tables from basic influence diagrams. We develop an algorithm derived from hypergraph theory [Eiter and Gottlob, 1995; Boros *et al.*, 2000; Eiter and Gottlob, 2002; Khachiyan *et al.*, 2005] to compute the liberal stable semantics and construct decision tables. As an illustration and application, we employ this method to construct a decision table for the game of poker [Mendelson, 2005].

The chapter is organised as follows. Section 3.1 introduces basic influence diagrams. Section 3.2 recalls background definitions for assumption-based argumentation and introduces an extension of this form of argumentation (needed for the purposes of this chapter). Section 3.3 shows how to transform basic influence diagrams into extended assumption-based argumentation frameworks. Then Section 3.4 introduces the liberal stable semantics and shows that under some mild assumptions this semantics allows to solve basic influence diagrams.

Sections 3.5 presents a hypergraph algorithm for computing the liberal stable semantics. Section 3.6 shows how to obtain a decision table once the liberal stable semantics has been computed. Section 3.7 applies the method to the game of poker. We discuss related work in Section 3.8. Section 3.9 summarises and concludes the chapter.

The results exposed in Sections 3.1 (basic influence diagrams), 3.2 (extended assumption-based argumentation), 3.3 (link between basic influence diagrams and extended assumption-based argumentation frameworks) and 3.4 (liberal stable semantics) have been published in [Matt and Toni, 2008a] at the Second International Conference on Computational Models of Argument.

3.1 Basic influence diagrams

As presented in Chapter 2, a decision table is a graphical representation of a mapping of the form $T : D \times \Omega \rightarrow X$ where $D = \{d_1, \ldots, d_m\}$ is a set of possible decisions to choose from, $\Omega = \{w_1, \ldots, w_n\}$ is a set of possible states of the world or scenarios and X is a set of outcomes. For every pair of decision d_i and scenario w_j, $T(d_i, w_j) = x_{i,j}$ represents the consequence of d_i under w_j. In decision under strict uncertainty, the decision maker does not know which scenario w_j will occur or is the true one. The set of scenarios Ω thus represents the uncertainty of the decision maker. Then, given a value function $U : X \rightarrow \mathbb{R}$ quantifying the desirability $U(x)$ of every outcome $x \in X$, the decision maker may choose a decision by applying one of the criteria listed in Chapter 2 (Wald's maximin return, Hurwicz's optimism-pessimism index, Savage's minimax regret, Laplace's principle of insufficient reason). In decision under risk, it is also assumed that the decision maker knows the probability distribution $P : \Omega \rightarrow [0, 1]$ of the scenarios. The best decision is then given by the expected utility criterion, also given in Chapter 2. We are not concerned

here with decision making (i.e. the adoption of decision criteria and computation of optimal decisions), but restrict our attention to the construction of T. The problem considered is to identify the set of possible decisions D, the set of scenarios Ω, the set of outcomes X and a mapping T that is consistent with the decision maker's knowledge of the decision domain.

We use *basic influence diagrams* to model the decision maker's objectives, decisions, beliefs, uncertainties, the causal dependencies between decisions and beliefs, and the possible conflicts between beliefs or decisions. Formally,

Definition 3.1. *A **basic influence diagram** is composed of two parts: a **diagram** and a set of **dependency rules**. The diagram is an annotated finite directed acyclic graph whose nodes are labelled by literals of the form p or $\neg p$ from a given language $\mathcal{L}$. Every node belongs to one of the following categories:*

- ***Goals** are the nodes that have no outgoing arcs. Goals may be positive or negative, graphically distinguished with a + and - symbol respectively.*
- ***Atomic decisions** form an arbitrary subset of the nodes that are not goals and that have no incoming arcs. Atomic decisions are graphically distinguished from the other nodes by a square box.*
- ***Beliefs** are the nodes which are neither goals nor atomic decisions. Beliefs that have no incoming arcs are called **fundamental beliefs**. The fundamental beliefs that are known to be true by the decision maker are called **facts** and are underlined. The other fundamental beliefs are called **unobservable fundamental beliefs** and are annotated with a ? symbol.*

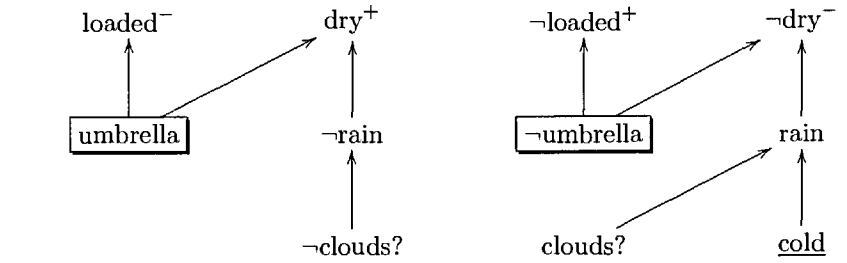
***Dependency rules** may be of the form 'if p_1 and ... and p_k then q' or 'q', where p_1,..., p_k and q are literals in $\mathcal{L}$. The arcs of the influence diagram and the dependency rules must conform in the following way. For each rule of the form 'if p_1 and ... and p_k then q' there must exist for every $i \in \{1, ..., k\}$*

an arc from p_i to q. For each dependency rule 'q', q must appear as a fact in the diagram (and be underlined). Conversely, for each arc from p to q in the diagram, there must exist at least one dependency rule 'if ... p ... then q' and for each fact q there must exist a dependency rule 'q'.

$\mathcal{L}$ is meant to be the set of all literals occurring in the diagram. Thus, $\mathcal{L}$ may not be closed under $\neg$ (negation). Goals serve to model the objectives of the decision maker and correspond to properties the decision maker desires to achieve (positive goals) or to avoid (negative goals). Unobservable fundamental beliefs represent the uncertainties of the decision maker. An arc from a node labelled by p to a node labelled by q means that the truth of q logically depends on the truth of p in the decision domain. All logical dependencies are expressed by dependency rules. These rules are strict, which mean that their conclusion is certain, as opposed to defeasible rules whose conclusion is only plausible. The use of defeasible rules amounts to filtering out abnormal or unlikely scenarios and thus prevents the construction of an exhaustive set of scenarios Ω. Finally, remark that according to their definition, basic influence diagrams may not be bipartite graphs [1].

Let us now see how to concretely represent a decision making problem using a basic influence diagram. The basic influence diagram shown in Figure 3.1 models a simple decision problem borrowed and adapted from [Amgoud and Prade, 2004] in which one has to decide whether to take an umbrella or not when going out for a walk. The decision maker has some elementary knowledge about the weather. He knows as a fact that the weather is cold and he believes that if the weather is cold and cloudy, then it will rain. Taking an umbrella guarantees him to be able to stay dry but has the disadvantage of making him

[1] A bipartite graph or bigraph is a set of graph vertices decomposed into two disjoint sets such that no two graph vertices within the same set are adjacent. For example, a basic influence diagram based on the following dependency rules is not bipartite: if a then b, if b then c, if a then c.



if umbrella then loaded
if umbrella then dry
if ¬rain then dry
if ¬umbrella then ¬loaded
if ¬umbrella and rain then ¬dry
if ¬clouds then ¬rain
if clouds and cold then rain
cold

Figure 3.1: Basic influence diagram for the umbrella problem.

loaded.

Our diagrams are similar to the type of diagrams used in the argumentation-based decision support system developed by [Morge and Mancarella, 2007]. Their diagrams also correspond to nodes connected by directed arcs modeling cause-effect dependencies and nodes correspond either to decisions, knowledge (equivalent to our beliefs) or value nodes (equivalent to our goals). However, our diagrams also include dependency rules, making explicit the dependencies between nodes. This is a specificity of basic influence diagrams. Without the dependency rules, the arcs of a diagram would be ambiguous. For instance, the arc cold → rain could be misunderstood as 'if cold then rain'. Note that even though the meaning of dependency rules is more precise than the one of arcs linking nodes in the diagram, the dependency rules on their own do not suffice to represent the decision problem and build a decision table, as for instance the diagram allows to identify the desirable (+) and undesirable (-) goals of the decision maker. This distinction is actually unnecessary for the construction

of a decision table. It is nevertheless quite useful to make a decision and in particular to build a value function, as we shall see in the next chapter. In the remainder of this chapter, we ignore the + and - symbols for goals. The diagram also allows us to distinguish decisions from unobservable fundamental beliefs. Finally, the diagram constitutes a sometimes precious graphical aid to the specification of dependency rules in complex decision domains.

Our final remark on basic influence diagrams concerns the use of the negation symbol $\neg$. Here $\neg$ corresponds to classical negation. Thus, whenever a node in the diagram is labelled by a literal $\neg p$, it will be assumed that in every possible outcome either p or $\neg p$ holds, and that p and $\neg p$ cannot hold simultaneously. Therefore, in a basic influence diagram, strong logical dependencies which we may see as *conflicts* exist between nodes of the form p and $\neg p$. However, we do not force $\mathcal{L}$ to be closed under $\neg$. If for some $p \in \mathcal{L}$, $\neg p$ does not occur in the diagram, then $\neg p$ does not belong to $\mathcal{L}$. Finally, for every $\neg p \in \mathcal{L}$, we will assume in the remainder of this chapter that $\neg\neg p = p$.

Given a basic influence diagram, we say that

Definition 3.2. *A **entails** q and write $A \vdash q$ if and only if A is a set of literals of the diagram composed only of atomic decisions and unobservable fundamental beliefs and there exists a sequence of literals $p_1, \ldots, p_n$ where $p_n = q$ such that for every $i \in \{1, \ldots, n\}$, either p_i is an element of A or p_i is the conclusion of a dependency rule whose premises all belong to $\{p_1, \ldots, p_{i-1}\}$.*

For example, in the umbrella problem, we would have $\{\text{umbrella}\} \vdash \text{loaded}$, $\{\neg\text{umbrella}, \text{clouds}\} \vdash \neg\text{dry}$ and $\{\neg\text{umbrella}, \text{clouds}\} \vdash \neg\text{loaded}$.

When reasoning with negation, rationality is based on two requirements, viz. consistency (or absence of conflicts) and closure under inference [Caminada and Amgoud, 2005; Toni, 2007]. These two properties are widely recognised as essential and we believe that a rational decision maker should conform to them.

When using a basic influence diagram and classical negation, it is important to strengthen this basic notion of rationality by adding a third requirement. Namely, in order to reflect the interpretation of $\neg$ as classical negation, for every $p, \neg p \in \mathcal{L}$, one must enforce the decision maker to choose to believe in either p or $\neg p$. We call this property *decidedness*. Some authors [Takahashi and Sawamura, 2004] insist on the philosophical importance of allowing descriptions of states of the world whereby both p and $\neg p$ hold, or even where neither of them holds. Such descriptions correspond for instance to logical formalizations of dialogues between parties sharing different views, but they are not meant to be taken as fully rational opinions of individuals as intended in this thesis. Therefore, we define

Definition 3.3. *An **outcome** is a subset of $\mathcal{L}$. An outcome O is **rational** if and only if*

- $\forall (p, \neg p) \in \mathcal{L}^2$*: it is not the case that* $p \in O$ *and* $\neg p \in O$ *(consistency)*
- $\forall (p, \neg p) \in \mathcal{L}^2$*:* $p \in O$ *or* $\neg p \in O$ *(decidedness)*
- *there exists a set of atomic decisions and unobservable fundamental beliefs* A *such that* $O = \{p \in \mathcal{L} \mid A \vdash p\}$ *(closure under dependency rules)*

We call *rational opinion* a set of atomic decisions and unobservable fundamental beliefs that entails a rational outcome.

Definition 3.4. *An **opinion** is a set of atomic decisions and unobservable fundamental beliefs. An opinion A is **rational** if and only if the set $O(A) = \{p \in \mathcal{L} \mid A \vdash p\}$ is a rational outcome.*

In order to construct a decision table from a basic influence diagram, we must find all the rational outcomes. The search for rational outcomes boils down to the search for rational opinions. This can be done by assumption-based argumentation, as discussed next.

3.2 Extended assumption-based argumentation

This section introduces a simple extension of assumption-based argumentation [Bondarenko *et al.*, 1997; Dung *et al.*, 2006; Toni, 2007]. But first, let us start with basic definitions from standard assumption-based argumentation.

3.2.1 Assumption-based argumentation

In [Dung *et al.*, 2006], a (flat) assumption-based argumentation framework is a tuple $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ where

- $\mathcal{L}$ is a language composed of a finite set of sentences.
- $\mathcal{R}$ is a set of inference rules of the form

$$\frac{p_1, \ldots, p_k}{q}$$

  where $p_1, \ldots, p_k, q \in \mathcal{L}$ and $k \geq 0$. The sentences $p_1, \ldots, p_k$ are called the *premises* of the rule and q the *conclusion* of the rule. The rule can be read 'if $p_1, \ldots, p_{k-1}$ and p_k hold, then q holds'. When $k = 0$, the rule represents the fact that q holds.

- $\mathcal{A} \subseteq \mathcal{L}$ is a set of assumptions. $\mathcal{A}$ must be such that no assumption appears as the conclusion of an inference rule.
- $\mathcal{C} : \mathcal{A} \mapsto \mathcal{L}$ is a total mapping associating to every assumption $x \in \mathcal{A}$ a contrary $\mathcal{C}(x) \in \mathcal{L}$.

We will assume in this thesis that all sentences in $\mathcal{L}$ are literals, although this is not necessarily the case in general assumption-based argumentation.

Definition 3.5. *A **tight argument** based on a set of assumptions A with conclusion q is denoted $A \vdash_t q$ and corresponds to a sequence of sets $S_1, \ldots, S_m \subseteq \mathcal{L}$,*

where $m \geq 1$, $S_1 = \{q\}$, $S_m = A$ *and for every* $i \in \{1, \ldots, m-1\}$, *there exists* $\sigma \in S_i$ *such that* $S_{i+1} = S_i - \{\sigma\} \cup S$ *for some inference rule of the form* $\frac{S}{\sigma}$.

In other words, a tight argument $A \vdash_t q$ is a backward deduction of q from A in the deductive system formed by the pair $(\mathcal{L}, \mathcal{R})$. Given any two sets of assumptions $A, B \subseteq \mathcal{A}$, we say that

Definition 3.6. *A* ***attacks*** *B if and only if $A' \vdash_t q$ and $q = C(b)$ for some $A' \subseteq A$ and $b \in B$.*

In other words, A attacks B if and only if there exists a tight argument based on A whose conclusion is the contrary of an assumption contained in B.

Given this notion of attack, various notions of acceptability can be defined. These notions capture different epistemic notions of rationality, as they can be used as criteria for accepting or rejecting sets of assumptions (referred to as *extensions*) in the same way the corresponding notions of acceptability from abstract argumentation (see definitions in background Section 2.2.4) can be used to accept or reject sets of arguments.

Definition 3.7. *A set of assumptions $A \subseteq \mathcal{A}$ is*

- ***conflict-free*** *iff A does not attack itself*
- ***naive*** *iff A is maximally conflict-free with respect to set inclusion*
- ***admissible*** *iff A is conflict-free and A attacks every set B of assumptions that attacks A*
- ***preferred*** *iff A is maximally admissible with respect to set inclusion*
- ***complete*** *iff A is admissible and contains all assumptions x such that A attacks all the sets of assumptions that attack $\{x\}$*
- ***stable*** *iff A is conflict-free and attacks every $\{x\}$ such that $x \in \mathcal{A} - A$*

- ***semi-stable*** *iff* A *is complete where* $\{A\} \cup \{B \subseteq \mathcal{A} | A \text{ attacks } B\}$ *is maximal with respect to the inclusion of sets of sets*
- ***grounded*** *iff* A *is minimally complete with respect to set inclusion*
- ***ideal*** *iff* A *is a maximally admissible set of assumptions with respect to set inclusion that is included in every preferred set of assumptions.*

The conflict-free, naive, admissible, preferred, complete, stable and grounded semantics for assumption-based argumentation have been given by [Bondarenko *et al.*, 1997]. The definition of the semi-stable semantics for assumption-based argumentation is here adapted from [Caminada, 2006]. The ideal semantics has been introduced by [Dung *et al.*, 2007].

3.2.2 Extended assumption-based argumentation

Extended assumption-based argumentation differs from standard assumption-based argumentation in the notion of contrary $\mathcal{C}$, the notion of attack between sets of assumptions, and the notions of stable and complete extensions.

Definition 3.8. *An* ***extended assumption-based argumentation framework*** *is a tuple* $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$*, where*

- $\mathcal{L}, \mathcal{R}$ *and* $\mathcal{A}$ *are defined as in standard assumption-based argumentation,*
- $\mathcal{C}$ *is a binary relation between sets of assumptions and sentences in* $\mathcal{L}$*:* $\mathcal{C} \subseteq 2^{\mathcal{A}} \times \mathcal{L}$.

For every pair $(A, p) \in \mathcal{C}$, p is a *contrary* of the set A of assumptions. However, p is not the contrary of any single assumption of A (unless A is a singleton of course). A very similar idea can be found in the work of [Nielsen and Parsons, 2006] who, in the context of abstract argumentation, consider

attacks from a set of arguments against an argument. In these joint attacks, no single argument in the attacking set forms an attack on its own.

The contrary relation $\mathcal{C}$ is by definition a subset of $2^{\mathcal{A}} \times \mathcal{L}$ and can thus be represented by a set (unordered list) of pairs of the form (P, q), where $P \subseteq \mathcal{A}$ and $q \in \mathcal{L}$. The maximum number of such pairs is $2^{|\mathcal{A}|}.|\mathcal{L}|$. In the next section, we will consider a special class of extended assumption-based frameworks, namely, those derived from basic influence diagrams. In these frameworks, the number of pairs contained in the contrary relation is in fact always bounded up by $N_{arg}.|\mathcal{L}|$, where N_{arg} is defined as the maximum number of tight arguments that support a sentence in $\mathcal{L}$.

We do not require that every set of assumptions admits a contrary, nor in particular that every singleton set $\{x\}$ admits a contrary in $\mathcal{L}$. The standard contrary relation is a special case of the extended contrary relation for which every pair $(A, p) \in \mathcal{C}$ is such that the set of assumptions A is a singleton and for every $x \in \mathcal{A}$, there exists a pair $(\{x\}, p) \in \mathcal{C}$.

For any sets of assumptions $A, B \subseteq \mathcal{A}$, we say that

Definition 3.9. *A **attacks** B if and only if there exists a pair $(P, q) \in \mathcal{C}$ such that $A' \vdash_t q$ and $P \subseteq B$ for some $A' \subseteq A$.*

In other words, we say that A attacks B if and only if there exists an argument supported by A whose conclusion is a contrary of some of the assumptions contained in B. The standard notion of attack is a special case of the extended relation of attack in which P must be a singleton.

As already said, all the semantics in extended assumption-based argumentation are defined as in standard assumption-based argumentation, except for the notions of complete and stable extensions. These modifications are necessary since sets of assumptions of the form $\{x\}$ (singletons) do not always have a contrary sentence in extended frameworks and thus may not be attacked.

Definition 3.10. *A set of assumptions A is said to be*

- ***complete*** *iff A is admissible and includes every set B such that A attacks all sets attacking B,*
- ***stable*** *iff A is conflict-free and attacks every set it does not include.*

The new definitions generalise the standard ones. Indeed, any standard framework can be seen as an extended one and, for standard frameworks, the standard notion of attack is equivalent to the new one. Then:

Property 3.1. *Given a standard assumption-based argumentation framework and a set A of assumptions*

1. *A is stable in the standard sense iff A is stable in the new sense, and*
2. *A is complete in the standard sense iff A is complete in the new sense.*

Proof. 1 $\Rightarrow$) If A is stable in the standard sense, then A is conflict-free. If A does not include some set B, i.e. $A \not\supseteq B$, then $\exists b \in B - A$. $b \notin A$, A attacks $\{b\}$ and therefore A attacks B. Hence, A is stable in the new sense.

1 $\Leftarrow$) If A is stable in the new sense, then A is conflict-free. If $x \notin A$, then $A \not\supseteq \{x\}$ so A attacks $\{x\}$ and A is stable in the standard sense.

2 $\Rightarrow$) If A is complete in the standard sense, then A is admissible. Assume B is defended by A (meaning that A attacks all sets attacking B) and let $b \in B$. If $\{b\}$ is not attacked, then A defends $\{b\}$ and therefore $b \in A$. Otherwise, if $\{b\}$ is attacked, A defends B so A defends $\{b\}$ and by completeness $b \in A$. So, $A \supseteq B$ and A is complete in the new sense.

2 $\Leftarrow$) If A is complete in the new sense, then A is admissible. If A defends $\{x\}$, then $A \supseteq \{x\}$, i.e. $x \in A$ and A is complete in the standard sense. □

The next section shows how basic influence diagrams can be seen as extended assumption-based argumentation frameworks. Then Section 3.4 discusses the

unsuitability of the notions of acceptability defined for extended assumption-based argumentation to characterise rational opinions. In the remainder, when clear from the context and by sheer convenience, we will often refer to extended assumption-based argumentation frameworks as assumption-based argumentation frameworks or even simply as argumentation frameworks.

3.3 Basic influence diagrams as argumentation frameworks

Basic influence diagrams can be easily transformed into argumentation frameworks.

Definition 3.11. *Given a basic influence diagram, take*

- $\mathcal{L}$ *as the set of all literals that appear in the diagram, or equivalently in the dependency rules,*
- $\mathcal{R}$ *as the set of all inference rules* $\frac{p_1,\ldots,p_n}{q}$ *where 'if* p_1 *and ... and* p_n *then* q*' is a dependency rule and* $\frac{}{q}$ *where '*q*' is a dependency rule,*
- $\mathcal{A}$ *as the set of all atomic decisions and unobservable fundamental beliefs in the diagram, and*
- $\mathcal{C}$ *as the set of all pairs* (P, q) *such that* $q, \neg q \in \mathcal{L}$ *and* $P \vdash_t \neg q$.

The resulting tuple $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ *is referred to as the extended (assumption-based argumentation) framework* ***derived from*** *the basic influence diagram.*

Here $\mathcal{L}$ is the set of literals labelling nodes of the diagram and $\mathcal{R}$ models both the arcs in the diagram and the dependency rules. $\mathcal{A}$ represents the uncertainties of the decision maker who does not know which atomic decisions to choose nor which unobservable fundamental beliefs hold true. Since atomic decisions and

unobservable fundamental beliefs have no incoming nodes, we can be sure that these assumptions are not the conclusions of any inference rule. Consequently,

Lemma 3.1. *The argumentation framework derived from a basic influence diagram always corresponds to a flat assumption-based argumentation framework.*

In the umbrella problem, the framework derived is given by

- $\mathcal{L}$ = {umbrella, loaded, ¬clouds, ¬rain, dry, clouds, ¬umbrella, ¬loaded, cold, rain, ¬dry}
- $\mathcal{R} = \{\frac{\text{umbrella}}{\text{loaded}}, \frac{\text{umbrella}}{\text{dry}}, \frac{\neg\text{rain}}{\text{dry}}, \frac{\neg\text{umbrella}}{\neg\text{loaded}}, \frac{\neg\text{umbrella,rain}}{\neg\text{dry}}, \frac{\neg\text{clouds}}{\neg\text{rain}}, \frac{\text{clouds,cold}}{\text{rain}}, \frac{}{\text{cold}}\}$
- $\mathcal{A}$ = {umbrella, ¬umbrella, clouds, ¬clouds}
- $\mathcal{C}$ = {({umbrella},¬loaded), ({¬umbrella},loaded), ({umbrella},¬dry), ({¬clouds},¬dry), ({¬umbrella,clouds},dry), ({¬umbrella},umbrella), ({umbrella},¬umbrella), ({clouds},¬rain), ({¬clouds},rain), ({¬clouds},clouds), ({clouds},¬clouds)}

The notion of entailment $\vdash$ for a basic influence diagram and the notion of tight argument $\vdash_t$ for its derived framework are linked in the following way.

Lemma 3.2. *If $A \vdash q$, then $A' \vdash_t q$ for some $A' \subseteq A$. If $A \vdash_t q$, then $A \vdash q$.*

Proof. If $A \vdash q$, then A is a set of assumptions (atomic decisions and unobservable fundamental beliefs) and there exists a sequence $p_1, \ldots, p_n \in \mathcal{L}$ where $p_n = q$ such that for every $i \in \{1, \ldots, n\}$, either p_i is an element of A or p_i is the conclusion of a dependency rule R_i whose premises P_i all belong to $\{p_1, \ldots, p_{i-1}\}$. For each dependency rule R_i there exists by construction an associated inference rule $\frac{P_i}{p_i}$. Define $S_1 = \{p_n\}$ and for every $i \in \{1, \ldots, n\}$ the set $S_{i+1} = S_i - \{p_{n+1-i}\} \cup P_{n+1-i}$ if $p_{n+1-i} \notin A$ and $S_{i+1} = S_i$ otherwise. Let us prove by induction that for every $i \in \{1, \ldots, n+1\}$ the property

$\Pi(i)$ holds, where $\Pi(i)$ means that $S_i \subseteq \{p_1, \ldots, p_{n+1-i}\} \cup A$. $\Pi(1)$ holds since $S_1 = \{p_n\} \subseteq \{p_1, \ldots, p_{n+1-1}\}$. If $\Pi(i)$ holds for $i \in \{1, \ldots, n\}$, then either

- $p_{n+1-i} \in A$, in which case $S_{i+1} = S_i \subseteq \{p_1, \ldots, p_{n+1-i}\} \cup A$ and consequently $S_{i+1} \subseteq \{p_1, \ldots, p_{n+1-(i+1)}\} \cup A$, or
- $p_{n+1-i} \notin A$, in which case $S_{i+1} = S_i - \{p_{n+1-i}\} \cup P_{n+1-i}$. By $\Pi(i)$ we have $S_{i+1} \subseteq \{p_1, \ldots, p_{n+1-i}\} \cup A - \{p_{n+1-i}\} \cup P_{n+1-i}$. Clearly we also have $S_{i+1} \subseteq \{p_1, \ldots, p_{n+1-(i+1)}\} \cup A \cup P_{n+1-i}$. Since $P_{n+1-i} \subseteq \{p_1, \ldots, p_{n+1-i-1}\}$ we finally get $S_{i+1} \subseteq \{p_1, \ldots, p_{n+1-(i+1)}\} \cup A$.

In both cases, $\Pi(i+1)$ holds. So $\Pi(i)$ is true for every $i \in \{1, \ldots, n+1\}$. Since in particular $\Pi(n+1)$ is true, we have $S_{n+1} \subseteq A$. Let $A' = S_{n+1}$. We have $A' \vdash_t q$ (by elimination in the sequence $S_1, \ldots, S_{n+1}$ of the sets S_i where $p_{n+1-i} \in \mathcal{A}$) and $A' \subseteq A$.

If $A \vdash_t q$, then by definition of $\vdash_t$ from Section 3.2.1 there exists a sequence $S_1, \ldots, S_m$ with $S_1 = \{q\}$, $S_m = A$ and for every $i \in \{1, \ldots, m-1\}$, there exists $\sigma_i \in S_i$ such that $S_{i+1} = S_i - \{\sigma_i\} \cup P_i$ for some inference rule $\frac{P_i}{\sigma_i}$. For every $i \in \{1, \ldots, m-1\}$ we have $P_i \subseteq S_{i+1}$ and since there exists a dependency rule corresponding to the inference rule $\frac{P_i}{\sigma_i}$ we also have [2] by construction $S_{i+1} \vdash \sigma_i$. Let us consider a sequence $p_1, \ldots, p_k, p_{k+1}, \ldots, p_{k+m-1}$ such that $\{p_1, \ldots, p_k\} = A$ and

$$(p_{k+1}, \ldots, p_{k+m-i}, \ldots, p_{k+m-1}) = (\sigma_{m-1}, \ldots, \sigma_i, \ldots, \sigma_1)$$

Let us show by induction that for every $i \in \{1, \ldots, m\}$ the property $\Pi(i)$ holds, where $\Pi(i)$ means that $S_i \subseteq \{p_1, \ldots, p_{k+m-i}\}$. $\Pi(m) : S_m = A \subseteq \{p_1, \ldots, p_k\}$ is true. We have $S_{i+1} = S_i - \{\sigma_i\} \cup P_i$ and consequently $S_i \subseteq S_{i+1} \cup \{\sigma_i\}$. Since $\sigma_i = p_{k+m-i}$ it is clear that $\Pi(i+1)$ implies $\Pi(i)$. So $\Pi(i)$ holds for every $i \in$

[2] $\vdash$ is strictly speaking an abuse of notation as S_{i+1} is not necessarily composed of decisions and unobservable fundamental beliefs. We simply mean here that σ_i can be derived from S_{i+1} by application of dependency rules.

$\{1,\ldots,m\}$. We have already shown that $S_{i+1} \vdash \sigma_i$ for every $i \in \{1,\ldots,m-1\}$. By combining these two results we get that $\{p_1,\ldots,p_{k+m-i-1}\} \vdash p_{k+m-i}$ for every $i \in \{1,\ldots,m-1\}$. Since $\sigma_1 \in S_1$ and $S_1 = \{q\}$, we have $\sigma_1 = q$. Moreover, observe that $p_{k+m-1} = \sigma_1 = q$. The sequence $p_1,\ldots,p_{k+m-1}$ constitutes a sequence of inferences from A to q. We conclude that $A \vdash q$. □

In the remainder, we will denote $A \vdash q$ when $A' \vdash_t q$ for some $A' \subseteq A$. This abuse of notation is justified by lemma 3.2. The notion of attack for the framework derived from a basic influence diagram allows to capture the conflicts between pairs of literals p and $\neg p$ in that diagram.

Lemma 3.3. *In the framework derived from a basic influence diagram and for any sets of assumptions $A, B \subseteq \mathcal{A}$, A attacks B if and only if there exists $q, \neg q \in \mathcal{L}$ such that $A \vdash q$ and $B \vdash \neg q$.*

Proof. By definition of an attack from Section 3.2.2, A attacks B if and only if there exists a pair $(P, q) \in \mathcal{C}$ such that $A' \vdash_t q$ and $P \subseteq B$ for some $A' \subseteq A$. By lemma 3.2, the existence of $A' \subseteq A$ and $A' \vdash_t q$ is itself equivalent to $A \vdash q$. By construction of the contrary relation, $(P, q) \in \mathcal{C}$ is true if and only if $P \vdash_t \neg q$. Again by lemma 3.2, $(P, q) \in \mathcal{C}$ and $P \subseteq B$ are true if and only if $B \vdash \neg q$. □

The following property ensures that the attack relation used the in framework derived from a basic influence diagram is always symmetric. This property will play a substantial role in the proofs for the results in Section 3.4 (when the attack relation is symmetric, every conflict-free set of assumptions is also admissible).

Property 3.2. *In the framework derived from a basic influence diagram and for any sets of assumptions A and B, if A attacks B, then B attacks A.*

Proof. If A attacks B, then by lemma 3.3 there exists $q, \neg q \in \mathcal{L}$ such that $A \vdash q$

and $B \vdash \neg q$. Let $p = \neg q$. We have $\neg p = \neg\neg q = q$ so $B \vdash p$ and $A \vdash \neg p$, with $p, \neg p \in \mathcal{L}$. Consequently, by lemma 3.3 again, B attacks A. □

The framework derived from a basic influence diagram models in fact all the information originally contained in the basic influence diagram, except for the positivity and negativity of goals, which are of no importance in the construction of decision tables. It is consequently possible to use assumption-based argumentation for studying basic influence diagrams. We now search for notions of acceptability that allow to exactly characterise all the rational opinions and by entailment all the rational outcomes of a basic influence diagram.

3.4 Rationality & the liberal stable semantics

We are after a notion of acceptability for extended assumption-based argumentation that allows to characterise rational opinions. We may first examine the existing notions of acceptability and start by observing that the conflict-freeness of A is already sufficient to guarantee the consistency of the outcome $O(A)$.

Lemma 3.4. *The outcome $O(A)$ is consistent if and only if the set of assumptions A is conflict-free.*

Proof. It is equivalent to prove that $O(A)$ is inconsistent if and only if A attacks itself. By lemma 3.3, A attacks itself if and only if there exists $q, \neg q \in \mathcal{L}$ such that $A \vdash q$ and $A \vdash \neg q$. By definition 3.4 of $O(A)$, if $A \vdash q$ and $A \vdash \neg q$, then $q \in O(A)$ and $\neg q \in O(A)$. Thus, if A attacks itself then $O(A)$ is inconsistent. Conversely, if $O(A)$ is inconsistent, there exists $q, \neg q \in \mathcal{L}$ such that $q \in O(A)$ and $\neg q \in O(A)$. By definition 3.4 of $O(A)$, $A \vdash q$ and $A \vdash \neg q$. Thus, A attacks itself. □

Unfortunately, all of the previously mentioned notions of acceptability fail to characterise the rationality of opinions.

Theorem 3.1. *Among the conflict-free, naive, admissible, stable, semi-stable, preferred, complete, grounded and ideal semantics, none is such that $O(A)$ is a rational outcome if and only if A is acceptable under the chosen semantics.*

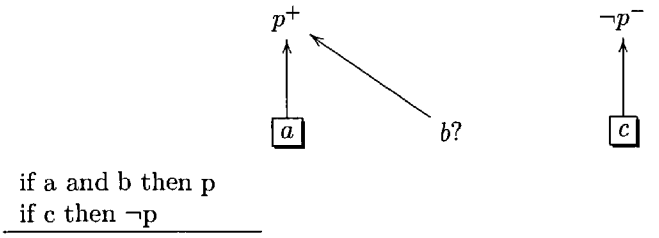


Figure 3.2: Basic influence diagram for theorem 3.1.

Proof. Let us consider the basic influence diagram of Figure 3.2. We obtain a generalised assumption-based framework with $\mathcal{L} = \{a, b, c, p, \neg p\}$, $\mathcal{R} = \{\frac{a,b}{p}, \frac{c}{\neg p}\}$, $\mathcal{A} = \{a, b, c\}$ and $\mathcal{C} = \{(\{a, b\}, \neg p), (\{c\}, p)\}$. The rational opinions are $\{c\}$, $\{a, b\}$, $\{a, c\}$ and $\{b, c\}$. $\emptyset$ is conflict-free but is not rational. $\{c\}$ is not naive (as $\{a, c\}$ is naive) but it is rational. $\emptyset$ is admissible (it is not attacked) but is not rational. $\{c\}$ is not stable (it attacks neither $\{a\}$ nor $\{b\}$) but it is rational. $\{c\}$ is not preferred (as $\{a, c\}$ is preferred) but it is rational. $\{c\}$ is not complete (it defends $\{a\}$ since $\{a\}$ is not attacked but does not contain a) but it is rational. $\{c\}$ is not grounded as $\{c\}$ is not complete but it is rational. $\{c\}$ is not semi-stable as it is not complete, but it is rational. $\emptyset$ is ideal (it is the only admissible set included in the empty intersection of the preferred sets $\{a, b\}$, $\{a, c\}$ and $\{b, c\}$) but is not rational. □

This justifies the need for a new semantics, corresponding to the rational opinions of basic influence diagrams. For any set of assumptions $A \subseteq \mathcal{A}$, let us first define

Definition 3.12. $A^+ = \{p \in \mathcal{L} \mid A \vdash p$ *and there exists* B *such that* $(B, p) \in \mathcal{C}\}$

Property 3.3. *In the framework derived from a basic influence diagram, it holds that* $A^+ = \{p \in \mathcal{L} \mid A \vdash p, \neg p \in \mathcal{L}\}$.

Proof. For any $p \in \mathcal{L}$, if there exists B such that $(B, p) \in \mathcal{C}$, then by construction of $\mathcal{C}$ in the derived framework we have $B \vdash_t \neg p$. Thus, $\neg p \in \mathcal{L}$. For any $p \in \mathcal{L}$, if $\neg p \in \mathcal{L}$, then we can construct a tight argument $B \vdash_t \neg p$ and again by construction of $\mathcal{C}$, $(B, p) \in \mathcal{C}$. □

For any set of assumptions A, we say that

Definition 3.13. *A is **liberal stable** if and only if*

- *A is conflict-free, and*
- *there is no conflict-free set B of assumptions such that $B^+ \supset A^+$.*

In the umbrella problem, the liberal stable sets of assumptions are

- {umbrella, clouds}
- {umbrella, ¬clouds}
- {¬umbrella, clouds}
- {¬umbrella, ¬clouds}

since these sets are conflict-free sets of assumptions where

$$\mathcal{A} = \{\text{umbrella}, \neg\text{umbrella}, \text{clouds}\neg\text{clouds}\}$$

and any set of assumptions B such that $B^+ \supset A^+$ must here be such that $B \supset A$ and thus would not be conflict-free. Here, the liberal stable sets of assumptions coincide with the rational opinions. Note that they also coincide with the naive sets of assumptions.

The liberal stable sets of assumptions also coincide with the rational opinions in the example of Figure 3.2. Indeed, note first that the conflict-free sets of assumptions are the strict subsets of $\mathcal{A}$. It is clear that in this example, if a set of assumptions A is such that both $A \nvdash p$ and $A \nvdash \neg p$ then $A^+ = \emptyset$. Such a set A is not liberal stable. Indeed, consider the conflict-free set $B = \{c\}$. We have $B^+ = \{\neg p\} \supset \emptyset = A^+$. So, it is necessary for a liberal stable set to be such that $A \vdash p$ or $A \vdash \neg p$. The only conflict-free set such that $A \vdash p$ is $A_1 = \{a, b\}$ and the only conflict-free sets such that $A \vdash \neg p$ are $A_2 = \{c\}$, $A_3 = \{a, c\}$ and $A_4 = \{b, c\}$. For all these sets, the only B such that $B^+ \supset A^+$ is $B = \mathcal{A}$ which is not conflict-free. Consequently, A_1, A_2, A_3 and A_4 are the liberal stable sets of assumptions. They exactly correspond to the rational opinions as seen in the proof of theorem 3.1.

Liberal stable sets of assumptions are guaranteed to exist as long as a conflict-free set of assumptions exists.

Property 3.4. *If there exists a conflict-free sets of assumptions, then there exists a liberal stable set of assumptions.*

Proof. If there exists a conflict-free set A of assumptions, then there exists a conflict-free set B such that B^+ is maximal with respect to set inclusion. By definition, B must be a liberal stable set of assumptions. □

However, in extended assumption-based argumentation, there may not always exist some conflict-free set of assumptions. For instance, if in a basic influence diagram q and $\neg q$ are facts, then in the derived framework, $\emptyset$ attacks itself, and consequently, all sets of assumptions also attack themselves.

We prove generally, in theorem 3.2 below, that the liberal stable sets correspond exactly to the rational opinions whenever the framework has a property referred to as *extensibility*.

Definition 3.14. *The framework derived from a basic influence diagram is said to be* ***extensible*** *if and only if for every naive set N of assumptions the outcome $O(N)$ is decided.*

Checking the property of extensibility thus requires the preliminary computation of all the naive sets. However, we may state that frameworks fail to be extensible notably when there exists $(p, \neg p) \in \mathcal{L}^2$ such that p and $\neg p$ are entailed by the same sets of assumptions. The existence of a simple but general characterisation of extensibility remains an open issue.

We have as an immediate corollary of the previous definition:

Corollary 3.1. *If the framework derived from a basic influence diagram is extensible, then every conflict-free set of assumptions is contained in at least one rational opinion.*

Proof. Let A be a conflict-free set. There exists a naive set N such that $A \subseteq N$. By extensibility, $O(N)$ is decided. Since N is conflict-free and by lemma 3.4, the outcome $O(N)$ is also decided. $O(N)$ is obviously closed under $\vdash$, so $O(N)$ is a rational outcome and N is a rational opinion. □

In other words, when extensibility holds, any conflict-free set of assumptions can be "extended" to a larger set of assumptions that constitutes a rational opinion. Hence the choice of the name "extensibility".

We now show that in extensible frameworks, liberal stability exactly characterises the rational opinions. Formally,

Theorem 3.2. *i) If A is a rational opinion, then A is a liberal stable set of assumptions. If the argumentation framework is extensible, we also have ii) if A is a liberal stable set of assumptions, then A is a rational opinion.*

Proof. According to lemma 3.4 and the definition of a rational outcome, if A is rational then $O(A)$ is consistent and A is conflict-free. Also, if A is liberal stable

then A is conflict-free and $O(A)$ is consistent. Consequently, it only remains to prove in i) that there is no conflict-free set B such that $B^+ \supset A^+$, and in ii) that for every pair $(p, \neg p) \in \mathcal{L}^2$, either $p \in O(A)$ or $\neg p \in O(A)$.

Proof of i): Let A be a rational opinion. Assume by contradiction that there is a conflict-free set B such that $B^+ \supset A^+$. Then $B^+ - A^+ \neq \emptyset$, so there exists $p \in B^+ - A^+$. Thus $B \vdash p$ and $A \nvdash p$. A is rational and therefore either $A \vdash \neg p$ or $A \vdash p$. We know that $A \nvdash p$ so by elimination $A \vdash \neg p$. By hypothesis, $B^+ \supset A^+$, therefore $B \vdash \neg p$. $B \vdash p$ and $B \vdash \neg p$ so by lemma 3.3, B attacks itself. This is absurd because B is conflict-free by hypothesis.

Proof of ii): Let A be a liberal stable set of assumptions and $p, \neg p \in \mathcal{L}$. Assume by contradiction that $p \notin O(A)$ and $\neg p \notin O(A)$. A is liberal stable and a fortiori conflict-free. Since we have assumed that the framework is extensible, by corollary 3.1, A can be extended to a set $A' \supseteq A$ that is rational. Clearly, $A' \supseteq A$ implies that $A'^+ \supseteq A^+$. By rationality of A' and lemma 3.4, A' is conflict-free and such that $A' \vdash p$ or $A' \vdash \neg p$. We may assume without loss of generality that $A' \vdash p$. Then $A'^+ \supset A^+$ since $p \in A'$ but $p \notin A$. This is absurd because A is liberal stable by hypothesis. □

Note that certain frameworks are not extensible, for instance if they admit at least one conflict-free set of assumptions but no rational outcome. As an example, consider the basic influence diagram of Figure 3.3 and its derived extended framework

- $\mathcal{L} = \{d, p, \neg p\}$
- $\mathcal{R} = \{\frac{d}{p}, \frac{p}{\neg p}\}$
- $\mathcal{A} = \{d\}$
- $\mathcal{C} = \{(\{d\}, p), (\{d\}, \neg p)\}$

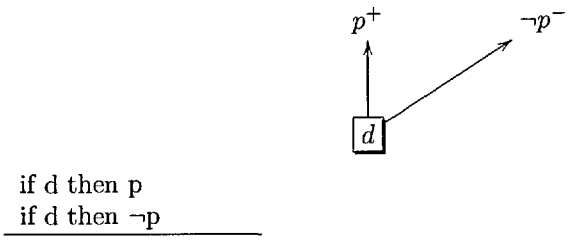


Figure 3.3: Pathological case of basic influence diagram.

in which $\emptyset$ is conflict-free but the only decided outcome $O(\{d\}) = \{d, p, \neg p\}$ is not consistent. Such a framework and its derived basic influence diagram are pathological as they do not conform to the idea that, in the decision domain, either p or $\neg p$ holds, but p and $\neg p$ cannot hold simultaneously. Thus, arguably, no such framework can be derived from a basic influence diagram.

In the remainder, we make the hypothesis that extensibility is satisfied in the framework derived from the basic influence diagram. We will not treat the case of basic influence diagrams that do not fulfil this natural property. Indeed, without this property, liberal stability would provide all rational opinions but could undesirably contain extensions that are not rational.

Clearly, the liberal stable semantics differs from the other ones. In fact, it relates to the other notions of acceptability in the following manner.

Theorem 3.3. *For any framework derived from a basic influence diagram, every stable set is liberal stable and every liberal stable set is conflict-free and admissible. For any framework derived from a basic influence diagram that fulfils the property of extensibility, every naive, stable or preferred set is liberal stable (see Figure 3.4).*

In Figure 3.4, $\surd$ denotes an implication that holds in general, E an implication that holds under extensibility and $\times$ an implication that does not generally hold, even under extensibility. The symbols in exponent indicate the paragraph

of the proof below in which the result is proved.

Semantics s	$s \Rightarrow$ liberal stable ?	liberal stable $\Rightarrow s$?
conflict-free	$\times^*$	$\surd^0$
naive	E^1	$\times^*$
admissible	$\times^*$	$\surd^2$
stable	$\surd^3$	$\times^*$
semi-stable	$\times^4$	$\times^*$
preferred	E^5	$\times^*$
complete	$\times^6$	$\times^*$
grounded	$\times^7$	$\times^*$
ideal	$\times^*$	$\times^8$

Figure 3.4: Links between liberal stability and other semantics.

Proof. *) These results derive from theorem 3.2 and the counter-examples of Figure 3.2 previously examined in context of the proof of theorem 3.1, where we have seen that a set of assumptions could be

- conflict-free, admissible or ideal but not rational, e.g. $\emptyset$,
- liberal stable but not naive, stable, semi-stable, preferred, or complete, e.g. $\{c\}$ or grounded e.g. $\{a, c\}$.

0) Every liberal stable set is conflict-free by definition.

1) Let A be a naive set. A is conflict-free. Assume there exists a conflict-free set B such that $B^+ \supset A^+$. Then, there would exist $p \in B^+ - A^+$. $A \vdash \neg p$ otherwise A could be extended by extensibility to a conflict-free set that entails p or $\neg p$ and that set would be strictly larger than A, which is absurd because A is maximally conflict-free. So, $\neg p \in A^+$. Since $B^+ \supset A^+$, $\neg p \in B^+$ and $B \vdash \neg p$. Thus, B attacks itself, which is absurd.

2) It is easy to see that when the relation of attack is symmetric, the notion of admissibility collapses with the notion of conflict-freeness. This happens in the frameworks derived from basic influence diagrams, as they are symmetric

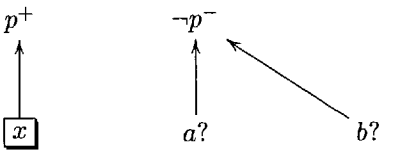
by property 3.2. Since liberal stability implies conflict-freeness, it is also clear that liberal stability implies admissibility.

3) Let A be a stable set. A is conflict-free. Assume by contradiction, that there exists a conflict-free set B such that $B^+ \supset A^+$. A does not include B, as we would otherwise have $B^+ \subseteq A^+$. Thus, by stability, A attacks B. So there exists $p \in \mathcal{L}$ such that $p \in A^+$ and $\neg p \in B^+$. Since $B^+ \supset A^+$, $p \in B^+$ so B attacks itself, which is absurd.

5) Preferred sets are defined as maximally admissible and by symmetry (property 3.2) actually coincide with the maximally conflict-free (naive) ones. The proof of 1) allows us to conclude that preferred sets are liberal stable.

6) Consider the basic influence diagram of Figure 3.5 and its derived framework where

- $\mathcal{L} = \{a, b, x, p, \neg p\}$,
- $\mathcal{R} = \{\frac{x}{p}, \frac{b}{\neg p}, \frac{a,b}{\neg p}\}$,
- $\mathcal{A} = \{a, b, x\}$ and
- $\mathcal{C} = \{(\{x\}, \{\neg p\}), (\{b\}, \{p\}), (\{a, b\}, \{p\})\}$.



if x then p
if b then ¬p
if a and b then ¬p

Figure 3.5: Basic influence diagram where a complete set of assumptions is not liberal stable.

Here $\{a\}$ is not attacked, so it is admissible. $\{a\}$ does not attack any set, so the only sets it defends are the sets of assumptions that are not attacked. The

sets $\{x\}$ and $\{b\}$ and their supersets of assumptions are all attacked. Therefore, $\{a\}$ is complete. However, it is not liberal stable since $\{a\}^+ = \emptyset \subset \{b\}^+$ and $\{b\}$ is conflict-free.

7) In the same framework as in 6), note that all complete sets must include $\{a\}$ (which is not attacked) and therefore $A = \{a\}$ is minimally complete (grounded). However, A is not liberal stable, as proved in 6).

4) Semi-stable sets must be complete. In the same framework as in 6), all complete sets contain a. So, $\{x\}$ is neither complete nor semi-stable. However, $\{x\}$ is liberal stable, since it is conflict-free and any set B such that $B^+ \supseteq \{x\}^+ = \{p\}$ would be such that $B^+ = \{p, \neg p\}$ and by lemma 3.3 would not be conflict-free.

8) Consider again the basic influence diagram of Figure 3.2 and its derived framework, where

- $\mathcal{L} = \{a, b, c, p, \neg p\}$,
- $\mathcal{R} = \{\frac{a,b}{p}, \frac{c}{\neg p}\}$,
- $\mathcal{A} = \{a, b, c\}$ and
- $\mathcal{C} = \{(\{a, b\}, \neg p), (\{c\}, p)\}$.

By symmetry of the framework (property 3.2) the admissible sets of assumptions correspond to the conflict-free sets of assumptions. Therefore, the preferred sets of assumptions correspond to the naive sets of assumptions. The preferred sets are thus $\{a, b\}$, $\{c, a\}$ and $\{c, b\}$. Ideal sets must be admissible and contained in the intersection $\{a, b\} \cap \{c, a\} \cap \{c, b\} = \emptyset$. So, the only ideal set of assumptions is $\emptyset$ but $\emptyset$ is not liberal stable. □

3.5 Computing the liberal stable semantics

We have shown in the previous theorem 3.3 that, under extensibility, every naive set of assumptions is liberal stable. In fact, we will see that the liberal stable sets of assumptions can easily be computed from the naive sets of assumptions. First, we explain how to compute the naive sets of assumptions using known results from hypergraph theory.

3.5.1 Computing the naive semantics

Let V be a finite set and let us refer to its elements as *vertexes*. In hypergraph theory [Eiter and Gottlob, 1995; Boros *et al.*, 2000; Eiter and Gottlob, 2002; Khachiyan *et al.*, 2005], one defines a *hyperedge* as a subset of V. A *hypergraph* is a pair $(V, \mathcal{H})$ where $\mathcal{H}$ is a set of hyperedges. A vertex set $S \subseteq V$ is said *independent* if it does not include any hyperedge of $\mathcal{H}$. Given the framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ derived from a basic influence diagram, we consider the hypergraph where $V = \mathcal{A}$ and $\mathcal{H}$ is the set of *minimally conflicting* sets of assumptions with respect to set inclusion. A set of assumptions is said to be conflicting when it attacks itself. This set of hyperedges $\mathcal{H}$ is obtained using Algorithm 1 below.

Algorithm 1 Computes the set $\mathcal{H}$ of minimally conflicting sets given as input an extended assumption-based argumentation framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ derived from a basic influence diagram.

1: **Initialise:** $\mathcal{H} \leftarrow \emptyset$
2: **for every** pair of the form $(p, \neg p) \in \mathcal{L}^2$ **do**
3: **for every** $A, B \subseteq \mathcal{A}$ such that $A \vdash_t p$ and $B \vdash_t \neg p$ **do**
4: $\mathcal{H} \leftarrow \mathcal{H} \cup \{A \cup B\}$
5: **end for**
6: **end for**
7: remove from $\mathcal{H}$ all its non minimal elements
8: **return** $\mathcal{H}$

This algorithm is correct since sets of assumptions X that attack themselves are such that there exists $p, \neg p \in \mathcal{L}$ such that $A \vdash_t p$ and $B \vdash_t \neg p$ for some

$A, B \subseteq X$. At line 7, all conflicting sets that are not minimal are removed.

In the example of Figure 3.2, the tight arguments considered are $\{a, b\} \vdash_t p$ and $\{c\} \vdash \neg p$ so $\mathcal{H} = \{\{a, b, c\}\}$. In the umbrella problem of Figure 3.1, the tight arguments are

- $\{\text{umbrella}\} \vdash_t \text{umbrella}$,
- $\{\text{umbrella}\} \vdash_t \text{loaded}$,
- $\{\text{umbrella}\} \vdash_t \text{dry}$,
- $\{\neg\text{clouds}\} \vdash_t \text{dry}$,
- $\{\neg\text{clouds}\} \vdash_t \neg\text{rain}$,
- $\{\neg\text{clouds}\} \vdash_t \neg\text{clouds}$,
- $\{\text{clouds}\} \vdash_t \text{clouds}$,
- $\{\text{clouds}\} \vdash_t \text{rain}$,
- $\{\neg\text{umbrella}\} \vdash_t \neg\text{umbrella}$,
- $\{\neg\text{umbrella}\} \vdash_t \neg\text{loaded}$ and
- $\{\neg\text{umbrella}, \text{clouds}\} \vdash_t \neg\text{dry}$.

The conflicting sets computed after the end of the loop in line 6 are

- $\{\text{umbrella}, \neg\text{umbrella}\}$,
- $\{\text{umbrella}, \neg\text{umbrella}, \text{clouds}\}$,
- $\{\neg\text{clouds}, \text{clouds}\}$ and
- $\{\neg\text{umbrella}, \text{clouds}, \neg\text{clouds}\}$.

After elimination of the non-minimal sets, we are left with

$$\mathcal{H} = \{\{\text{umbrella}, \neg\text{umbrella}\}, \{\neg\text{clouds}, \text{clouds}\}\}$$

Observe that the conflict-free sets of assumptions correspond exactly to the independent sets of the hypergraph obtained from the derived framework. Therefore, trivially:

Lemma 3.5. *The naive sets of assumptions correspond exactly to the maximal independent sets of the hypergraph $(V, \mathcal{H})$, where $V = \mathcal{A}$ and $\mathcal{H}$ is the set of minimally conflicting sets of assumptions in the framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ derived from a basic influence diagram.*

The problem of finding the set $\mathcal{I}(\mathcal{H})$ of all maximal independent sets of a hypergraph has been tackled in the literature on hypergraphs [Eiter and Gottlob, 1995; Boros *et al.*, 2000; Khachiyan *et al.*, 2005]. We can use Algorithm 2 below to compute $\mathcal{I}(\mathcal{H})$ and thus compute the naive semantics. The algorithm involves objects called *transverses* of $\mathcal{H}$, which are sets of vertexes that intersect every hyperedge of $\mathcal{H}$. A transverse is called *minimal* if no strict subset of it is a transverse. The set of all minimal transverses of $\mathcal{H}$ is denoted $Tr(\mathcal{H})$. For every transverse T, $\overline{T}$ denotes the complement $V - T$.

After line 13 of the algorithm, $\mathcal{G}$ is equal to $Tr(H)$ as this part of the algorithm corresponds to the algorithm proposed by [Eiter and Gottlob, 1995] to compute the set of all minimal transverse of a hypergraph. This part of the program (still up to line 13) can alternatively be replaced by the algorithm of [Boros *et al.*, 2000]. At line 14, $\mathcal{I}(\mathcal{H})$ is equal to the complement of $Tr(\mathcal{H})$ and according to [Khachiyan *et al.*, 2005] is equal to the set $\mathcal{I}(\mathcal{H})$ of all maximal independent sets of $\mathcal{H}$. The maximal independent sets of the argumentation hypergraph exactly corresponds to the set of naive sets of assumptions, as sanc-

Algorithm 2 Computes the set $\mathcal{N}$ of all naive sets of assumptions given as input the hypergraph $(V, \mathcal{H})$, where $V = \mathcal{A}$ and $\mathcal{H}$ is the set of all minimally conflicting sets of $\mathcal{A}$ in the extended framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ derived from a basic influence diagram.

1: $k \leftarrow max(|E| : E \in \mathcal{H})$; $\mathcal{G} \leftarrow \emptyset$
2: **loop**
3: **if** $\exists T \in 2^V, |T| \leq k : T \in Tr(G) \wedge T \notin \mathcal{H}$ **then**
4: reduce $\overline{T}$ to a minimal transverse T_1
5: **else**
6: **if** $\exists \mathcal{G}' \subseteq \mathcal{G}, |\mathcal{G}'| = k+1 : \forall E \in \mathcal{G} : E \not\subseteq \{x | x \in V, |\{E \in \mathcal{G}', x \in E\}| > 1\}$ **then**
7: reduce $\{x | x \in V, |\{E \in \mathcal{G}', x \in E\}| > 1\}$ to a minimal transverse T_1
8: **else**
9: **exit loop**
10: **end if**
11: **end if**
12: $\mathcal{G} \leftarrow \mathcal{G} \cup \{T_1\}$;
13: **endloop**
14: $\mathcal{I}(\mathcal{H}) \leftarrow \{V - H | H \in \mathcal{G}\}$
15: $\mathcal{N} \leftarrow \mathcal{I}(\mathcal{H})$
16: **return** $\mathcal{N}$

tioned by lemma 3.5. Thus Algorithm 2 is correct.

This algorithm has been implemented in Maple version 9.5. In the example of Figure 3.2, we have $V = \{a, b, c\}$ and $\mathcal{H} = \{\{a, b, c\}\}$. We obtain the naive sets $\mathcal{N} = \{\{a, b\}, \{a, c\}, \{b, c\}\}$. In the umbrella problem, we have $V = \{\text{umbrella}, \neg\text{umbrella}, \text{clouds}, \neg\text{clouds}\}$ and the set of hyperedges $\mathcal{H} = \{\{\text{umbrella}, \neg\text{umbrella}\}, \{\text{clouds}, \neg\text{clouds}\}\}$. We obtain $\mathcal{N} = \{\{\text{umbrella}, \text{clouds}\}, \{\text{umbrella}, \neg\text{clouds}\}, \{\neg\text{umbrella}, \text{clouds}\}, \{\neg\text{umbrella}, \neg\text{clouds}\}\}$.

3.5.2 From the naive to the liberal stable semantics

In the remainder, we assume that the argumentation framework obtained by translation of the basic influence diagram satisfies the property of extensibility. Then, by theorem 3.3, we know that every naive set of assumptions is liberal stable. The potentially remaining liberal stable sets can then be identified by exploiting the following theorem.

Theorem 3.4. *Under extensibility, A is liberal stable if and only if there exists a naive set of assumptions N such that $A \subseteq N$ and $A^+ = N^+$.*

Proof. Let A be a liberal stable set of assumptions. A is conflict-free, so there exists a naive set N such that $A \subseteq N$. This implies $A^+ \subseteq N^+$. By liberal stability, it is impossible that $A^+ \subset N^+$, so $A^+ = N^+$.

Conversely, assume that A is a set of assumptions such that $A \subseteq N$ and $A^+ = N^+$ where N is a naive set. Since N is conflict-free and $A \subseteq N$, A is also conflict-free. Assume, by contradiction, that there exists a conflict-free set B such that $B^+ \supset A^+$. Since $A^+ = N^+$, we have $B^+ \supset N^+$. Clearly, B^+ cannot be empty, which implies that there exists at least one literal $q \in \mathcal{L}$ (entailed by B) such that $\neg q \in \mathcal{L}$. By extensibility, since N is naive, $O(N)$ is decided, but then there must exist a pair $(p, \neg p) \in \mathcal{L}^2$ such that $B \vdash p$ and $B \vdash \neg p$. Thus, B attacks itself, which is absurd. □

The set of liberal stable sets of assumptions can be consequently derived from the set of naive sets of assumptions by using Algorithm 3 below.

Algorithm 3 Computes the set $\mathcal{S}$ of all liberal stable sets of assumptions, given as input the set $\mathcal{N}$ of all naive sets of assumptions in an extensible framework derived from a basic influence diagram.

1: **Initialise:** $\mathcal{S} \leftarrow \emptyset$
2: **for every** $N \in \mathcal{N}$ **do**
3: $C \leftarrow N^+$ and $k \leftarrow |C|$
4: **for every** $p_j \in \{p_1, \ldots, p_k\} = C$ **do**
5: $R_j \leftarrow \{P_{i,j} \mid P_{i,j} \vdash_t p_j, P_{i,j} \subseteq N\}$
6: **end for**
7: **for every** $(P_{\bullet,1}, ..., P_{\bullet,k}) \in R_1 \times ... \times R_k$ **do**
8: $P \leftarrow P_{\bullet,1} \cup \ldots \cup P_{\bullet,k}$
9: $\mathcal{S} \leftarrow \mathcal{S} \cup \{X \mid P \subseteq X \subseteq N\}$
10: **end for**
11: **end for**
12: **return** $\mathcal{S}$

This algorithm finds all the conflict-free sets A such that $A^+ = N^+$ and $A \subseteq N$ for some naive set of assumptions N. For each such N, it first computes

N^+ (line 3) and gathers all the subsets of N (guaranteed to be conflict-free) supporting any argument with conclusion $p \in N^+$ for every $p \in N^+$ (lines 4 to 6). The sets P obtained in line 8 fulfill $P^+ \supseteq N^+$ by construction, and are in fact such that $P^+ = N^+$ since $P \subseteq N$ implies $P^+ \subseteq N^+$. Under extensibility, the previous theorem 3.4 guarantees that $\mathcal{S}$ corresponds exactly to the set of all liberal stable sets of assumptions. Thus, Algorithm 3 is correct.

In the example of Figure 3.2, the naive sets of assumptions are $N_1 = \{a, b\}$, $N_2 = \{b, c\}$ and $N_3 = \{a, c\}$. We have $N_1^+ = \{p\}$, $N_2^+ = \{\neg p\}$ and $N_3^+ = \{\neg p\}$. The only tight argument supporting p is $\{a, b\} \vdash_t p$ and the only tight argument supporting $\neg p$ is $\{c\} \vdash_t \neg p$. Therefore the liberal stable sets are all the sets of assumptions that contain either $\{a, b\}$ or $\{c\}$ but not both. These are $\{a, b\}$, $\{c\}$, $\{b, c\}$ and $\{a, c\}$. Here, $\{c\}$ is liberal stable and its outcome is rational, although $\{c\}$ is not naive.

As shown at the end of Section 3.5.1 for the umbrella problem, the naive sets of assumptions are

- $N_1 = \{\text{umbrella}, \text{clouds}\}$,
- $N_2 = \{\text{umbrella}, \neg\text{clouds}\}$,
- $N_3 = \{\neg\text{umbrella}, \text{clouds}\}$ and
- $N_4 = \{\neg\text{umbrella}, \neg\text{clouds}\}$.

We have then

- $N_1^+ = \{\text{umbrella}, \text{clouds}, \text{loaded}, \text{rain}, \text{dry}\}$,
- $N_2^+ = \{\text{umbrella}, \neg\text{clouds}, \text{loaded}, \neg\text{rain}, \text{dry}\}$,
- $N_3^+ = \{\neg\text{umbrella}, \text{clouds}, \neg\text{loaded}, \text{rain}, \neg\text{dry}\}$ and
- $N_4^+ = \{\neg\text{umbrella}, \neg\text{clouds}, \neg\text{loaded}, \neg\text{rain}, \text{dry}\}$.

The tight arguments are

- $\{\text{umbrella}\} \vdash_t \text{umbrella}$,
- $\{\text{umbrella}\} \vdash_t \text{loaded}$,
- $\{\text{umbrella}\} \vdash_t \text{dry}$,
- $\{\neg\text{clouds}\} \vdash_t \text{dry}$,
- $\{\neg\text{clouds}\} \vdash_t \neg\text{rain}$,
- $\{\neg\text{clouds}\} \vdash_t \neg\text{clouds}$,
- $\{\text{clouds}\} \vdash_t \text{clouds}$,
- $\{\text{clouds}\} \vdash_t \text{rain}$,
- $\{\neg\text{umbrella}\} \vdash_t \neg\text{umbrella}$,
- $\{\neg\text{umbrella}\} \vdash_t \neg\text{loaded}$ and
- $\{\neg\text{umbrella}, \text{clouds}\} \vdash_t \neg\text{dry}$.

The liberal stable sets are

- $\{\text{umbrella}, \text{clouds}\}$,
- $\{\text{umbrella}, \neg\text{clouds}\}$,
- $\{\neg\text{umbrella}, \text{clouds}\}$ and
- $\{\neg\text{umbrella}, \neg\text{clouds}\}$.

Here, the liberal stable sets exactly coincide with the naive sets, as already observed in Section 3.4. More generally, we have

Property 3.5. *In an extensible framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ derived from a basic influence diagram, if the set of assumptions $\mathcal{A}$ is closed under $\neg$, then the liberal stable sets of assumptions correspond to the naive sets of assumptions.*

Proof. By theorem 3.4, a set of assumptions A is liberal stable if and only if there exists a naive set N of assumptions such that $A \subseteq N$ and $A^+ = N^+$. When $\mathcal{A}$ is closed under $\neg$ and for any set of assumptions $X \subseteq \mathcal{A}$, if $x \in X$, then $x \in X^+$. Since the framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ derived from a basic influence diagram is always flat by lemma 3.1, for any set of assumptions A, if $A \subset N$, then $A^+ \subset N^+$. Thus, any liberal stable set A must be such that $A = N$ for some naive set N of assumptions. Conversely, we know by theorem 3.3 that under extensibility every naive set of assumptions is liberal stable. □

In such situations, it is perfectly unnecessary to use Algorithm 3, as the naive sets of assumptions are already given by Algorithm 2.

3.6 Computing decision tables

Having computed the liberal stable semantics, we finally may generate a full decision table. The scenarios can be simply obtained by considering the unobservable fundamental beliefs that underlie the possible outcomes. Denoting by B the set of unobservable fundamental beliefs in the basic influence diagram and $\mathcal{S}$ the set of liberal stable sets of assumptions of the derived extended assumption-based argumentation framework,

Definition 3.15. *The set of scenarios or* ***universe*** *is*

$$\Omega = \{w \mid w = S \cap B, S \in \mathcal{S}\}$$

Note that each scenario w is a set (of unobservable fundamental beliefs) and the universe Ω is a set of sets.

The possible decisions correspond to the sets of atomic decisions X which always lead to a rational outcome (given by the liberal stable sets of assump-

tions), no matter the scenario $w \in \Omega$ that occurs. Thus, if we denote by Δ the set of atomic decisions of the basic influence diagram, then

Definition 3.16. *The set of* ***possible decisions*** *is*

$$D = \{X \mid X \subseteq \Delta, \forall w \in \Omega : X \cup w \in \mathcal{S}\}$$

Possible decisions that contain more than one atomic decision are interpreted as combined decisions. Combined decisions are practically 'possible' in the sense that they always lead to a rational outcome, no matter the scenario which occurs. Good decisions in practise often need to be the combination of elementary (atomic) decisions [3].

Finally, the outcome $T(d, w)$ of every possible decision $d \in D$ and scenario $w \in \Omega$ is given by the set of all literals entailed by d and w. Thus,

Definition 3.17. *The* ***decision table*** $T : D \times \Omega$ *is defined* $\forall (d, w) \in D \times \Omega$ *as*

$$T(d, w) = \{p \mid p \in \mathcal{L}, d \cup w \vdash p\}$$

In the example of Figure 3.2, we have $\mathcal{S} = \{\{a, b\}, \{c\}, \{b, c\}, \{a, c\}\}$, $B = \{b\}$ and $\Delta = \{a, c\}$. Thus $\Omega = \{\{b\}, \emptyset\}$, $D = \{\{c\}\}$ and T as shown in the decision table of Figure 3.6.

.	$w_1 = \emptyset$	$w_2 = \{b\}$
$d_1 = \{c\}$	$T(d_1, w_1) = \{c, \neg p\}$	$T(d_1, w_2) = \{c, b, \neg p\}$

Figure 3.6: Decision table for the example of Figure 3.2.

In the umbrella problem, the liberal stable sets are $\mathcal{S} = \{S_1, S_2, S_3, S_4\}$ with $S_1 = \{\text{umbrella}, \text{clouds}\}$, $S_2 = \{\text{umbrella}, \neg\text{clouds}\}$, $S_3 = \{\neg\text{umbrella}, \text{clouds}\}$

[3] For instance, a doctor may choose to prescribe two drugs in combination in order to cure a patient and prevent the manifestation of undesirable secondary effects.

and $S_4 = \{\neg\text{umbrella}, \neg\text{clouds}\}$. We have $B = \{\text{clouds}, \neg\text{clouds}\}$ and $\Delta = \{\text{umbrella}, \neg\text{umbrella}\}$. Thus, we get the universe $\Omega = \{\{\text{clouds}\}, \{\neg\text{clouds}\}\}$, the set of possible decisions $D = \{\{\text{umbrella}\}, \{\neg\text{umbrella}\}\}$ and T is shown in the decision table of Figure 3.7.

.	$w_1 = \{\text{clouds}\}$	$w_2 = \{\neg\text{clouds}\}$
$d_1 = \{\text{umbrella}\}$	$T(d_1, w_1) =$ {umbrella, loaded, dry, clouds, cold, rain}	$T(d_1, w_2) =$ {umbrella, loaded, dry, ¬clouds, cold, ¬rain}
$d_2 = \{\neg\text{umbrella}\}$	$T(d_2, w_1) =$ {¬umbrella, ¬loaded, clouds, cold, rain, ¬dry}	$T(d_2, w_2) =$ {¬umbrella, ¬loaded, ¬clouds, cold, ¬rain, dry}

Figure 3.7: Decision table for the umbrella problem.

3.7 Application to the construction of a decision table for poker

Poker is perhaps one of the few games which is still considered a challenge in the artificial intelligence community [Billings *et al.*, 1999], unlike chess, backgammon or othello. In its most popular version, called Texas Hold'Em poker [Mendelson, 2005], each player is given by a dealer two cards called *pocket cards* which he keeps secret from the other players. There can be from two to eight players at the table. Other cards known as *community cards* are shown progressively (there are 5 cards in total) with which the players can compose their *hand* (best possible combination of 5 cards drawn from the pool of pocket and community cards). At poker, players must constantly bet and take risks under uncertainty to win. Playing a hand is hard, as only partial information about the hand of the other players is available. Players decide how they want to bet but their decisions trigger behaviours from the other players which can be predicted only

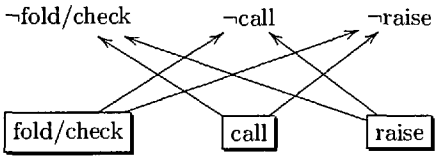
very qualitatively. Good players use poker tells [Caro, 2003] to better judge the strength of their opponents. Understanding tells involves understanding some psychological mechanisms and body language and is not an exact science. Often, the tells picked up by a player may conflict, because good opponents know how to mislead those who observe them. Poker players are actors sending wrong and contradictory signals to their opponents to confuse them, but some tells are more reliable than others. It is important for poker players and poker playing programs to quickly identify the possible scenarios and their corresponding outcomes before taking any decision.

There are essentially four possible moves at poker

- fold: withdraw from further participation in the current pot
- check: make no bet, but still hold your cards
- call: match a bet
- raise: increase the bet

These are the decisions to choose from. We may however treat folding and checking as one decision. Indeed, when checking is possible, a player should always check rather than fold. Checking allows the player to stay in the hand without betting more money. This is clearly always preferable to folding. Nevertheless, a player can only check if his bet already equals the maximum betting amount of the other players. We therefore reduce the problem to three decisions: d_1 = fold/check, d_2 = call and d_3 = raise. These decision are mutually exclusive. This needs to be expressed logically and we can simply do so using the following portion of basic influence diagram shown in Figure 3.8.

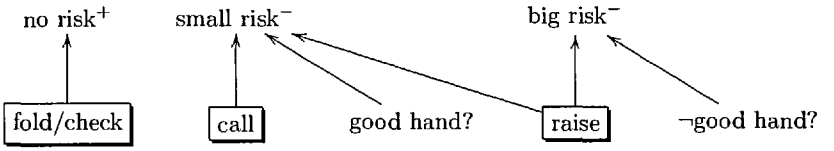
All decisions have their pros and cons. For instance some moves involve taking risk, others don't. The risk in fact also very much depends on the quality



if fold/check then ¬call
if fold/check then ¬raise
if call then ¬fold/check
if call then ¬raise
if raise then ¬fold/check
if raise then ¬call

Figure 3.8: Poker basic influence diagram (part 1).

of the player's hand. The amount of risk involved by each decision can be defined by another portion of the basic influence diagram (see Figure 3.9).

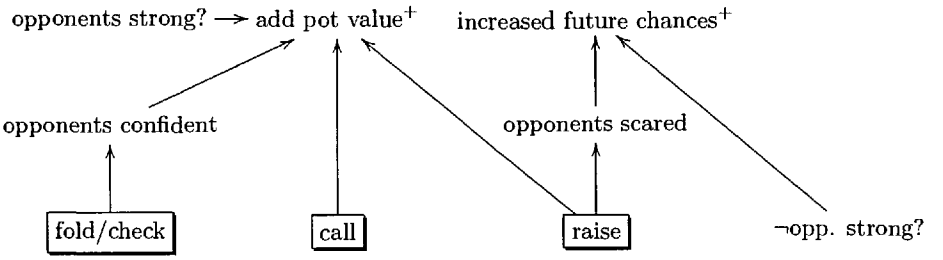


if fold/check then no_risk
if call then small_risk
if raise and good_hand then small_risk
if raise and ¬good_hand then big_risk

Figure 3.9: Poker basic influence diagram (part 2).

By refusing to take risk, one acts weakly and this makes the opponents feel quite confident. Confident opponents want to stay in the hand and usually want to bet more. As a result, the competition between the players increases. On the opposite, if one is willing to take risks, one acts strongly and this tends to scare opponents, chasing them sometimes away. This increases the chances to win the hand as fewer opponents remain, but this also does not help increasing the pot size. This can be modeled by adding to the basic influence diagram

another part (see Figure 3.10).



if fold/check then opponents_confident
if call then add_pot_value
if raise then add_pot_value
if raise then opponents_scared
if opponents_confident and opponents_strong then add_pot_value
if opponents_scared and ¬opponents_strong then increased_future_chances

Figure 3.10: Poker basic influence diagram (part 3).

Players also need to reason about their chances to win or lose. In poker, a player wins the pot either because all the others have folded or because he has the strongest hand amongst those that remain. All hands are ranked, from the strongest which is the straight flush (e.g. $A\spadesuit$-$K\spadesuit$-$Q\spadesuit$-$J\spadesuit$-$10\spadesuit$) to a pair (e.g. $J\clubsuit$-$J\diamondsuit$-X-X-X) and high cards (e.g. $Q\heartsuit$-X-X-X-X) as follows.

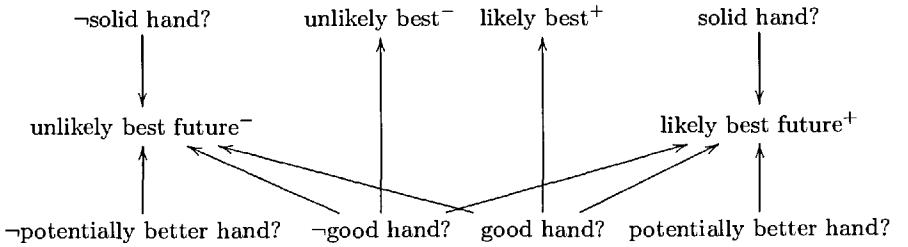
1. Straight Flush: five cards in sequence, all of the same suit.
2. Four of a Kind: four cards of one rank, and an unmatched card of another rank.
3. Full House: three matching cards of one rank, and two matching cards of another rank.
4. Flush: five cards of the same suit, not in rank sequence.
5. Straight: five cards of sequential rank but in more than one suit.
6. Three of a Kind: three cards of the same rank, plus two unmatched cards.

7. Two Pairs: two cards of the same rank, plus two cards of another rank (that match each other but not the first pair), plus one unmatched card.

8. One Pair: two cards of the same rank, plus three other unmatched cards.

9. High Cards: five cards in which no two cards have the same rank, the five cards are not in sequence, and the five cards are not all the same suit.

Players try to guess whether they currently hold the best hand or not, or if they eventually can have it as more community cards get revealed. For instance, a player with a pair of pocket tens and a ten on the flop (the three first community cards) is very likely to have the best hand and to have it until the end. But a player with only a pocket jack that has flopped a pair of jacks is in danger since another player may get a higher pair such as queens QQ when the fourth and fifth community cards are displayed. We model this in Figure 3.11.



if ¬good_hand then unlikely_best
if good_hand then likely_best
if ¬good_hand and ¬potentially_better_hand then unlikely_best_future
if ¬good_hand and potentially_better_hand then likely_best_future
if good_hand and solid_hand then likely_best_future
if good_hand and ¬solid_hand then unlikely_best_future

Figure 3.11: Poker basic influence diagram (part 4).

Our poker basic influence diagrams is composed of parts 1 to 4. More parts could be added and the current part could be extended if needed, for instance, to model the opponents individually and reason more precisely about hand

strength. Restricting ourselves to these four parts, we obtain the set of scenarios shown in Figure 3.12, where for example

$$w_1 = \{\text{good_hand}, \text{opponents_strong}, \neg\text{potentially_better_hand}, \text{solid_hand}\}$$

and a set of possible decisions initially expected $D = \{d_1, d_2, d_3\}$. The corresponding decision table is shown in Figure 3.13, where due to lack of space we have only shown the goals in each outcome and have employed the following abbreviations:

- NR stands for 'no risk',
- APV for 'added pot value',
- UB for 'unlikely the best',
- UBF for 'unlikely the best in the future',
- LB for 'likely the best',
- LBF for 'likely to be the best in the future',
- SR for 'small risk',
- BR for 'big risk' and
- IFC for 'increased future chances'.

Note that instead of having 16 columns, we only use 8. Columns which have identical outcomes can naturally be merged. This however requires using disjunctions of scenarios instead of single scenarios as column labels.

We believe that such a poker decision table could in the future serve as the basis for the development of a decision-theoretic poker program. The idea being that the program would gather information continuously from the game and

·	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
good hand	√	√	√	√	√	√	√	√
not good hand	×	×	×	×	×	×	×	×
opponents strong	√	√	√	√	×	×	×	×
not opponents strong	×	×	×	×	√	√	√	√
not potentially better hand	√	√	×	×	√	√	×	×
potentially better hand	×	×	√	√	×	×	√	√
solid hand	√	×	√	×	√	×	√	×
not solid hand	×	√	×	√	×	√	×	√

·	w_9	w_{10}	w_{11}	w_{12}	w_{13}	w_{14}	w_{15}	w_{16}
good hand	×	×	×	×	×	×	×	×
not good hand	√	√	√	√	√	√	√	√
opponents strong	√	√	√	√	×	×	×	×
not opponents strong	×	×	×	×	√	√	√	√
not potentially better hand	√	√	×	×	√	√	×	×
potentially better hand	×	×	√	√	×	×	√	√
solid hand	√	×	√	×	√	×	√	×
not solid hand	×	√	×	√	×	√	×	√

Figure 3.12: Scenarios for the poker decision table.

·	$w_1 \vee w_3$	$w_2 \vee w_4$	$w_5 \vee w_7$
d_1	NR^+, APV^+, LB^+, LBF^+	NR^+, APV^+, LB^+, UBF^-	NR^+, LB^+, LBF^+
d_2	SR^-, APV^+, LB^+, LBF^+	SR^-, APV^+, LB^+, UBF^-	SR^-, APV^+, LB^+, LBF^+
d_3	SR^-, APV^+, LB^+, LBF^+	SR^-, APV^+, LB^+, UBF^-	$SR^-, APV^+, IFC^+, LB^+, LBF^+$

·	$w_6 \vee w_8$	$w_9 \vee w_{10}$	$w_{11} \vee w_{12}$
d_1	NR^+, LB^+, UBF^-	NR^+, APV^+, UB^-, UBF^-	NR^+, APV^+, UB^-, LBF^+
d_2	SR^-, APV^+, LB^+, UBF^-	SR^-, APV^+, UB^-, UBF^-	SR^-, APV^+, UB^-, LBF^+
d_3	$SR^-, APV^+, IFC^+, LB^+, UBF^-$	BR^-, APV^+, UB^-, UBF^-	BR^-, APV^+, UB^-, LBF^+

·	$w_{13} \vee w_{14}$	$w_{15} \vee w_{16}$
d_1	NR^+, UB^-, UBF^-	NR^+, UB^-, LBF^+
d_2	SR^-, APV^+, UB^-, UBF^-	SR^-, APV^+, UB^-, LBF^+
d_3	$BR^-, APV^+, IFC^+, UB^-, UBF^-$	$BR^-, APV^+, IFC^+, UB^-, LBF^+$

Figure 3.13: Poker decision table.

use it to build a probability distribution of the scenarios (for instance using the method that will be presented in Chapter 5) and a value function for the outcomes (this will actually be done in Chapter 4), and then apply (a generalisation of) the expected utility criterion to make decisions from that table (this will be the issue addressed in Chapter 6). The decision table would be constructed first and then fixed during the game, so as not to be manipulated by disinformation being disseminated by the program's opponents.

3.8 Related work

An influence diagram [Howard and Matheson, 1981] is a directed acyclic graph containing decision nodes, uncertainty/deterministic nodes and a value node.

Decision nodes represent the possible decisions, uncertainty nodes represent any kind of uncertainty about logical or numeric variables and the value node models a von Neumann-Morgenstern utility function. The arcs are used to represent probabilistic, deterministic and functional dependencies between nodes. Influence diagrams can thus be used to completely model and solve decision making problems. In comparison, basic influence diagrams involve atomic decisions (which may be different from the possible decisions as discussed in Section 3.6), unobservable fundamental beliefs which are logical uncertainty nodes, deterministic nodes which are beliefs, and goals instead of the value node. Basic influence diagrams do not model the value function of the decision maker. They do not have probabilistic and functional arcs, they only have deterministic arcs modeling logical dependencies. Basic influence diagrams cannot be used to solve decision making problems, they can only be used to build decision tables. Note that influence diagrams do not allow to identify the scenarios and outcomes of decisions and that they are not equipped with any notion of rationality. Influence diagrams allow to make decisions without relying on any (explicit) decision table. Thus, basic influence diagrams are more specific, less informative and serve a different purpose.

Bayesian networks [Pearl, 1986] are a special kind of influence diagrams which are only used for representing and solving statistical inference problems. Bayesian networks are widely used for knowledge representation [Brachman and Levesque, 2004] in expert systems design. Bayesian networks model probabilistic dependencies, which is not the case with basic influence diagrams. Nevertheless, basic influence diagrams are used as a tool to represent (expert) knowledge in decision domains. Recently, influence diagrams have started to be used to structure argumentation-based systems for decision support [Morge and Mancarella, 2007] and our contribution is very much in line with this type of approach to

decision making (a comparison with the work of Morge and Mancarella can be found in Section 3.1).

Most of the research work on the construction of decision tables has so far focused on a simple yet extremely useful type of decision table called *decision logic tables* [Schmidt and Kavanagh, 1964; King, 1968]. These are defined as functional descriptions that map conditions (scenarios) to actions (decisions). In this approach, it is implicitly assumed that only one goal must be satisfied. In other words, the set of possible outcomes is restricted to $X = \{\text{Yes}, \text{No}\}$, the Yes being interpreted as 'the goal is achieved' and No as 'the goal is not achieved'. It becomes then possible to represent the decision table as a collection of *decision rules* of the form "if scenario w is true, then choose action d". A decision table is said to be: *complete* if, for every scenario, there is a rule that recommends a certain course of action to satisfy the goal; *redundant* when a rule appears more than once within the decision table; *ambiguous* when the conditions of any two rules are identical but the recommended decisions are different; *inconsistent* when the condition of any two rules are identical but the recommended decisions contradict each other; and to have a *logically impossible ambiguity* when there exists a decision rule whose conditions contradict each other (the condition of the rule is unsatisfiable). Note that inconsistency is a special form of ambiguity. [Verhelst, 1980] proposed a manual method for constructing decision tables based on the following stages: 1) obtain conditions, condition states and actions of the decision situation, 2) specify the problem in terms of decision rules, 3) fill the decision table based on the rules, 4) check for completeness, redundancy, ambiguity, etc. and 5) simplify the decision table and display it. Our method proceeds in a very similar fashion, starting with the elicitation of beliefs and decisions, the description of a decision situation in terms of a diagram and if-then rules, the construction of the decision table

via computation of the liberal stable semantics and finishing with the potential simplification of the table by merging identical columns (as done in the poker Section 3.7). The differences are that steps 1) and 2) can be done using a basic influence diagram as a graphical aid, step 3) is completed automatically by computation of the liberal stable semantics and step 4) is obsolete, as the outcomes are guaranteed to be rational under extensibility.

[Vanthienen and Dries, 1993] promoted the idea of automating the construction process and developed the PROLOGA software. We have implemented our algorithms for computing decision tables in Maple version 9.5. [Vanthienen and Robben, 1993] developed applications e.g. for the verification and visualisation of legal procedures, checking consistency in medical treatments, help desk applications for computer networks, validation of knowledge-based systems, rates and premiums in banks and insurance companies. Many other applications of decision logic tables exist, from management data processing, medicine, quality control to engineering problems [Moore and Cullinane, 1969]. [Milner, 1977] proposed for example decision tables as a form of tabular diagnostics for discovering reasons for breakdowns of manufacturing systems. Generally speaking, decision tables can be used as a way to define and document the logic of decision making. Thus, our poker decision table could also be used as tool for teaching poker.

An important number of algorithms have been developed to make decision logic tables complete, non redundant, non ambiguous, consistent and to detect logical errors [King, 1968; Milner, 1977; Shwayder, 1975; Lew and Tamanaha, 1976; Ibramsha and Rajaraman, 1978; Lew, 1984]. The type of decision tables we have considered in this chapter is different and more general than the one on which decision table theory focuses, as a) we consider more than one goal to be achieved and b) we also consider goals that need to be avoided (such as being

loaded when going out for a walk in the umbrella problem). The decision tables we construct could at best be used for deriving *practical arguments*, i.e. heuristic arguments either in favour or against decisions, but these arguments still would not have the force of decision rules (i.e. lead with certainty to the satisfaction of all positive goals and avoidance of all negative goals). The decision tables we construct are thus neither complete (there may not exist an ideal decision under every scenario), nor unambiguous (more than one decision may be ideal or optimal under a given scenario) or consistent (two mutually exclusive decisions may perfectly be optimal under the same scenario). The notions of redundancy and logical impossible ambiguity do not even apply to our approach.

[Shwayder, 1975] studies methods for reducing the number of decision rules used to represent decision logic tables. The basic intuition is that any two decision rules recommending the same decision under different conditions can be combined into one rule recommending that decision under the disjunction of these conditions. In the case where outcomes are described by more than one goal, such direct simplifications are not possible anymore. However, when two columns in a decision table are equal, the columns can be merged into one only by preserving the outcomes and taking the disjunction of the conditions or union of the associated events, as we have done in the decision table for poker in Section 3.7.

Since decision logic tables are based upon decision rules of the form 'if condition then action', [Reilly *et al.*, 1987] has advocated logic programming techniques [Kowalski, 1979b; 1979a] to define and implement them. The resulting decision tables can then be easily queried using Prolog, allowing the user to ask if a certain goal is achievable and if so to know by which course of action, under given specific circumstances. We have also chosen to rely upon if-then rules to specify decision tables and have advocated the use of a generalisation of logic pro-

gramming, namely assumption-based argumentation [Bondarenko *et al.*, 1997; Dung *et al.*, 2006]. The rules considered are however of a different type: 'if condition and optionally action then condition or goal'. We query the assumption-based argumentation frameworks derived from basic influence diagrams using the liberal stable semantics to rationally identify all the possible outcomes of decisions (Section 3.4).

3.9 Summary and conclusion

We have proposed *basic influence diagrams* as a tool to describe qualitatively practical situations of decision making under strict uncertainty. These diagrams give a logical structure to the decision domain and reveal the decision maker's most fundamental uncertainties. Basic influence diagrams can be analysed and their rational outcomes be identified using an *extension of assumption-based argumentation.* In the argumentation framework derived from a diagram and under a quite natural hypothesis called *extensibility*, we have proved that the possible outcomes of decisions are in one-to-one correspondence with the consequences of *liberal stable sets of assumptions.* We have studied the relationship between the liberal stable semantics and existing notions of acceptability in assumption-based argumentation. The liberal stable semantics can be computed using and adapting existing hypergraph algorithms. Computing the liberal stable semantics directly allows to derive the scenarios, possible decisions, rational outcomes and compute decision tables. These tables are by construction rational, simple to understand and fit naturally in applications of decision theory.

Chapter 4

Argument strength & value functions

The objective in this chapter is to investigate the practical construction of value functions in decision theory [Fishburn, 1964; Raiffa, 1968; French, 1987]. Value functions serve to quantify the desirability of the outcomes of decisions from the point of view of a decision maker. In the previous chapter, we have shown how to use basic influence diagrams, extended assumption-based argumentation frameworks and the liberal stable semantics to identify the possible outcomes of decisions. In this chapter, we use abstract argumentation for constructing value functions.

The chapter is organised as follows. Section 4.1 introduces the problem of constructing value functions and motivates the use of argumentation to solve this problem. Section 4.2 presents an abstract argumentation model for representing the uncertain knowledge of a decision maker concerning the desirability of his objectives. Section 4.3 introduces a new measure of argument strength suitable for abstract argumentation. This measure is of general purpose and will

thus also be used in the next chapter for the construction of imprecise probabilities. In Section 4.4, we introduce a mechanism for aggregating arguments and constructing value functions from arguments based on linear social welfare theory [Fleming, 1952; Harsanyi, 1955; Dyer and Sarin, 1978]. In Section 4.5, we apply this method to generate value functions for poker programs having various psychological and emotional features. We discuss related work in Section 4.6 and conclude in Section 4.7.

The argument strength measure presented in Section 4.3 has been published in the proceedings of the Eleventh European Conference on Logics in Artificial Intelligence [Matt and Toni, 2008b].

4.1 On the construction of value functions

As a result of Chapter 3, we may now assume given

- a set of possible decisions $D = \{d_1, \ldots, d_m\}$,
- a set of scenarios $\Omega = \{w_1, \ldots, w_n\}$, and
- a mapping $T : D \times \Omega \rightarrow X$ from decisions and scenarios to outcomes.

We denote by $x_{i,j}$ the outcome of decision d_i under scenario w_j and X the set of all such outcomes. A *value function* is a real-valued function $V : X \rightarrow \mathbb{R}$ such that, for every outcome $x \in X$, the value $V(x)$ represents the desirability of x from the decision maker's viewpoint. We have seen in Chapter 2 that value functions are necessary for decision making both in situations of strict uncertainty and risk.

We are now concerned with the issue of measuring the values $V(x)$, for all $x \in X$. Existing measures in *value theory* [Fishburn, 1964] include the partial ordinal measure, the ordinal measure, the first, second and third ordered metric measures, the first, second and third bounded interval measures and the

interval measure. The partial ordinal measure for instance denotes the class of value functions that are compatible with a given partial order [1] on X, the ordinal measure denotes the class of value functions that are compatible with a total order [2]. In fact, all of these measures are derived from statements of preference concerning the elements of X that need to be elicited from the decision maker. Such preference statements are amenable to relations of the form $V(x) \geq V(y)$, $V(a)-V(b) \geq x(V(c)-V(d))$ or $V(a) = pV(x)+(1-p)V(y)$, where x, y, a, b, c and d are outcomes in X, $x \in \mathbb{R}$ and $p \in [0,1]$. The decision maker's preferences are thus usually modeled as a system of linear inequalities. Existence and possibly uniqueness of solutions of these systems can be usually guaranteed under conditions concerning the rationality of the decision maker, e.g. the von Neumann-Morgenstern postulates [von Neumann and Morgenstern, 1944] (see utility functions in Chapter 2). Sometimes, solutions do not exist and one then usually tries to find a value function V that best satisfies the system (minimising the error by the least squares method).

[Fishburn, 1964] mentioned a number of important issues that may arise when applying these methods: a) the decision maker may lack rationality and provide conflicting statements, b) he may be uncertain about his objectives, c) his preferences may change over time, d) he might be dishonest, e) the number of outcomes can be very large, making the elicitation process quite lengthy and f) the decision maker may be unable to give preference statements concerning complex outcomes. To address issues a) and b), we envisage to incorporate argumentation as in [Dung, 1995] in the measurement method, as argumentation indeed allows to represent uncertain and conflicting knowledge. In this chapter, again, arguments will be treated as justified claims [Krause *et al.*, 1995]. To address issue b), we assume given a set of objectives (positive and negative)

[1] A partial order is a binary relation that is reflexive, transitive and antisymmetric.

[2] A total order is a partial order that is total, i.e. such that every pair of elements of the set can be compared.

goals G and assume that each outcome is a set of goals. These assumptions are compatible with the use of basic influence diagrams in Chapter 3. Indeed, the goals of an influence diagram are the only literals that have an influence on the overall satisfaction of the decision maker. This is the reason why we may restrict our attention to goals. We also restrict the scope of arguments to single goals (instead of outcomes in the powerset of G). By doing so, we resolve issues e) and f). By forcing the decision maker to use arguments justifying his claims, we understandably also put a limit to dishonesty, thus partially resolving issue d). Finally, although issue c) cannot really be resolved, we will show that, by representing the decision maker's opinion in the simple form of an abstract argumentation framework [Dung, 1995], it will be quite easy for the decision maker to revise his beliefs (modeled as arguments) and update his value function over time. In a nutshell, we believe that argumentation can bring practical value and add substantive goodness to the measurement process.

4.2 Modeling opinions about goals by abstract argumentation

In the previous chapter, we have seen how to compute outcomes and, in particular, the positive or negative goals satisfied in each one of them. Goals in the terminology of basic influence diagrams (Chapter 3) are the literals which directly affect the decision maker's satisfaction, so that the value $V(x)$ of an outcome x should logically only depend on the set of goals satisfied in x. We thus restrict our attention to value functions defined over sets of goals. We assume that G is the set of all literals of goal nodes from a basic influence diagram. The set of positive goals (corresponding to nodes labelled with the symbol $+$) is denoted G^+ and the set of negative goals (corresponding to nodes labelled

labelled with the symbol -) is denoted G^-. We define [3]

Definition 4.1. *Every* ***outcome*** $x \in X$ *is a subset of* G*:* $x \subseteq G$.

For instance in the umbrella problem from Chapter 3, the set of goals is

$$G = \{\text{dry}, \neg\text{dry}, \text{loaded}, \neg\text{loaded}\}$$

and the set of outcomes is

$$X = \{\{\text{dry}, \text{loaded}\}, \{\text{dry}, \neg\text{loaded}\}, \{\neg\text{dry}, \neg\text{loaded}\}\}$$

Note that not all subsets of G are outcomes. For instance in the previous example, $\{\}$ and $\{\neg\text{dry}, \text{loaded}\}$ are not outcomes.

In order to ease and ground practically the process of reasoning about the desirability of goals, the decision maker is asked to specify a discrete *scale of desirability*, defined as

Definition 4.2. *A* ***scale of desirability*** *is a set of the form*

$$S = \{v_1, v_2, \ldots, v_N\}$$

where $N \geq 1$, $v_1, \ldots, v_N \in \mathbb{R}$ *(or simply* $\mathbb{Q}$ *or* $\mathbb{Z}$*) and* $v_1 < v_2 < \ldots < v_N$.

N corresponds to the number of levels of desirability needed to describe the goals in G. By convention, positive levels of desirability must correspond to desirable goals in G^+ and negative levels to undesirable goals in G^-. Each v_i is given an interpretation in natural language, such as slightly, fairly, mildly, quite and extremely desirable/undesirable and the absolute value of v_i must be set by the decision maker so as to reflect the intensity in desirability/undesirability.

[3] In Chapter 3, outcomes of decision tables are subsets of $\mathcal{L}$. The outcomes in this chapter are simply the intersections of the outcomes of the decision table with the set of goals G.

Thus, the scale must be such that if $v_i < v_j$, then v_i is interpreted as less desirable than v_j. For example, we may adopt a basic scale of desirability $S = \{v_1, v_2, v_3, v_4\}$ with values in $\mathbb{Q}$, where

- $v_1 = -1$ is interpreted as quite undesirable,
- $v_2 = -0.25$ is interpreted as fairly undesirable,
- $v_3 = +0.25$ is interpreted as fairly desirable, and
- $v_2 = +1$ is interpreted as quite desirable.

As we will see in definition 4.3 below, we will use positive elements of the scale S to assess the desirability of positive goals in G^+ and negative elements to assess the desirability of negative goals in G^-.

Once the scale is in place, the decision maker is invited to formulate a collection of arguments. Every argument is interpreted as a justified claim. We assume given some language L (set of literals closed under negation $\neg$) suitable for expressing justifications for arguments. To build a value function on X, we may restrict ourselves to only two types of claims, leading us to distinguish between two categories of arguments, viz. *value and mitigation arguments.* Value arguments serve to express the desirability of goals, and mitigation arguments to express doubts concerning the justification of value arguments.

Definition 4.3. *A **value argument** a is a pair (J_a, C_a), where*

- *C_a the claim in the form of a pair $(g_a, v_a) \in G \times S$ stating that goal g_a has desirability v_a and which verifies $v_a \geq 0$ if $g_a \in G^+$ and $v_a < 0$ if $g_a \in G^-$.*
- *$J_a \in L$ is a justification for the claim C_a.*

Definition 4.4. *A **mitigation argument** m is a pair (J_m, C_m), where*

- *C_m is a claim of the form $\neg J$, where J is the justification of some value or mitigation argument.*

- *$J_m \in L$ is a justification for the claim C_m.*

Intuitively, given a mitigation argument (J_m, C_m), $C_m = \neg J$ states that the justification J is invalid or erroneous.

The decision maker must provide a set A of value arguments and a set M of mitigation arguments. In the umbrella problem, the set of value arguments A could be $\{a_1, a_2, a_3, a_4\}$ as defined below

- $a_1 = (\text{ill}, (\text{dry}, +1))$: "it is quite desirable to stay dry during the walk because I am ill"

- $a_2 = (\text{ill}, (\neg\text{dry}, -1))$: "it is quite undesirable to be wet during the walk because I am ill"

- $a_3 = (\text{tired}, (\text{loaded}, -0.25))$: "it is fairly undesirable to be loaded during the walk because I am going to be tired"

- $a_4 = (\text{tired}, (\neg\text{loaded}, 0.25))$: "it is fairly desirable not to be loaded during the walk because I am going to be tired"

The set of mitigation arguments M could be $\{m_1, m_2, m_3\}$ as defined below

- $m_1 = (\text{pills}, \neg\text{ill})$: "I won't be ill anymore during the walk because I have taken pills to recover from my flu"

- $m_2 = (\text{bad_flu}, \neg\text{pills})$: "the pills I have taken will be ineffective because I have a really bad flu"

- $m_3 = (\text{breaks}, \neg\text{tired})$: "I won't be tired during the walk because I am planning to have several breaks"

Given any sets of value and mitigation arguments A and M, we can construct an abstract argumentation framework [Dung, 1995] (see Chapter 2) as follows

Definition 4.5. *$F = (Arg, att)$, where $Arg = A \cup M$ and*

$$att = \{(x, y) \mid x = (J_x, \neg J_y) \in M, y = (J_y, C_y) \in A \cup M\}$$

The abstract argumentation framework corresponding to our running example is shown in Figure 4.1.

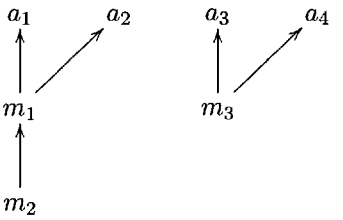


Figure 4.1: Value and mitigation arguments for the umbrella problem.

The framework is constructed so that mitigation arguments attack value arguments and mitigation arguments can attack one another or value arguments. However, value arguments do not attack one another. This means that we use attacks to undermine arguments but do not have rebuttals. Indeed, the framework is an abstract model of the information available to the decision maker. Value arguments represent information that directly relates to the value of his objectives, and mitigation arguments capture the decision maker's uncertainties about his objectives. Attacks by mitigation arguments against value arguments cast doubt on the validity of value arguments. In the next section, we explain how to precisely obtain an estimation of the strength of arguments in any abstract argumentation framework. Then, in Section 4.4, we will use this notion of strength on value arguments to determine value functions.

4.3 Strength of arguments in abstract argumentation frameworks

Several notions of argument strength have been proposed in the literature on argumentation. One distinguishes between the so-called *intrinsic* and *interaction-based* measures. The term "intrinsic" is used to refer to approaches such as [Pollock, 1992; Krause *et al.*, 1995; Ambler, 1996; Parsons, 1997; Prakken and Sartor, 1997; Amgoud and Cayrol, 1998; Pollock, 2001; Kohlas *et al.*, 2002] whereby the strength of an argument is independent of its interactions with other arguments, i.e. the arguments that attack or are attacked by it. Those approaches have been discussed in Chapter 2. On the other hand, the term "interaction-based" refers to measures whereby the strength of an argument depends on the arguments attacking it (the attackers), the attackers of its attackers (the defenders), etc. Amongst interaction-based measures, one can distinguish between qualitative [Dung, 1995; Jakobovits and Vermeir, 1999] and quantitative [Besnard and Hunter, 2000; Cayrol and Lagasquie-Schiex, 2005; Matt and Toni, 2008b] measures. We refer to interaction-based measures that are quantitative as *dialectical measures* and define a new dialectical measure exhibiting a number of intuitively appealing properties.

Our dialectical measure is defined for any given abstract argumentation framework $F = (Arg, att)$ and some fixed argument $x \in Arg$, as the value of a two-person, zero-sum game of strategy with imperfect information called (F, x) *game of argumentation strategy* and played according to rules described in the next subsection. Games of strategy were first introduced by [Borel, 1921]. It was then von Neumann who defined the concept of value of a game through his famous *minimax theorem* [von Neumann, 1928]. We first give the rules of a game of argumentation strategy (Section 4.3.1). We then explain how the

value of this game and the strength of an argument are defined (Section 4.3.2) and computed (Section 4.3.3). Finally, we study the properties of this measure (Section 4.3.4).

4.3.1 Rules of the (F, x) game of argumentation strategy

The game involves two players called *proponent* and *opponent* of x. The argument x as well as the set of arguments Arg and the attack relationship att between arguments are revealed to both players. The role of the proponent is to form an opinion P to defend x. The opponent must form an opinion O to defeat the proponent. We assume that P and O are sets of arguments. We call in this chapter [4]

Definition 4.6. *An **opinion** is a subset of Arg.*

The players must choose their opinion secretly. The proponent of x is forced to choose an opinion embracing x, which means that P must contain at least x.

Definition 4.7. *The **proponent's strategy** is an opinion P that contains x and the **opponent's strategy** is any opinion.*

The players then *simultaneously* reveal their opinions. A certain amount of money called indifferently payoff or reward and denoted $\phi(P, O)$ is then paid by the opponent to the proponent. This amount depends on the attacks existing between the arguments contained in P and O. Thus, for any sets of arguments A and B, we write $B_F^{\leftarrow A}$ to denote the set of attacks from A against B in the abstract argumentation framework F:

$$B_F^{\leftarrow A} = \{(a, b) \mid a \in A, b \in B, (a, b) \in att\}$$

[4] This is strictly speaking an abuse of notation, as the term opinion was used in Chapter 3 to denote sets of assumptions.

We also denote $|B_F^{\leftarrow A}|$ the number of attacks from A against B. To calculate the payoff, we use the function $f : \mathbb{N} \rightarrow [0,1]$ defined $\forall n \in \mathbb{N}$ as

$$f(n) = \frac{n}{n+1}$$

This function f has the advantage of being monotonic increasing and to range from 0 to 1 as the number n of attacks considered ranges from $n = 0$ to infinity. As a result, the following notion of reward (and subsequently of strength) remains bounded between 0 and 1, no matter how many attacks exist between the strategies of the two players.

Definition 4.8. *The **proponent's reward/payoff** $\phi(P, O)$ is determined as follows:*

- *if P is not a conflict-free extension in F, then $\phi(P, O) = 0$*
- *if P is conflict-free and O does not attack P in F, then $\phi(P, O) = 1$*
- *otherwise*

$$\phi(P, O) = \frac{1}{2}\left[1 + f\left(|O_F^{\leftarrow P}|\right) - f\left(|P_F^{\leftarrow O}|\right)\right]$$

The function ϕ chosen is the only affine function in the variables $f\left(|O_F^{\leftarrow P}|\right)$ and $f\left(|P_F^{\leftarrow O}|\right)$ which is equal to

- 1 when $f\left(|O_F^{\leftarrow P}|\right) = 1$ and $f\left(|P_F^{\leftarrow O}|\right) = 0$
- 1/2 when $f\left(|O_F^{\leftarrow P}|\right) = f\left(|P_F^{\leftarrow O}|\right)$
- 0 when $f\left(|O_F^{\leftarrow P}|\right) = 0$ and $f\left(|P_F^{\leftarrow O}|\right) = 1$

In the abstract argumentation framework F used for the umbrella problem (see Section 4.2) and the argument $x = a_1$, the proponent may choose the strategy $P = \{a_1, m_3\}$ and the opponent the strategy $O = \{m_1\}$. P is conflict-free and O attacks P. There is one attack (m_3, m_1) from P to O and one attack

(m_1, a_1) from O to P, so the reward for the proponent is

$$\phi(P, O) = \frac{1}{2}[1 + f(1) - f(1)] = \frac{1}{2}$$

Note that the rules and payoff of the (F, x) game are set so as to drive the players according to three natural principles of rationality in dialectic:

1. to properly defend the argument x, its proponent should form an opinion embracing x and avoid self-contradiction,

2. the role of the opponent is to form an opinion which attacks the proponent's opinion, and

3. in a dispute, it is better for each party to maximise the number of attacks against his adversary and to minimise the number of attacks from him, for the simple reason that opinions are weakened by attacks.

4.3.2 Value of the (F, x) game of argumentation strategy

Let $\{P_1, ..., P_m\}$ denote the set of possible strategies of the proponent (all the subsets of Arg containing x) and $\{O_1, ..., O_n\}$ the set of possible strategies of the opponent (all the subsets of Arg). Let us denote $r_{i,j} = \phi(P_i, O_j)$ the payoff when the proponent plays with the strategy P_i and the opponent with O_j, as given in definition 4.8. Assume that, each time they play the game (a round), the players choose their strategy randomly according to some probability distributions X and Y fixed in time, where $X = (x_1, \ldots, x_m)$ is the vector whose component i corresponds to the probability x_i for the proponent of playing P_i and $Y = (y_1, \ldots, y_n)$ is the vector whose component j corresponds to the probability y_j for the opponent of playing O_j. X and Y are called *mixed strategies*. The proponent's expected payoff is equal to the opponent's expected loss (this is why the game of argumentation strategy is said to be a *zero-sum* game) and,

according to the definition of expectation in probability theory and given the probabilistic independence of the strategies chosen by the players, is equal to

$$E = \sum_{i=1}^{m} \sum_{j=1}^{n} x_i r_{i,j} y_j$$

Using the vectors $X = (x_i)$ and $Y = (y_j)$ and matrix $R = ((r_{i,j}))$ we can write more compactly that $E = X'RY$, where X' is the transpose of X. If the proponent randomises his strategies according to X, he can expect at least $\min_Y X'RY$. The optimal X will make his expectation at least greater or equal to $\max_X \min_Y X'RY$. Similarly, the opponent can reduce the expectation of his loss to at most $\min_Y \max_X X'RY$. The minimax theorem [von Neumann, 1928] states that

Theorem 4.1. *There exists a unique value $v \in \mathbb{R}$ such that*

$$\max_X \min_Y X'RY = \min_Y \max_X X'RY = v$$

v is called the *value* of the game. This value can be interpreted both as the expected payoff that is guaranteed to the proponent and as the maximal expected loss of the opponent. We adopt

Definition 4.9. *The **strength** $s_F(x)$ of argument x in F is the value of the (F, x) game of argumentation strategy.*

4.3.3 Computing the value of the (F, x) game of argumentation strategy and the strength of x

The technique used for computing the value of two-person zero-sum games with imperfect information is described in textbooks on operations research [Hillier and Lieberman, 1995] and is based upon linear optimisation. Linear optimisa-

tion problems can be solved using the simplex algorithm [Dantzig *et al.*, 1955] or more efficient polynomial time methods [Khachiyan, 1979; Karmarkar, 1984]. Using this technique, the strength $s_F(x)$ of x in the abstract argumentation framework F is the maximal value that the variable x_{m+1} can take when it is subject to the following linear inequality constraints:

$$\forall j \in \{1, \ldots, n\} : \sum_{i=1}^{m} r_{i,j}\, x_i - x_{m+1} \geq 0$$

$$\sum_{i=1}^{m} x_i = 1$$

$$x_1, \ldots, x_m \geq 0$$

A mixed strategy $X = (x_1, \ldots, x_m)$ for which x_{m+1} is maximal is referred to as *optimal mixed strategy.*

As an illustration, let us examine in detail the case of $F = (Arg, att)$ where $Arg = \{a, b\}$, $att = \{(a, b)\}$ and $x = b$. The proponent strategies are $\{P_1, P_2\}$ with $P_1 = \{b\}$ and $P_2 = \{a, b\}$. The opponent strategies are $\{O_1, O_2, O_3, O_4\}$ with $O_1 = \emptyset$, $O_2 = \{b\}$, $O_3 = \{a, b\}$ and $O_4 = \{a\}$. The proponent's payoff matrix is

$$R = \begin{pmatrix} 1 & 1 & 0.25 & 0.25 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

and the optimisation problem to solve to compute $s_F(b)$ is to maximise x_3

subject to the constraints

$$\begin{cases} x_1 - x_3 \geq 0 \\ x_1 - x_3 \geq 0 \\ 0.25x_1 - x_3 \geq 0 \\ 0.25x_1 - x_3 \geq 0 \\ x_1 + x_2 = 1 \\ x_1, x_2 \geq 0 \end{cases}$$

which can be simplified as the problem of maximising x_3 subject to the constraints

$$\begin{cases} 0.25x_1 \geq x_3 \\ x_1 + x_2 = 1 \\ x_1, x_2 \geq 0 \end{cases}$$

with obvious solution $x_3 = 0.25$ obtained for the mixed strategy $X = (x_1, x_2) = (1, 0)$. Thus $s_F(b) = 0.25$. X is an optimal mixed strategy, thus the proponent of b should always play $P_1 = \{b\}$ and never play $P_2 = \{a, b\}$. Indeed, $\{a, b\}$ is not a conflict-free opinion and always leads to a null reward.

Other examples of argument strength in elementary argumentation frameworks are provided in Table 4.1. In order to obtain the values in this table, for each one of these frameworks F and each argument x in F, we have constructed the (F, x) game payoff matrix $R = ((r_{i,j}))$ where $r_{i,j} = \phi(P_i, O_j)$ and computed the game's value using the simplex algorithm. Figure 4.2 represents graphically the results obtained with the last framework of the table. The proponent's optimal mixed strategies corresponding to these values are omitted for simplicity, but can be summarised in short as follows:

- playing $\{a, f\}$ with probability 1 for argument a (i.e. all the elements x_i^* of the optimal mixed strategy $X^* = (x_i^*)$ are equal to zero except for the index i such that $P_i = \{a, f\}$ for which $x_i^* = 1$),
- $\{b\}$ with probability 1 for argument b,
- $\{a, c, f\}$ with probability $\frac{2}{3}$ and $\{a, c\}$ with probability $\frac{1}{3}$ for argument c,
- $\{b, d\}$ with probability 1 for argument d,
- $\{b, e\}$ with probability 1 for argument e and
- $\{f\}$ with probability 1 for argument f.

The qualitative ranking obtained is in decreasing order of strength

$$s_F(f) > s_F(a) = s_F(b) = s_F(d) > s_F(c) > s_F(e)$$

This ranking is different from the one obtained using the dialectical measure by [Besnard and Hunter, 2000] and [Cayrol and Lagasquie-Schiex, 2005], where

$$s_F(f) > s_F(d) > s_F(a) = s_F(b) > s_F(e) > s_F(c)$$

In the next section, we will identify properties of our notion of strength with respect to qualitative interaction-based measures of acceptability as in Dung [Dung, 1995]. The dialectical measures of [Besnard and Hunter, 2000] and [Cayrol and Lagasquie-Schiex, 2005] do not exhibit these properties, as, in the earlier example they rank argument e higher than argument c despite the fact that c is contained in the stable extension $\{a, c, f\}$ and e is not contained in any admissible extension of F.

arguments	attacks	strength values for arguments given in alphabetical order
{a}	{}	1
{a, b}	{(a,b)}	1, 0.25
{a, b}	{(a,b), (b,a)}	0.5, 0.5
{a,..., d}	{(a,b), (c,b), (d,c)}	1, 0.25, 0.25, 1
{a,..., d}	{(a,b), (c,b), (d,b), (b,d)}	1, 0.167, 1, 0.625
{a,..., e}	{(a,b), (c,b), (d,a), (e,c)}	0.25, 0.5, 0.25, 1, 1
{a,..., f}	{(a,b), (b,c), (c,d), (d,e), (e,f)}	1, 0.25, 0.5, 0.386, 0.5, 0.425
{a,..., f}	{(a,b), (b,a), (b,c), (c,d), (e,c), (f,e)}	0.5, 0.5, 0.417, 0.5, 0.25, 1

Table 4.1: Examples of argument strength.

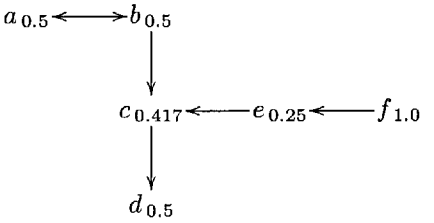


Figure 4.2: Argumentation framework with strength labels.

4.3.4 Properties of the argument strength measure

We consider three kinds of properties of our dialectical measure of argument strength (definition 4.9):

1. elementary properties which directly derive from the rules of the game of argumentation strategy and the minimax theorem 4.1,

2. properties related to some notions of acceptability as in abstract argumentation, and

3. dynamic properties, when the given abstract argumentation framework changes.

While presenting these properties we take as given an abstract argumentation framework $F = (Arg, att)$ and an argument $x \in Arg$.

Elementary properties

The strength of an argument is comprised between zero and one.

Property 4.1. $0 \leq s_F(x) \leq 1$

Proof. Trivially, by definition 4.8, $\forall (i,j), r_{i,j} = \phi(P_i, O_j) \in [0,1]$. Thus, for every mixed strategies X and Y, since X and Y are probability distributions, we also have $X'RY \in [0,1]$, which implies $0 \leq \min_Y X'RY$ and $\max_X X'RY \leq 1$. Therefore, we have

$$0 \leq \max_X \min_Y X'RY \quad \text{and} \quad \min_Y \max_X X'RY \leq 1$$

By the minimax theorem 4.1, $0 \leq v \leq 1$, and thus $v = s_F(x) \in [0,1]$. □

The strength is minimal and null when the argument attacks itself.

Property 4.2. $s_F(x) = 0$ *if and only if* x *attacks itself.*

Proof. $\Rightarrow$: $s_F(x) = v = \min_Y \max_X X'RY = 0$ implies that $\exists Y^* = (y_j^*)$ such that $\forall X$, $X'RY^* \leq 0$. This holds notably for any $X = e_i$ (the vector whose components are all equal to 0 except the ith one which is equal to 1). Hence, $\forall i$, $\sum_j r_{i,j} y_j^* \leq 0$. Since $r_{i,j} y_j^* \geq 0$, it is clear that $\forall (i,j)$, $r_{i,j} y_j^* = 0$. Y^* is a probability distribution, so there exists k such that $y_{j_k}^* > 0$. It is then necessary that, $\forall i$, $r_{i,j_k} = 0$. Since trivially, by definition 4.8, $\phi(P,O) = 0$ if and only if P is not conflict-free, then $\forall i$, P_i attacks itself. In particular, $P_i = \{x\}$ attacks itself, which means that x attacks itself.

$\Leftarrow$: If x attacks itself, then all proponent strategies in the (F,x) game are non-conflict-free sets of arguments. Since P attacks itself implies $\phi(P,O) = 0$ for any O, we have $R = ((0))$, $v = 0$ and $s_F(x) = 0$. □

The strength of an argument is maximal and equal to one when the argument is not attacked.

Property 4.3. *$s_F(x) = 1$ if and only if there exists no $y \in Arg$ such that $(y, x) \in att$.*

Proof. $\Rightarrow$: If $s_F(x) = v = 1$, then we have $\max_X \min_Y X'RY = 1$. Y ranges over the set of all real-valued probability distributions which is larger than the set $\mathcal{Z}$ of all zero-one valued probability distributions. Thus, $\forall X$,

$$\min_{Y \in \mathcal{Z}} X'RY \geq \min_{Y} X'RY$$

Therefore, $\max_X \min_{Y \in \mathcal{Z}} X'RY \geq \max_X \min_Y X'RY = 1$. This can be rewritten as $\max_X \min_j \sum_i r_{i,j} x_i \geq 1$. $\exists X^* = (x_i^*)$ such that $\min_j \sum_i r_{i,j} x_i^* \geq 1$, i.e., $\forall j$, $\sum_i r_{i,j} x_i^* \geq 1$. Since, $\forall (i,j)$, $r_{i,j} \leq 1$ (trivially by definition 4.8) and X^* is a probability distribution, then, $\forall j$, $\sum_i r_{i,j} x_i^* \leq 1$, so that in fact, $\forall j$, $\sum_i r_{i,j} x_i^* = 1$. This may only hold if, $\forall (i,j)$, $r_{i,j} < 1 \Rightarrow x_i^* = 0$. X^* is a probability distribution, so there exists k such that $x_k^* > 0$. By contraposition of the previous implications, $\forall j$, $\neg(r_{k,j} < 1)$, i.e. $r_{k,j} \geq 1$. Since ϕ is bounded up by 1 by definition 4.8, then $\forall j$, $r_{k,j} = 1$. Since $\phi(P, O) = 1$ if and only if P is conflict-free and O does not attack P (by definition 4.8), P_k is conflict-free and $\forall j$, O_j does not attack P_k. $x \in P_k$ by definition 4.7 of a proponent's strategy, so there is no opponent strategy, and thus argument, that attacks x.

$\Leftarrow$: By selecting strategy $\{x\}$ with probability 1, then by definition 4.8, the proponent has a guaranteed payoff of 1 irrespective of what the opponent does. Therefore, $s_F(x) = v \geq 1$. by property 4.1, $s_F(x) = 1$. □

Properties related to the acceptability of the argument

Here, we relate the strength of an argument to the acceptability of the extensions that contain the argument, where an extension is a set of arguments. Such extensions can be used by the proponent as strategies to defend the argument. We refer the reader to Chapter 2 for the definitions of the notions of acceptability

for abstract argumentation that we use in the remainder of this section. Here, we denote by $n(x)$ the number of attacks against x and by k the *out-degree* of F, defined as

Definition 4.10. *The* ***out-degree*** *of F is the number k of maximal out-going attacks of an argument in F, i.e.*

$$k = \max_{a \in Arg} |\{b | (a, b) \in att\}|$$

For instance, in the framework of Figure 4.2, the out-degree is $k = 2$, since b attacks a and c and all the other arguments attack at most one argument.

When an argument x does not attack itself, its strength is bounded below by a value that is inversely proportional to $n(x)$, which means that x may have a strength closer to zero as $n(x)$ becomes large. Its strength is also bounded above by a decreasing function of $n(x)$ which takes the value 1 for $n(x) = 0$ and converges towards $1/2$ as $n(x) \to +\infty$. Thus, the strength of x may still be close to $1/2$ even if $n(x)$ is large. This is formalised in the following

Property 4.4. *If x is contained in a conflict-free extension, then*

$$\frac{n(x)+2}{2(n(x)+1)} \geq s_F(x) \geq \frac{1}{2(n(x)+1)}.$$

Proof. If there exist $n(x)$ attacks against x, then there exists an opponent strategy O with $n(x)$ arguments attacking x. For this strategy O and any proponent strategy P, there exists at least $n(x)$ attacks from O against P. Thus, denoting by $f(\infty) = \lim_{n \to +\infty} f(n) = 1$, we have

$$\phi(P, O) < \frac{1}{2}[1 + f(\infty) - f(n(x))] = 1 - \frac{1}{2} f(n(x))$$

Thus, by always playing O, the opponent can secure a strict maximum loss of

$$1 - \frac{1}{2}f(n(x)) = \frac{n(x)+2}{2(n(x)+1)}$$

On the other hand, by always playing with the strategy $P = \{x\}$ which is conflict-free, the proponent can secure a minimum reward of

$$\frac{1}{2}[1 + f(0) - f(n(x))] = \frac{1}{2}[1 - \frac{n(x)}{n(x)+1}] = \frac{1}{2(n(x)+1)}$$

□

As an example, in the framework of Figure 4.2, $x = e$ is contained in a conflict-free extension and $n(x) = 1$ so

$$\frac{3}{4} \geq s_F(e) \geq \frac{1}{4}$$

Indeed, we have computed $s_F(e) = 0.25$.

With respect to standard extension semantics in abstract argumentation frameworks, $n(x)$ is without loss of generality always reducible to at most 2 [Dunne, 2007]. Such reductions not only lead to changes in the inequalities (obtained by the previous result) that bound the strength of arguments, but also have an impact on the strength of the arguments. For example, the arguments a and a are equivalent with respect to the frameworks on the left (F) and the right (F') hand side of Figure 4.3, where $n(x) = 4$ and $n(x) = 2$ respectively, although $s_F(a) = 0.125$ and $s_{F'}(a) = 0.216$. By reducing the number of attacks against a from 3 in F to 2 in F', one increases its strength.

If an argument can be defended by an admissible opinion, then we can obtain a better lower bound on the strength of that argument than in the case where we only know that the argument is contained in a conflict-free extension. That

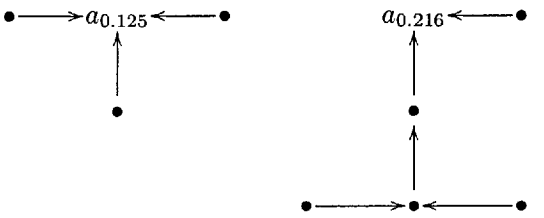


Figure 4.3: Equivalent arguments with distinct strength.

lower bound is remarkably equal to $1/2$ when the out-degree is $k = 1$.

Property 4.5. *If x is contained in an admissible extension, then*

$$s_F(x) \geq \max\left(\frac{1}{2(1+n(x))}, \frac{1}{1+\sqrt{k}}, \frac{k+3}{4(k+1)}\right).$$

Proof. There exists an admissible extension P that contains x and that the proponent can choose. Thus, for every opponent strategy O, P is conflict-free and attacks all the arguments contained in O that attack P. Namely, for every argument of O that attacks P, there is at least one counter-attack from P against that argument. In the worst case for the proponent, every argument in $O - P$ attacks P with k different attacks. Thus,

$$\phi(P, O) \geq \frac{1}{2}[1 + f(|O - P|) - f(k|O - P|)]$$

For every $t \geq 0$, let

$$g_k(t) = \frac{1}{2}[1 + f(t) - f(kt)] = \frac{1}{2}[1 + \frac{t}{t+1} - \frac{kt}{kt+1}]$$

We have

$$g'_k(t) = \frac{dg_k}{dt} = \frac{1}{2}[\frac{1}{(t+1)^2} - \frac{k}{(kt+1)^2}]$$

The minimum of g_k is thus attained for t such that $g'_k(t) = 0$, i.e.

$$\frac{1}{(t+1)^2} - \frac{k}{(kt+1)^2} = 0$$

whose positive solution is $t = 1/\sqrt{k}$. Thus, for every $t \geq 0$, we have $g_k(t) \geq g_k(1/\sqrt{k})$ and consequently

$$g_k(t) \geq \frac{1}{2}[1 + \frac{1/\sqrt{k}}{1+1/\sqrt{k}} - \frac{k*1/\sqrt{k}}{k*1/\sqrt{k}+1}] = \frac{1}{1+\sqrt{k}}$$

By playing always P the proponent of x can secure a payoff of at least

$$\frac{1}{1+\sqrt{k}}$$

Notice that $t = |O - P|$ is always an integer. If $t = 0$, then O does not attack P and $\phi(P, O) = 1$. Otherwise $t \geq 1$. In that case, $k \geq 1$. Then $1/\sqrt{k} \leq t$. Since g_k is monotonic increasing for the values of $t \geq 1/\sqrt{k}$, we also have

$$\phi(P, O) \geq \frac{1}{2}[1 + f(1) - f(k)] = \frac{k+3}{4(k+1)}$$

So for that same strategy P, the payoff is also guaranteed to be at least

$$\frac{k+3}{4(k+1)}$$

Since P is conflict-free, $\{x\}$ does not attack itself and by playing always $\{x\}$, the proponent can also secure a payoff of at least

$$\frac{1}{2(n(x)+1)}$$

□

The same lower bound holds *a fortiori* if we know that x is contained in a preferred, stable, complete, grounded or ideal extension, since all of these extensions are necessarily also admissible (again, see definitions in Chapter 2). For instance, in the framework of Figure 4.2, $x = c$ is contained in the stable extension $\{a, c, f\}$, $k = 2$ and $n = 1$ so

$$s_F(c) \geq \max\left(\frac{1}{4}, \frac{1}{1+\sqrt{2}}, \frac{5}{12}\right) = \frac{5}{12} = 0.417$$

We have indeed computed that $s_F(c) = 0.417$.

This dialectical measure of argument strength can be said to be "compatible" with the classical notions of acceptability of abstract argumentation insofar as the strength of arguments contained in admissible extensions is always guaranteed to be above a certain level (given by the previous property). However, the measure is finer-grained in that it allows to rank the arguments that are contained in a conflict-free extension without being contained in any admissible extension.

Dynamic properties

For any argument $a \in Arg$ such that $(a, x) \notin att$, we denote

$$F_{+(a,x)} = (Arg, \{(a, x)\} \cup att)$$

We also say that an argument a is *superfluous* with respect to x if forbidding the proponent of x to play with opinions containing a does not decrease the value of the (F, x) game. Given two frameworks $F = (Arg, att)$ and $F' = (Arg', att')$, we say that F and F' are *disjoint* when $Arg \cap Arg' = \emptyset$. We finally define for

any frameworks F and F' the sum

$$F + F' = (Arg \cup Arg', att \cup att')$$

We consider below a number of properties defined in terms of these notions.

As sanctioned by the property 4.6 below, adding an attack against an argument reduces its strength. For example, in Figure 4.4, the attack (b, a) reduces the strength of a from 1 to 1/2.



Figure 4.4: Example of attack reducing the strength of an argument.

Property 4.6. $s_{F_{+(a,x)}}(x) \leq s_F(x)$

Proof. The sets of strategies available to the proponent and opponent are the same in the (F, x) and $(F_{+(a,x)}, x)$ games. Let P and O be some proponent and opponent strategies. Note that $P_F^{\leftarrow O} \subseteq P_{F_{+(a,x)}}^{\leftarrow O}$ and either $O_F^{\leftarrow P} = O_{F_{+(a,x)}}^{\leftarrow P}$ (if $a \notin P$) or P attacks itself in $F_{+(a,x)}$ (if $a \in P$). By monotonicity of f, in any case ($a \in P$ or $a \notin P$), $\phi_{F_{+(a,x)}}(P, O) \leq \phi_F(P, O)$. It follows that $s_{F_{+(a,x)}}(x) \leq s_F(x)$. □

A sanctioned by property 4.7 below, an argument becomes stronger (respectively weaker) when it attacks more arguments, if the newly attacked arguments are superfluous (respectively not superfluous) for its defence. For example, in Figure 4.5, e is superfluous with respect to c since we know that the optimal mixed strategy for defending c consists in playing with only $\{a, c, f\}$ or $\{a, c\}$ (see Section 4.3.3) and e is not contained in any of these two extensions. The newly added attack from c against e (right hand side of Figure 4.5) does not decrease the strength of c (this strength is 0.417 before and after adding the attack). However, a is not superfluous with respect to c since a is contained in

all of the strategies ($\{a, c, f\}$ and $\{a, c\}$) optimally played to defend c, and one can observe in Figure 4.6 that the newly added attack from c against a (right hand side of Figure 4.6) reduces indeed the strength of c from 0.417 to 0.25.

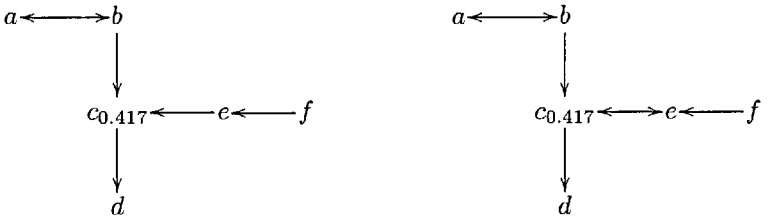


Figure 4.5: Adding an attack against a superfluous argument.

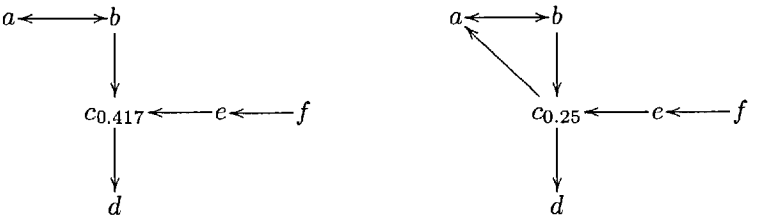


Figure 4.6: Adding an attack against a non-superfluous argument.

Property 4.7. *It holds that*

- *if a is superfluous with respect to x, then $s_{F_{+(x,a)}}(x) \geq s_F(x)$, and*
- *if a is not superfluous with respect to x, then $s_{F_{+(x,a)}}(x) \leq s_F(x)$.*

Proof. If a is superfluous with respect to x then there exists an optimal mixed strategy $X^* = (x_i^*)$ for the (F, x) game such that, $\forall i$, $x_i^* > 0 \Rightarrow a \notin P_i$. Let then P be an "active" strategy under X^*, i.e. $P = P_i$ and $x_i^* > 0$. Then, $\forall O$, we have $O_F^{\leftarrow P} \subseteq O_{F_{+(x,a)}}^{\leftarrow P}$, $P_F^{\leftarrow O} = P_{F_{+(x,a)}}^{\leftarrow O}$ (if it is not the case that $x \in O$ and $a \in P$) or P attacks itself in $F_{+(x,a)}$ (if $x \in O$ and $a \in P$). The last case cannot occur, indeed a is assumed to be superfluous with respect to x, P is assumed to be an active strategy and thus $a \notin P$. So, by monotonicity

of f, $\phi_F(P,O) \leq \phi_{F_{+(x,a)}}(P,O)$. Since $a \notin P$, P is conflict-free in F iff P is conflict-free in $F_{+(x,a)}$ and O attacks P in F iff O attacks P in $F_{+(x,a)}$. Therefore, we have proved that for every active strategy P under X^* we have $\phi_F(P,O) \leq \phi_{F_{+(x,a)}}(P,O)$. By playing with X^* in the $(F_{+(x,a)},x)$ game, the proponent can secure a payoff of at least $s_F(x)$. Hence, $s_{F_{+(x,a)}}(x) \geq s_F(x)$.

If a is not superfluous with respect to x, then the proponent of x must play strategies containing x but not a, as all strategies containing both x and a are non-conflict-free in $F_{+(x,a)}$ and result by definition 4.8 in a payoff that is zero. Since the proponent of x is restricted to playing only with strategies containing x but not a and a is not superfluous with respect to x, the proponent's payoff $s_F(x)$ is reduced by definition of the term superfluous. □

As sanctioned by property 4.8 below, adding an attack against an attacker of some argument increases the strength of that same argument. For example, in Figure 4.7, the addition of the attack (c,a) against the attacker a of b increases the strength of b from 0.25 to 0.5.



Figure 4.7: Adding a defender to an argument.

Property 4.8. *If y attacks x in F, then $s_{F_{+(z,y)}}(x) \geq s_F(x)$.*

Proof. The sets of strategies of the players are the same in the (F,x) and $(F_{+(z,y)},x)$ games. We have $O_F^{\leftarrow P} \subseteq O_{F_{+(z,y)}}^{\leftarrow P}$ (if $z \in P$ and $y \in O$) and $O_F^{\leftarrow P} = O_{F_{+(z,y)}}^{\leftarrow P}$ (otherwise). We also have $P_F^{\leftarrow O} \subseteq P_{F_{+(z,y)}}^{\leftarrow O}$ (if $y \in P$ and $z \in O$) and $P_F^{\leftarrow O} = P_{F_{+(z,y)}}^{\leftarrow O}$ (otherwise). Note that if $y \in P$ then P attacks itself in both F and $F_{+(z,y)}$. So, $\phi_F(P,O) \leq \phi_{F_{+(z,y)}}(P,O)$ and $s_F(x) \leq s_{F_{+(z,y)}}(x)$. □

The strength of an argument only depends on the arguments that are related to it in the dialectical sense. Formally,

Property 4.9. *If* $F = (Arg, att)$ *and* $F' = (Arg', att')$ *are disjoint and* $x \in Arg$, *then* $s_{F+F'}(x) = s_F(x)$.

Proof. Let us denote by

- v_F the value of the (F, x) game
- $v_{F+F'}$ the value of the $(F + F', x)$ game
- $v_{F+F'}(P \subseteq Arg)$ the value of a variant of the $(F + F', x)$ game where the proponent is restricted to play only with strategies $P \subseteq Arg$
- $v_{F+F'}(O \subseteq Arg)$ the value of a variant of the $(F + F', x)$ game where the opponent is restricted to play only with strategies $O \subseteq Arg$

The gain of a player cannot strictly increase by restriction of his set of possible strategies. Thus,

- for the proponent's gain, we have $v_{F+F'}(P \subseteq Arg) \leq v_{F+F'}$, and
- for the opponent's gain, we have $-v_{F+F'}(O \subseteq Arg) \leq -v_{F+F'}$, or equivalently, $v_{F+F'}(O \subseteq Arg) \geq v_{F+F'}$.

In the $(F+F', x)$ game, if one player plays only with strategies that are subsets of Arg, then his adversary wins or loses nothing by also playing with strategies that are subsets of Arg, since the arguments in Arg' neither attack these strategies nor are attacked by them. So, we also have

- $v_{F+F'}(P \subseteq Arg) = v_F$ and
- $v_{F+F'}(O \subseteq Arg) = v_F$.

In summary,

$$v_{F+F'} \geq v_{F+F'}(P \subseteq Arg) = v_F = v_{F+F'}(O \subseteq Arg) \geq v_{F+F'}$$

so that $v_{F+F'} = v_F$. □

This last property is quite useful, as it allows to distribute the computation of the strength in frameworks that have disjoint components.

4.4 Social welfare-based aggregation of value arguments

In the previous section, we have defined an argument strength measure s_F for any abstract argumentation framework F. Here, we consider frameworks of the kind described in Section 4.2. There, arguments are either value or mitigation arguments. Figure 4.8 shows the strength of arguments given in Figure 4.1, Section 4.2.

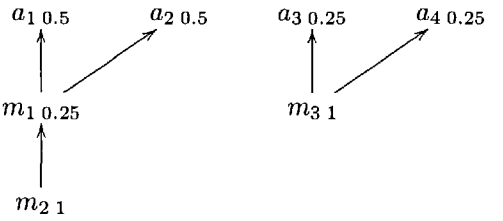


Figure 4.8: Strength of arguments in the umbrella problem.

In this section, we define a method for aggregating value arguments using this notion of strength to weigh arguments.

Conceptually, the issue of aggregating value arguments is similar to the one of aggregating the individual preferences of a society or group of individuals $\{a_1, \ldots, a_k\}$. In the areas of Social Welfare and Economics, influential works such as [Fleming, 1952], [Harsanyi, 1955] and [Dyer and Sarin, 1978] have identified rational preference aggregation rules. Fleming and Harsanyi were the first to argue (both mathematically and philosophically) in favour of simple classes of social welfare functions. These social welfare functions are linear weighted combinations of individual preference functions over outcomes $V_{a_1}, \ldots, V_{a_k} : X \rightarrow \mathbb{R}$,

defined as

$$V = \sum_{i=1}^{k} \lambda_{a_i} V_{a_i}$$

The weights $\lambda_{a_1}, \ldots, \lambda_{a_k} \in \mathbb{R}^+$ are called *interpersonal weights*. Each weight λ_{a_i} represents the social importance of a_i, i.e. the value taken by the social welfare function V when only the preference of individual a_i is taken into account.

We suggest to draw a parallel between individual preferences and value arguments, as well as between interpersonal weights and the strength of arguments. These analogies suggest that value functions can be constructed in the same way as social welfare functions. In this perspective, value and mitigation arguments constitute a society in which the social importance of value arguments corresponds to their dialectical strength and mitigation arguments play the role of counter-powers that regulate the importance (given by the strength) of value arguments.

Given a set A of value arguments as in Section 4.2, we define

Definition 4.11. *For each value argument $a = (J_a, C_a) \in A$ where $C_a = (g_a, v_a)$, the **individual value function** $V_a : 2^G \to \mathbb{R}$ is defined, $\forall o \subseteq G$, by*

$$V_a(o) = \begin{cases} v_a & \text{if} \quad g_a \in o \\ 0 & \text{otherwise} \end{cases}$$

Equivalently, $V_a(o) = v_a I_o(g_a)$, where $I_o : G \to \{0, 1\}$ is the characteristic function of the set o. Since mitigation arguments do not constitute statements directly concerning outcomes, we define $V_m = 0$ for every $m \in M$. Choosing as interpersonal weight λ_a for each argument $a \in Arg = A \cup M$ the strength $s_F(a)$ of a, we obtain for every outcome $o \subseteq G$

$$V(o) = \sum_{a \in Arg} s_F(a) V_a(o) = \sum_{a \in A} s_F(a) V_a(o)$$

This finally gives us

Definition 4.12. *The **dialectical value function** V is given for every outcome o by*

$$V(o) = \sum_{a \in A} s_F(a) V_a(o)$$

Applying this formula to the umbrella problem, we obtain

- $V(\{\text{dry}, \text{loaded}\}) = 0.5 * (+1) * 1 + 0.5 * (-1) * 0 + 0.25 * (-0.25) * 1 + 0.25 * (+0.25) * 0 = 0.5 - 0.25 * 0.25 = 0.4375$
- $V(\{\text{dry}, \neg\text{loaded}\}) = 0.5 * (+1) * 1 + 0.5 * (-1) * 0 + 0.25 * (-0.25) * 0 + 0.25 * (+0.25) * 1 = 0.5 + 0.25 * 0.25 = 0.5625$
- $V(\{\neg\text{dry}, \neg\text{loaded}\}) = 0.5 * (+1) * 0 + 0.5 * (-1) * 1 + 0.25 * (-0.25) * 0 + 0.25 * (+0.25) * 1 = -0.5 + 0.25 * 0.25 = -0.4375$

As intuitively expected,

$$V(\{\text{dry}, \neg\text{loaded}\}) > V(\{\text{dry}, \text{loaded}\}) > V(\{\neg\text{dry}, \neg\text{loaded}\})$$

We are now going to study the properties of dialectical value functions.

Let V be a dialectical value function and let us consider the preference relation $\succ$ induced by V. This relation is defined $\forall o, o' \subseteq G$ as $o \succ o'$ if and only if $V(o) > V(o')$. For every value argument a, we also use the individual preference $\succ_a$ defined for any outcomes o and o' as $o \succ_a o'$ if and only if $V_a(o) > V_a(o)$ and read 'o is strictly preferred to o' according to a'. Since V is by construction a linear social welfare function, the following three properties hold (refer to [Fleming, 1952] for the proofs of these results).

Property 4.10. *The preference relation $\succ$ is a total order on X.*

Namely, dialectical value functions induce rational preferences.

Property 4.11. *If* $o \succ_a o'$ *and* $\neg(o' \succ_b o)$ *for every* $b \neq a$, *then* $o \succ o'$.

Namely, if the preference induced by an argument over two outcomes does not conflict with the preference induced by any other argument, then the preference induced by the dialectical value function agrees with that argument.

Property 4.12. *If* $o \succ_a o'$ *and* $o' \succ_b o$, $V_a(o) - V_a(o') = V_b(o') - V_b(o)$, *and neither* $o \succ_c o'$ *nor* $o' \succ_c o$ *for every* $c \neq a \neq b$, *then* $o \succ o'$ *iff* $s_F(a) > s_F(b)$.

Namely, when there is a conflict between the preferences induced by two arguments a and b and the preferences induced by all other value arguments c are indifferent, the preference induced by the dialectical value function depends exclusively on the relative strength of a and b.

The mathematical properties of s_F imply additional properties for the dialectical value function V. These properties are expressed in terms of

- an argumentation framework F composed of a set A of value arguments and a set M of mitigation arguments,
- an argumentation framework F' disjoint from F and composed of only one value argument a' and a set M' of mitigation arguments.

We will denote $V_{F+F'}$ the value function obtained from $F+F'$ and V_F the value function obtained from F only. For every outcome o, we also define the set of value arguments in favour or pro o.

Definition 4.13. *The set of* ***arguments pro*** *the outcome* o *is*

$$Pro(o) = \{a = (J_a, (g_a, v_a)) \in L \times (G \times S) \mid g_a \in o, v_a \geq 0\}$$

Similarly, we define the set of value arguments against or con o.

Definition 4.14. *The set of* ***arguments con*** *the outcome o is*

$$Con(o) = \{a = (J_a, (g_a, v_a)) \in L \times (G \times S) \mid g_a \in o, v_a < 0\}$$

An argument pro an outcome o increases the value of o. Formally,

Property 4.13. *If $a' \in Pro(o)$, then $V_{F+F'}(o) = V_F(o) + s_{F'}(a')v_{a'} \geq V_F(o)$.*

Proof.

$$\begin{aligned}
V_{F+F'}(o) &= \sum_{a \in A \cup \{a'\}} s_{F+F'} V_a(o) = \sum_{a \in A} s_{F+F'}(a) V_a(o) + s_{F+F'}(a') V_{a'}(o) \\
&= \sum_{a \in A} s_F(a) V_a(o) + s_{F'}(a') V_{a'}(o) \text{ (by prop. 4.9)} \\
&= \sum_{a \in A} s_F(a) V_a(o) + s_{F'}(a') v_{a'} \text{ (since } a' \in Pro(o) \text{ and thus } g_{a'} \in o) \\
&\geq \sum_{a \in A} s_F(a) V_a(o) \text{ ($s_{F'}(a') \geq 0$ by prop. 4.1, $v_{a'} \geq 0$ since } a' \in Pro(o))
\end{aligned}$$

□

Similarly, an argument con o decreases the value of o. Formally,

Property 4.14. *If $a' \in Con(o)$, then $V_{F+F'}(o) = V_F(o) + s_{F'}(a')v_{a'} \leq V_F(o)$.*

Proof.

$$\begin{aligned}
V_{F+F'}(o) &= \sum_{a \in A \cup \{a'\}} s_{F+F'} V_a(o) = \sum_{a \in A} s_{F+F'}(a) V_a(o) + s_{F+F'}(a') V_{a'}(o) \\
&= \sum_{a \in A} s_F(a) V_a(o) + s_{F'}(a') V_{a'}(o) \text{ (by prop. 4.9)} \\
&= \sum_{a \in A} s_F(a) V_a(o) + s_{F'}(a') v_{a'} \text{ (since } a' \in Con(o) \text{ and thus } g_{a'} \in o) \\
&\leq \sum_{a \in A} s_F(a) V_a(o) \text{ ($s_{F'}(a') \geq 0$ by prop. 4.1, $v_{a'} < 0$ since } a' \in Con(o))
\end{aligned}$$

□

Note that the increase or decrease of value caused by the value argument a' in F' depends on its strength with respect to the mitigation arguments in M'.

The value of an outcome can be bounded below by the sum of the values of its arguments con and above by the sum of the values of its arguments pro.

Property 4.15.

$$\sum_{a \in A \cap Con(o)} v_a \leq V_F(o) \leq \sum_{a \in A \cap Pro(o)} v_a$$

Proof. Since $V_a(o) = 0$ whenever $g_a \notin o$, we have

$$V_F(o) = \sum_{a \in A \cap Pro(o)} s_F(a) V_a(o) + \sum_{a \in A \cap Con(o)} s_F(a) V_a(o)$$

We have $V_a(o) = v_a \geq 0$ for every argument $a \in A \cap Pro(o)$ and $V_a(o) = v_a < 0$ for every argument $a \in A \cap Con(o)$. By property 4.1, we also have $0 \leq s_F(a) \leq 1$, consequently

$$V_F(o) \geq \sum_{a \in A \cap Con(o)} s_F(a) v_a \geq \sum_{a \in A \cap Con(o)} v_a$$

$$V_F(o) \leq \sum_{a \in A \cap Pro(o)} s_F(a) v_a \leq \sum_{a \in A \cap Pro(o)} v_a$$

□

The following four properties are deduced from the dynamic properties of our dialectical strength measure. Assuming $a \in A$ is a value argument of F, $b \in M$ is a mitigation argument of F that does not attack a and denoting as in Section 4.3.4 by $F_{+(b,a)}$ the framework in which the attack (b, a) is added to F, then we also have

Property 4.16. *If $a \in Pro(o)$, then $V_{F_{+(b,a)}}(o) \leq V_F(o)$.*

Proof. This derives directly from the fact that $g_a \in o$, $v_a \geq 0$ and $s_{F_{+(b,a)}}(a) \leq$

$s_F(a)$ which holds by property 4.6. □

Property 4.17. *If* $a \in Con(o)$, *then* $V_{F_{+(b,a)}}(o) \geq V_F(o)$.

Proof. This derives directly from the fact that $g_a \in o$, $v_a < 0$ and $s_{F_{+(b,a)}}(a) \leq s_F(a)$ which holds by property 4.6. □

Intuitively, these last two properties mean that adding an attack against an argument pro an outcome reduces the value of the outcome, and that adding an attack against an argument con an outcome increases the value of the outcome.

In the next two properties, we assume that a is a value argument in F, that m and m' are two mitigation arguments in F such that m attacks a and m' does not attack m in F. We naturally denote by $F_{+(m',m)}$ the framework in which the attack (m', m) is added to F. We have in this case

Property 4.18. *If* $a \in Pro(o)$, *then* $V_{F_{+(m',m)}}(o) \geq V_F(o)$.

Proof. This derives directly from the fact that $g_a \in o$, $v_a \geq 0$ and $s_{F_{+(m',m)}}(a) \geq s_F(a)$ which holds by property 4.8. □

Property 4.19. *If* $a \in Con(o)$, *then* $V_{F_{+(m',m)}}(o) \leq V_F(o)$.

Proof. This derives directly from the fact that $g_a \in o$, $v_a < 0$ and $s_{F_{+(m',m)}}(a) \geq s_F(a)$ which holds by property 4.8. □

Intuitively, these last two properties mean that adding a defender to an argument pro an outcome increases the value of the outcome, and that adding a defender to an argument con an outcome decreases the value of the outcome.

Now, we consider that $F_1, \ldots, F_k$ are mutually disjoint argumentation frameworks composed of value and mitigation arguments. We show that

Property 4.20. *If* $F_1, \ldots, F_k$ *are disjoint and* $F = F_1 + \ldots + F_k$, *then*

$$V_F = V_{F_1} + \ldots + V_{F_k}$$

Proof. This derives directly from the additivity of V_F with respect to value arguments and property 4.9 concerning the strength of arguments contained in mutually disjoint frameworks. □

This last property means that the value functions obtained from dialectically independent bodies of uncertain knowledge concerning the desirability of outcomes simply add up.

Now that we have defined dialectical value functions and studied their properties, we may use them for the practical construction of value functions. In the next section, we construct value functions for poker playing programs.

4.5 Value functions for poker playing programs

As mentioned in Chapter 3, and unlike games of perfect information such as chess, checkers, othello and backgammon where brute-force search techniques have proved to be quite effective, the game of poker is nowadays still considered as a challenge to artificial intelligence, chiefly because of the difficulty of assessing the expected value of betting decisions. [Billings *et al.*, 1999] provides an excellent account of state-of-the-art game of poker playing program architectures. Programs able to play a strong game of poker rely on knowledge and observation based probabilistic models of opponent players, techniques to initialise and re-weight probabilities after each round, stochastic procedures such as selective sampling to estimate the expected value of each betting decision, ways to randomise playing strategies, etc.

Statistical studies show that the strategies adopted by human players are very much dependent on their psychological and emotional states [Caro, 2003]. For instance, risk averse players tend to fold too many hands pre-flop, greedy players often make exaggerated bets, non-opportunistic players easily waste pre-

mium pocket cards, patient players are better able to trap their opponents, aggressive players overly rely on bluffing, etc. We envisage the possibility of enriching existing poker playing programs with human-like psychological and emotional features. This can be done by re-designing the value functions used by poker playing programs to assess the expected value of betting decisions. Current programs such as Poki 2 [Billings *et al.*, 1999] only value an outcome by considering the amount of money or chips gained. Following Section 4.2, we suggest to substitute the monetary scale with a desirability scale implicitly given by the last column of Table 4.2 and we model the influence of a player's psychological and emotional state on its valuation using the value arguments defined in that same table. Justifications here relate to various possible features of the player's psychological and emotional state. The claims of the argument express how much satisfaction or enjoyment a player gets in different circumstances. For instance, risk averse players most enjoy playing hands that do not put them at risk, as opposed to risk neutral or risk seeking players, greedy players like to see other players increasing the pot size (as they see it as their potential gain), and aggressive players have a strong desire to increase their chances of winning pots by frightening their opponents (bluffing).

The mental states and strategies of poker players constantly change in time. Within each game, the number of opponents decreases, the size pot increases and the uncertainty is reduced as more community cards get revealed. Game after game, more players get eliminated (in tournaments), the blinds (minimum betting amounts) periodically increase, putting players under time pressure and forcing short-stacked players to take more risk, players may gain a lot of knowledge about each other, etc. All these changes in the game scene have a direct influence on the players' mental state. These influences are modeled as arguments mitigating the natural or default profile of each player.

justification J_a	goal g_a	desirability v_a
RA (risk averse)	NR^+ (no risk)	0.5 (desirable)
RA (risk averse)	SR^- (small risk)	−0.25 (fairly undesirable)
RA (risk averse)	BR^- (big risk)	−1 (quite undesirable)
SRA (slightly risk averse)	NR^+ (no risk)	0.5 (desirable)
SRA (slightly risk averse)	SR^- (small risk)	0.25 (fairly desirable)
SRA (slightly risk averse)	BR^- (big risk)	−0.5 (undesirable)
RN (risk neutral)	NR^+ (no risk)	0.5 (desirable)
RN (risk neutral)	SR^- (small risk)	0.5 (desirable)
RN (risk neutral)	BR^- (big risk)	0 (null/indifferent)
NG (non-greedy)	APV^+ (add pot value)	−0.5 (undesirable)
G (greedy)	APV^+ (add pot value)	0.5 (desirable)
VG (very greedy)	APV^+ (add pot value)	1 (quite desirable)
O (opportunistic)	LB^+ (likely best hand)	1 (quite desirable)
O (opportunistic)	UB^- (unlikely best hand)	−1 (quite undesirable)
NO (not opportunistic)	LB^+ (likely best hand)	0.5 (desirable)
NO (not opportunistic)	UB^- (unlikely best hand)	−0.5 (undesirable)
P (patient)	LBF^+ (likely best hand in future)	0.5 (desirable)
P (patient)	UBF^- (unlikely best hand in future)	−0.5 (undesirable)
NP (not patient)	LBF^+ (likely best hand in future)	0.25 (fairly desirable)
NP (not patient)	UBF^- (unlikely best hand in future)	−0.25 (fairly undesirable)
NA (non aggressive)	IFC^+ (increase future chances)	−0.25 (slightly undesirable)
A (aggressive)	IFC^+ (increase future chances)	0.5 (desirable)
VA (very aggressive)	IFC^+ (increase future chances)	1 (quite desirable)

Table 4.2: Value arguments depending on the psychological and emotional state of the player.

We propose the mitigation arguments as defined in Table 4.3. The top half of the table gives mitigation arguments that attack value arguments. The bottom half gives mitigation arguments that attack other mitigation arguments. Two levels of attacks seem sufficient here to express the conflicts and uncertainties of a game of poker player concerning the desirability of his objectives. The resulting argumentation frameworks can only be composed of smaller graphs disconnected with one another, allowing us to distribute the computation of the strength of the arguments (see property 4.20). Consider for example a risk averse (RA) but greedy (G) and very aggressive (VA) player being short-stacked and almost out of the table (KO), and playing against loose opponents (LO), some of them being also greedy players (GO). Suppose the player's hand is really bad, making him almost sure to lose (L). The player considers the idea of bluffing, but believes that due to his relative lack of experience, the bluff would most likely fail (BF). This entire situation would be modeled by the argumentation framework shown in Figure 4.9.

justification J_m	claim $\neg J_a$
KO (short-stacked and almost out)	¬RA
LO (only loose opponents left)	¬RA
H (heads-up or very close to end of tournament)	¬SRA
SET (short-stacked but enough time to rebuild chip stack)	¬RN
CO (a conservative opponent still in competition after flop)	¬RN
DO (need to damage stack of very threatening opponent)	¬NG
GO (several greedy opponents left)	¬G
T (wish to trap opponent or look friendly)	¬VG
KO	¬O
KOO (opportunity to kick-out a short-stacked opponent)	¬NO
SA (only slim advantage over opponents)	¬P
W (almost sure to win pot)	¬NP
SB (attempt to semi-bluff)	¬NA
SB (attempt to semi-bluff)	¬VA
B (attempt to bluff may succeed)	¬NA
BF (attempt to bluff is likely to fail)	¬NA
BF (attempt to bluff is likely to fail)	¬VA
L (almost sure to lose hand)	¬KO
MC (some opponents are getting more conservative)	¬LO
SHO (opponent seems to have a strong hand)	¬H
ST (very strong table with experienced players)	¬SET
COB (conservative opponent seems to bluff)	¬CO
ALLIN (another player has gone all in against the threatening opponent)	¬DO
W	¬GO
SHO	¬T
WP (the opponent that is almost out is a weak player)	¬KOO
NOB (number of outs is big)	¬SA
TOP (an opponent could have a straight, flush, full house or four of a kind)	¬W
NOS (number of outs is small)	¬SB
CALL (an inexperienced/upset opponent is likely to call you)	¬B
KO	¬BF

Table 4.3: Mitigation arguments depending on the playing scene.

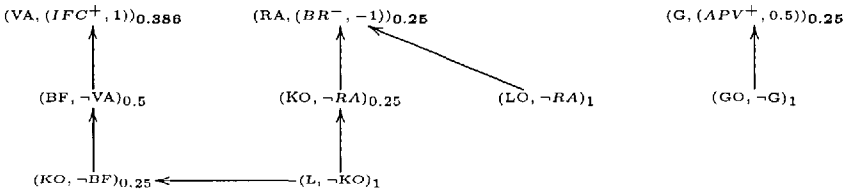


Figure 4.9: Framework modeling the uncertainties of a desperate bluffer.

Once the argumentation framework F is constructed, one can apply the game-theoretic method of Section 4.3 to estimate the strength of value (and optionally mitigation) arguments. The strength of each argument is indicated in subscript and printed in bold font in the case of value arguments in Figure 4.9. The resulting dialectical value of the outcome $x = \{BR^-, APV^+, IFC^+\}$ would be here $V_F(x) = 0.25 * (-1) + 0.25 * (0.5) + 0.386 * (1) = 0.261 \approx 0.25$. So, despite the risk involved, bluffing would still appear to be a slightly desirable betting decision for that desperate aggressive player.

We recall in Figure 4.10 the general decision table generated in Section 3.7,

where d_1 is fold/check, d_2 is call and d_3 is raise. We may now assess the value of each outcome given by that table. Suppose an opponent has raised before the bluffer. Then the possibility to check is discarded and d_1 becomes fold. The results for the value of the outcomes in the corresponding table are computed in Figure 4.11. Of course, these results are specific to the case of a desperate bluffer and should not be used for different types of players nor bluffers playing in different circumstances.

·	$w_1 \vee w_3$	$w_2 \vee w_4$	$w_5 \vee w_7$
d_1	NR^+, APV^+, LB^+, LBF^+	NR^+, APV^+, LB^+, UBF^-	NR^+, LB^+, LBF^+
d_2	SR^-, APV^+, LB^+, LBF^+	SR^-, APV^+, LB^+, UBF^-	SR^-, APV^+, LB^+, LBF^+
d_3	SR^-, APV^+, LB^+, LBF^+	SR^-, APV^+, LB^+, UBF^-	$SR^-, APV^+, IFC^+, LB^+, LBF^+$

·	$w_6 \vee w_8$	$w_9 \vee w_{10}$	$w_{11} \vee w_{12}$
d_1	NR^+, LB^+, UBF^-	NR^+, APV^+, UB^-, UBF^-	NR^+, APV^+, UB^-, LBF^+
d_2	SR^-, APV^+, LB^+, UBF^-	SR^-, APV^+, UB^-, UBF^-	SR^-, APV^+, UB^-, LBF^+
d_3	$SR^-, APV^+, IFC^+, LB^+, UBF^-$	BR^-, APV^+, UB^-, UBF^-	BR^-, APV^+, UB^-, LBF^+

·	$w_{13} \vee w_{14}$	$w_{15} \vee w_{16}$
d_1	NR^+, UB^-, UBF^-	NR^+, UB^-, LBF^+
d_2	SR^-, APV^+, UB^-, UBF^-	SR^-, APV^+, UB^-, LBF^+
d_3	$BR^-, APV^+, IFC^+, UB^-, UBF^-$	$BR^-, APV^+, IFC^+, UB^-, LBF^+$

Figure 4.10: Poker decision table.

·	$w_1 \vee w_3$	$w_2 \vee w_4$	$w_5 \vee w_7$	$w_6 \vee w_8$	$w_9 \vee w_{10}$	$w_{11} \vee w_{12}$	$w_{13} \vee w_{14}$	$w_{15} \vee w_{16}$
d_1	0.125	0.125	0	0	0.125	0.125	0	0
d_2	0.125	0.125	0.125	0.125	0.125	0.125	0.125	0.125
d_3	0.125	0.125	0.511	0.511	−0.125	−0.125	0.261	0.261

Figure 4.11: Value function for the desperate bluffer after a raise.

In the light of this value function and if no probabilistic information concerning the scenarios is available, the decisions should be compared using one of the criteria for decision under strict uncertainty (see Chapter 2). Nevertheless, it is already clear that a rational bluffer would not fold in this case, as no matter which scenario occurs, the outcome resulting from a call is always at least as desirable as the one resulting from folding. Observe from the table in Figure 4.12 that all of the four classical criteria reject that decision. In fact, Wald's pessimistic criterion leads to a call, while Hurwicz, Savage and Laplace's criteria all lead to a raise.

	fold/check	call	raise	best decision
Wald's maximin return	0	0.125	−0.125	call
Hurwicz optimism-pessimism index ($x = 0.5$)	0.0625	0.125	0.193	raise
Savage's minimax regret	0.511	0.386	0.25	raise
Laplace's principle of insufficient reason	0.0625	0.125	0.193	raise

Figure 4.12: Best decision for the desperate bluffer (strict uncertainty).

In the case of decision under risk, the best decision may vary between call and raise. To be fully rigorous, the decision of folding could be equally desirable as checking, if the probability of the event $E = \{w_5, w_7, w_{13}, w_{14}, w_{15}, w_{16}\}$ was null. No matter what the probability distribution P of the scenarios is, the expected value of checking is always equal to $E(\text{check}) = 0.125$. Whether a raise should be preferred to a call very much depends here on the probability of $E' = \{w_9, w_{10}, w_{11}, w_{12}\}$ (ending the turn with a bad hand and facing strong opponents) since raising under this event has undesirable consequences (with an associated negative value of -0.125 in units of desirability) for the player. If the bluffer thinks he has a fifty percent chance of avoiding this risk (occurrence of E') and that all other scenarios are equally likely, the expected value of a raise would be $E(\text{raise}) = 0.125 * P(\{w_1, w_3, w_2, w_4\}) + 0.511 * P(\{w_5, w_7, w_6, w_8\}) - 0.125 * P(E') + 0.216 * P(\{w_{13}, w_{14}, w_{15}, w_{16}\}) = 0.125 * 4 * 0.5/12 + 0.511 * 4 * 0.5/12 - 0.125 * 0.5 + 0.216 * 4 * 0.5/12 = 0.0208\overline{3} + 0.0851\overline{6} + 0.0625 + 0.036 = 0.2045$. Since $E(\text{raise}) > E(\text{check})$ the bluffer should again prefer to raise.

4.6 Related work

The vast majority of existing works on value functions is in fact concerned with utility functions. A *utility function* U is a value function $U : X \to \mathbb{R}$ that has the property of representing the preferences of a subject (decision maker). Various types of utility functions have been studied, depending on the type (i.e. structure and properties) of the preference relation considered. We consider these here-after and then relate them to our dialectical value function.

In game theory [von Neumann and Morgenstern, 1944] for instance, rational players are assumed to be able to compare uncertain outcomes (risky alternatives). The preference relation employed is a binary relation $\succeq$ over the set of lotteries, where in a lottery $\langle p_1, x_1 ; \ldots ; p_n, x_n \rangle$, the subject gets by definition the price $x_i \in X$ with probability p_i. The relation is assumed to be complete, transitive, Archimedean and to satisfy the property of independence to irrelevant alternatives. The expected utility theorem shows that there exists a value function u, unique up to scaling changes, whose expectation represents the preference $\succeq$ in the sense that for every lotteries L and L', $L \succeq L'$ if and only if $\sum_{i=1}^{n} u(x_i)p_i \geq \sum_{i=1}^{n'} u(x'_i)p'_i$. A value function u satisfying this property is called a *von Neumann-Morgenstern* utility function.

In decision analysis [Keeney and Raiffa, 1976; Farquhar, 1977; Fishburn, 1979], marketing science [Hauser and Urban, 1975; Eliashberg, 1980] and psychology [Goodman *et al.*, 1979; Lee, 1971; Tversky, 1967], outcomes often correspond to continuous variables such as wealth, capital or labour. The infiniteness of the set of outcomes X renders the measurement of von Neumann-Morgenstern utility functions from a subject practically impossible (the subject indeed cannot compare infinitely many lotteries). For this reason, behavioural assumptions are made concerning the subject and used to reduce the utility function to a special analytical form. Under such assumptions, the measurement of a von-Neumann Morgenstern utility function boils down to the estimation of one or several parameters. [Keeney and Raiffa, 1976] provide an excellent account of these approaches. For example, the utility of a sum of gains $x \geq 0$ is of the form $u_r(x) = 1 - e^{-rx}$ for a subject with constant absolute risk aversion $r > 0$. The parameter r can be estimated by asking the subject which probability value p makes the lottery $\langle 1 - p, x ; p, z \rangle$ equally preferable to getting a certain reward of y, for some fixed amounts $x < y < z$.

In qualitative decision theory [Tan and Pearl, 1994; Bacchus and Grove, 1996; Boutilier *et al.*, 1999; McGeachie and Doyle, 2002], outcomes are represented in a logical form, e.g. $x = \{p, q, \neg r\}$, $y = \{\neg p, q, \neg r\}$, $z = \{\neg p, \neg q, \neg r\}$, etc. The preference relations take the form of reflexive, transitive and not necessarily complete relations (i.e. pre-orders) over outcomes. These preferences are represented in a quite compact fashion by collections of statements of the form $f \rhd g$, meaning that any outcome x satisfying formula f is preferred to any outcome satisfying formula g, all literals of x and y being equal outside the support of the formulae f and g. The utility functions used to represent such preferences are called *ceteris paribus* utility functions.

Dialectical value functions are not derived from a preference relation and as a consequence do not constitute utility functions. Dialectical value functions are derived from judgements of desirability, expressed in terms of value and mitigation arguments. Value arguments differ from preference statements in the sense that a) the subject does not need to be certain about their validity, b) they relate to the absolute desirability of goals, rather than than the relative desirability of outcomes or lotteries, and c) they allow the subject to examine his/her own uncertainties. Like ceteris paribus preferences, sets of arguments can be incomplete, making dialectical value functions always measurable. Our knowledge model based on abstract argumentation and linear value function makes a behavioural assumption concerning the decision maker, as done for the practical construction of von Neumann-Morgenstern utility functions. We thus essentially assume that the decision maker is able to express arguments, provide justifications, identify their conflicts and to assign each argument the importance that it deserves from a dialectical standpoint.

In argumentation for decision making [Amgoud, 2005; 2009] (see Chapter 2), the arguments employed concern either the pros and cons of decisions (practical

arguments) or beliefs (epistemic arguments). In our approach, we focus on the desirability of outcomes (via value arguments) rather than decisions and thus do not make use of practical arguments. Value arguments can like practical arguments also be either pro or con. More precisely, value arguments associate numerical values of desirability to goals. Obviously, our goals play the same role as "values" in value-based argumentation [Bench-Capon and Dunne, 2002; Bench-Capon, 2003] and our scale of desirability plays the same role as the preference relation assumed over these "values". We also importantly rely on mitigation arguments, which naturally can be seen as epistemic arguments. A remark that may be of interest to practitioners of argumentation for decision making is that Fleming's preference aggregation mechanism [Fleming, 1952] and the dialectical strength measure used to aggregate value arguments (see definition 4.12) can serve as a non-polar decision principle [Amgoud and de Saint-Cyr, 2008] and so be used in argumentation for decision making to value decisions from sets of (accepted) practical arguments in favour or against decisions.

Other dialectical measures of argument strength such as [Besnard and Hunter, 2000] and [Cayrol and Lagasquie-Schiex, 2005] may also prove useful for the practical construction of value functions as these measures provide different rankings of the arguments, as observed in Section 4.3. When the arguments considered by a subject are those advanced by others, it is important, in order to evaluate the strength of these arguments, not only to look at their dialectical structure, but also to adopt the viewpoint (knowledge-base) of the subject. To this effect, [Hunter, 2004] designed a measure of strength called *believability* that is an interaction-based measure, in the sense that it depends on the attacks (undercuts) between logical arguments, but which is at the same time intrinsic, insofar as its computation involves the degree of entailment of the subject's knowledge base for the arguments considered. This measure has fur-

thermore the great advantage of being computationally much simpler than our game-theoretic notion of strength, as it can be computed by iteration on the finite tree of attackers and defenders of the argument that needs to be evaluated, rather than by linear optimisation.

4.7 Summary and conclusion

We have shown that abstract argumentation [Dung, 1995] offers an elegant framework for modeling the uncertain knowledge of a decision maker concerning the desirability of his objectives, and game theory [von Neumann, 1928] offers a rational way to quantify the strength of arguments in abstract argumentation. Moreover, we have applied Fleming's social welfare theory [Fleming, 1952] to the case of a "society" of value arguments to construct (by aggregation) a dialectical value function. We have discussed some of the practical advantages offered by this approach, shown the normative goodness of dialectical value functions, their adequacy with the decision maker's beliefs and have used them to enrich the strategy of poker playing programs.

Chapter 5

Capturing uncertainty with imprecise probabilities

The term *uncertainty* is generally used to denote the mental state of a subject who knows in which states a given object of interest can possibly be, but who ignores which state precisely is the true one. In the context of decision making under strict uncertainty or risk, the possible states $w_1, \ldots, w_n$ are referred to as *scenarios* and the set $\Omega = \{w_1, \ldots, w_n\}$ of all scenarios as *universe*. The universe is assumed to be *exhaustive* (containing the true state) and its scenarios to be *mutually exclusive* (no two scenarios can both be the true state). Subsets $E \subseteq \Omega$ are referred to as *events* and we say that an event E *occurs* when the true scenario belongs to E. The powerset 2^Ω of Ω is the set of all events. Uncertainty about the true scenario is most often described in terms of a *Bayesian probability distribution* quantifying the chance of occurrence of each scenario in the universe and allowing to estimate the probability of events.

Broadly speaking, probability has been defined in two ways. *Frequency probability* [Neyman, 1950; Fishburn, 1964] interprets the probability of a scenario

as the limit of its relative frequency in a large number of observations. *Subjective probability* interprets the probability of a scenario as the confidence an individual has in the occurrence of that scenario [Savage, 1972] and is subjective in that respect. Subjective probability is not necessarily based on any precise computation but is often a reasonable assessment made by an expert or knowledgeable person. The purpose of this chapter is to present a general method for estimating the probabilities of scenarios and events given statistical information about their occurrence in the past and the opinions of experts about their occurrence in the future. We will thus be concerned with both the frequency and the subjective approaches to probability. We will consider opinions in the form of arguments.

In practical decision making, the information available to a subject may be incomplete, imprecise, fragmentary, unreliable, vague or contradictory [Klir, 2006]. The notion of Bayesian probability has been enlarged to cope with such situations. The notion of *imprecise probability* is used as a generic term to cover all mathematical models which measure uncertainty without sharp numerical probabilities. The probabilities obtained using the method introduced in this chapter will in general be imprecise.

The chapter is organised as follows. Section 5.1 presents background on Bayesian probability, confidence intervals, imprecise probabilities and a special kind of imprecise probability called *belief function* [Shafer, 1976]. The problem of constructing reliable belief functions based on statistical information only is studied in Section 5.2. We call the resulting imprecise probabilities *robust belief functions.* Section 5.3 presents our model based on abstract argumentation to represent expert opinions. Section 5.4 proposes formulae for constructing belief functions in the general case where expert opinions are available, in addition to statistics. We call these belief functions *dialectical probabilities.* Together with

robust belief functions dialectical probabilities constitute the main contribution of this chapter. In Section 5.5, we formally justify the mathematical coherence of dialectical probabilities. Section 5.6 provides a detailed study of the substantive and normative properties of dialectical probabilities concerning the information-theoretic indeterminacy and uncertainty of the subject. In Section 5.7, we test the effectiveness of dialectical probabilities in trust computing. In Section 5.8 we discuss related work and conclude in Section 5.9.

5.1 Background on imprecise probabilities

5.1.1 Bayesian probability

Bayesian probability is concerned with the probability of events, that is to say, with the estimation of the likelihood of their occurrence. A *Bayesian probability distribution* is a function $P : \Omega \to [0, 1]$ that is normalised

$$\sum_{w \in \Omega} P(w) = 1$$

The probability of an event $E \subseteq \Omega$ is by abuse of notation denoted $P(E)$ and is defined as

$$P(E) = \sum_{w \in E} P(w)$$

It is often not necessary to consider all events E in 2^{Ω}. Any family of events is acceptable provided that it contains Ω and is closed under *complementation*

$$E^c = \Omega - E$$

and finite union. For any events A and B, we say that A and B are disjoint if and only if $A \cap B = \emptyset$. For any pair of such disjoint events, we have

$$P(A \cup B) = P(A) + P(B)$$

This basic property of Bayesian probability distributions is referred to as *additivity*. If A and B are not disjoint, we have instead

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

Given any real-valued function $X : \Omega \to \mathbb{R}$, the value

$$E_P(X) = \sum_{w \in \Omega} X(w)P(w)$$

is called the *expected value* of X. Given two events A and B with $P(B) > 0$, the *conditional probability* of A given B is denoted $P(A|B)$ and defined as

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

In the frequency approach to probability, a probability distribution can be constructed assuming as given a sample of N observations (often taken of size larger than 30) of the occurrence of the true scenario in the past

$$w_{i_1}, \ldots, w_{i_N} \in \Omega$$

The probability derived from the sample is called *statistical prior*, denoted $\hat{p}$ and defined by the *relative frequencies*

$$\hat{p}(w) = \frac{1}{N} \sum_{k=1}^{N} I_{\{w\}}(w_{i_k})$$

where $I_{\{w\}} : \Omega \to \{0,1\}$ is the indicator function of the event $\{w\}$, i.e.

$$I_{\{w\}}(w_{i_k}) = \begin{cases} 1 & \text{if } w_{i_k} = w \\ 0 & \text{otherwise} \end{cases}$$

In the subjective approach to probability, P is often not calculated by any precise mean but simply guessed or derived from some knowledge about the underlying mechanics of the random system or object of interest.

5.1.2 Confidence intervals for relative frequencies

In statistics [Shapiro and Gross, 1981], the reliability of an estimated relative frequency very much depends on the size N of the random sample available and is indicated using a *confidence interval.* The methods we will introduce in Sections 5.2 and 5.4 to build imprecise probabilities rely upon confidence intervals. This subsection explains how to obtain confidence intervals for the relative frequencies $\hat{p}(w)$, where $w \in \Omega$.

Each relative frequency $\hat{p}(w)$ is defined as the arithmetic average of N identical and independent random variables $I_{\{w\}}$. Each relative frequency $\hat{p}(w)$ is therefore also random. Consider now the central limit theorem [de Moivre, 1718; Laplace, 1812; Ljapunov, 1901]:

Theorem 5.1. *If $X_1, \ldots, X_N$ is a sequence of N independent and identically distributed random variables with finite expectation μ and variance $\sigma^2 > 0$, then for $N \geq 30$, the random variable*

$$Z = \sqrt{N}\frac{\overline{X} - \mu}{\sigma}$$

where $\overline{X}$ *denotes the average*

$$\overline{X} = \frac{X_1 + \ldots + X_N}{N}$$

is approximately distributed as the standard normal distribution $\mathcal{N}(0,1)$.

Here, each random variable $I_{\{w\}}$ follows a Bernoulli distribution with probability success $p(w)$ (the true probability of w), which means that it takes the value 1 with probability $p(w)$ and the value 0 with probability $1 - p(w)$. The expectation and variance of $I_{\{w\}}$ are thus simple to assess

$$\begin{aligned}
\mu &= E_p(I_{\{w\}}) = \sum_{i=1}^{n} p(w_i) I_{\{w\}}(w_i) = p(w) \\
\sigma^2 &= E_p((I_{\{w\}} - \mu)^2) = \sum_{i=1}^{n} p(w_i)(I_{\{w\}}(w_i) - p(w))^2 \\
&= (1 - p(w))p^2(w) + p(w)(1 - p(w))^2 \\
&= p(w)(1 - p(w))
\end{aligned}$$

Note that since $E_p(I_{\{w\}}) = p(w)$, we have by linearity of E_p [Neyman, 1950; Fishburn, 1964]

Theorem 5.2. *The relative frequency* $\hat{p}(w)$ *is an unbiased estimator for the true probability* $p(w)$, *i.e.*

$$E_p\left(\frac{1}{N}\sum_{k=1}^{N} I_{\{w\}}\right) = p(w)$$

The central limit theorem 5.1 allows us to compute approximate confidence intervals for relative frequencies when $N \geq 30$. These intervals are known as *Wilson (score) intervals* [Wilson, 1927].

Theorem 5.3. *Let* x *be a parameter comprised between* 0 *and* 1. *Let* $\underline{p}(w)$ *and*

$\overline{p}(w)$ *be the roots of the quadratic equation in the variable* p

$$\left(1+\frac{z^2}{N}\right)p^2-\left(2\hat{p}(w)+\frac{z^2}{N}\right)p+\hat{p}(w)^2=0$$

where we assume that $\underline{p}(w)\leq\overline{p}(w)$ *and* z *corresponds to the* $1-x/2$ *percentile of the standard normal distribution* $\mathcal{N}(0,1)$. *Then,* $\mathcal{I}(w)=\left[\underline{p}(w),\overline{p}(w)\right]$ *is an approximate* $(1-x)100\%$ *confidence interval for the true probability* $p(w)$.

This means that the probability that the true probability $p(w)$ belongs to $\mathcal{I}(w)$ is equal to $(1-x)100\%$. The proof of this result is given in the appendix. Since the bounds of the Wilson intervals are the roots of a simple quadratic equation, we can in fact easily calculate their value:

$$\underline{p}(w)=\frac{\hat{p}(w)+\frac{1}{2N}z^2-z\sqrt{\frac{\hat{p}(w)(1-\hat{p}(w))}{N}+\frac{z^2}{4N^2}}}{1+\frac{1}{N}z^2}$$

$$\overline{p}(w)=\frac{\hat{p}(w)+\frac{1}{2N}z^2+z\sqrt{\frac{\hat{p}(w)(1-\hat{p}(w))}{N}+\frac{z^2}{4N^2}}}{1+\frac{1}{N}z^2}$$

In practise, one often computes 95% confidence intervals with the 97.5% percentile value $z\approx 1.96$. Note that the Wilson scores are not defined for $\hat{p}(w)=0$ or $\hat{p}(w)=1$. In the remainder, we will consequently need to assume that each scenario $w\in\Omega$ has been observed at least one time (so that $\hat{p}(w)>0$) and that more than one scenario has been observed (so that $\hat{p}(w)\neq 1$).

5.1.3 Imprecise probabilities

It has been observed by the numerous statisticians and practitioners of probability theory that the information needed to estimate probability could in practise be deficient. There exist many forms of information deficiency. Information may for instance be incomplete, imprecise, fragmentary, unreliable, vague or

contradictory [Klir, 2006]. The term *incomplete* is used to refer to missing data or to describe information that concerns only a fragment of reality. *Imprecise* information corresponds to numerical data that are not sharp and are often represented by intervals of values. *Fragmentary* information corresponds to small, heterogeneous or disconnected pieces of evidence. *Unreliable* information denotes information that is susceptible to be erroneous or misleading. *Vague* information relates to classes of objects whose boundaries are not clearly or explicitly stated. Finally, information is said to be *contradictory* when it is composed of mutually exclusive pieces of evidence. Notably, one may say that statistical information is fragmentary, unreliable and contradictory.

Confidence interval analysis reveals that more than one probability distribution is compatible with a given statistical sample. Generally, more than one probability distribution can be compatible with the information available to a subject due to information deficiencies [Klir, 2006]. Sets of possible Bayesian probability distributions compatible with some information are called *credal sets* [Levi, 1980].

Definition 5.1. *A* ***credal set*** *is a non-empty set $\mathcal{D}$ of Bayesian probability distributions on Ω.*

The minimum of an event's probability when the distribution ranges over a credal set is called a *lower probability*, and its maximum an *upper probability*.

Definition 5.2. *The* ***lower*** *and* ***upper probabilities*** *associated with a credal set $\mathcal{D}$ are the functions $\underline{P} : 2^{\Omega} \rightarrow [0,1]$ and $\overline{P} : 2^{\Omega} \rightarrow [0,1]$ defined $\forall E \subseteq \Omega$ as*

$$\underline{P}(E) = \inf\{P(E) \mid P \in \mathcal{D}\} \quad \overline{P}(E) = \sup\{P(E) \mid P \in \mathcal{D}\}$$

For every probability P in the credal set $\mathcal{D}$ and any event E, it holds that

$$\underline{P}(E) \leq P(E) \leq \overline{P}(E)$$

and thus lower and upper probabilities are referred to as *imprecise probabilities*. For a comprehensive presentation of imprecise probabilities, the reader may refer to [Walley, 1991; 1996; 2000]. Imprecise probabilities are appropriate for representing uncertainty in a variety of problems where the available information is incomplete [Zaffalon, 2002a], imprecise [Walley and Fine, 1979; Zaffalon, 2002b; Kozine and Utkin, 2005], vague [Zadeh, 1978; de Cooman, 1997; Walley and de Cooman, 2001] or contradictory [Walley, 1996]. In general, confidence intervals are not imprecise probabilities [Walley, 1991].

Note that Bayesian probability is a special case of imprecise probability. When the credal set is a singleton $\mathcal{D} = \{P\}$, the lower and upper probabilities are equal to P, so that for every event E

$$\underline{P}(E) = \overline{P}(E) = P(E)$$

Thus Bayesian probability distributions exactly correspond to the imprecise probabilities that are precise. In general, however, lower and upper probabilities are distinct. Various types of imprecise probabilities which extend the Bayesian approach have been proposed over the past decades, such as notably

- ϵ-contamination models [Huber, 1981; Berger and Berliner, 1986] in robust statistics,
- possibility and necessity measures [Zadeh, 1978; Dubois and Prade, 1988; 1995; de Cooman, 1997] and belief and plausibility functions [Shafer, 1976; Sentz and Scott, 2002] in generalised uncertainty and information theory and artificial intelligence,

- Choquet capacities [Choquet, 1953] and imprecise probabilities induced by convex or norm-bounded credal sets [Ben-Tal and Nemirovski, 2001; Bertsimas and Brown, 2005; Natarajan *et al.*, 2005] in robust optimisation.

In the rest of this chapter, we will be essentially concerned with belief functions. Belief functions form a subclass of Choquet capacities, which in turn belong to the family of monotone measures, as we shall see below.

A *monotone measure* is a function $\mu : 2^\Omega \rightarrow [0,1]$ such that $\mu(\emptyset) = 0$, $\mu(\Omega) = 1$ and for every pair of events (A, B), if $A \subseteq B$ then $\mu(A) \leq \mu(B)$. Given an integer $k \geq 2$, a *Choquet capacity of order k* [Choquet, 1953] is a monotone measure μ that satisfies the inequalities

$$\mu\left(\bigcup_{j=1}^{k} A_j\right) \geq \sum_{\substack{K \subseteq \{1,\ldots,k\} \\ K \neq \emptyset}} (-1)^{|K|+1} \mu\left(\bigcap_{j \in K} A_j\right)$$

for all families of k events $A_1, \ldots, A_k \subseteq \Omega$.

Property 5.1. *Every Choquet capacity of order $k > 2$ is also a Choquet capacity of order $k' = 2, 3, \ldots, k$.*

The proof of this result by [Choquet, 1953] is given in the appendix. Choquet capacities of order 2 satisfy by definition the inequality

$$\mu(A_1 \cup A_2) \geq \mu(A_1) + \mu(A_2) - \mu(A_1 \cap A_2)$$

for all pairs of events (A_1, A_2) and are according to the previous property the most general capacities. These capacities form a rather large class of lower probabilities.

Property 5.2. *Every Choquet capacity of order 2 is a lower probability.*

The proof of this result by [Walley, 1991] is given in the appendix. Monotone

measures μ which are capacities of order k for all integers $k \geq 2$ are called Choquet capacities of order ∞ (infinite order) and are the least general ones. These capacities are more commonly known as *belief functions*. For any belief function Bel, there exists a credal set $\mathcal{D}$ such that, for every event E,

$$Bel(E) = \inf\{P(E) \mid P \in \mathcal{D}\}$$

The corresponding upper envelope is called the *plausibility function* Pl and is defined for every event E as

$$Pl(E) = 1 - Bel(E^c)$$

Consequently, for that same credal set $\mathcal{D}$, we may write

$$Pl(E) = \sup\{P(E) \mid P \in \mathcal{D}\}$$

Fortunately, there exists an equivalent, simpler and more constructive characterisation of belief and plausibility functions, that we overview next.

5.1.4 Belief and plausibility functions

This section introduces belief functions, which form the basis of the so-called *Dempster-Shafer* or *evidence theory* [Shafer, 1976]. Note that, in this thesis, we are not concerned with the problem of evidence combination, nor do we need to rely on the interpretation given to belief functions in the Dempster-Shafer theory. Therefore, this section is not an introduction to Dempster-Shafer theory. The reader interested in evidence combination may refer to the survey by [Sentz and Scott, 2002].

As seen in Section 5.1.3, belief functions form the most specific class of Cho-

quet capacities and as such constitute imprecise probabilities. Belief and plausibility functions admit an equivalent but more constructive definition, which explains the important role they have played in the field of artificial intelligence, especially regarding the design of expert systems. Belief functions can be defined in terms of evidence mass function.

Definition 5.3. *An **evidence mass function** is a function* $m : 2^{\Omega} \to [0,1]$ *that verifies* $m(\emptyset) = 0$ *and*

$$\sum_{X \subseteq \Omega} m(X) = 1$$

Definition 5.4. *A **belief function** is a function* $Bel : 2^{\Omega} \to [0,1]$ *such that, for every* $E \subseteq \Omega$,

$$Bel(E) = \sum_{X \subseteq E} m(X)$$

for some evidence mass function m.

Definition 5.5. *Given an evidence mass function* m *and its associated belief function Bel, the corresponding **plausibility function*** $Pl : 2^{\Omega} \to [0,1]$ *is the dual of Bel defined for every event* $E \subseteq \Omega$ *as* $Pl(E) = 1 - Bel(E^c)$, *or equivalently*

$$Pl(E) = \sum_{X \subseteq \Omega, X \cap E \neq \emptyset} m(X)$$

The events for which $m(E) \neq 0$ are called *focal sets.* When all the focal sets of an evidence mass function are singletons, the associated belief and plausibility functions are equal and constitute a Bayesian probability distribution. Thus, Bayesian probability distributions are special belief functions, which in turn are imprecise probabilities.

5.1.5 Coherence of imprecise probabilities

When imprecise probabilities are employed for statistical inference or decision making, it is useful to interpret them as rational betting prices on events, called *coherent previsions* [Walley, 1991]. Let us now introduce the notions of prevision and coherence.

A *gamble* is any function $X : \Omega \rightarrow \mathbb{R}$ and represents a ticket for a lottery, where the true scenario w is revealed and a price $X(w)$ in units of money or utility is won by the participant. The true scenario is not known in advance by the subject, so the outcome of the lottery is uncertain, hence the name "gamble". Each event E can in fact be seen as a lottery with a price equal to 1 if E occurs, and 0 if E does not occur. Formally, it is easy to see that the gamble corresponding to an event E is the *indicator function* $I_E : \Omega \rightarrow \{0,1\}$ of the set E, defined $\forall w \in \Omega$ as

$$I_E(w) = \begin{cases} 1 & \text{if } w \in E \\ 0 & \text{otherwise} \end{cases}$$

By convenience, we use de Finetti's notation and systematically write E instead of I_E. For any gambles X and Y, we write $X \leq Y$ if and only if $\forall w \in \Omega$ we have $X(w) \leq Y(w)$. We define $\max(X) = \max\{X(w) \mid w \in \Omega\}$, $\min(X) = \min\{X(w) \mid w \in \Omega\}$. Note, in particular, that

- if $E \subset \Omega$, then $\min(E) = \min_{w \in \Omega} I_E(w) = 0$
- $\min(\Omega) = \min_{w \in \Omega} I_\Omega(w) = 1$
- if $E \supset \emptyset$, then $\max(E) = \max_{w \in \Omega} I_E(w) = 1$
- $\max(\emptyset) = \max_{w \in \Omega} I_\emptyset(w) = 0$

We define the operations of addition $+$, subtraction $-$ and multiplication $*$

of gambles in exactly the same way as for real-valued functions, for instance, $\forall w \in \Omega, (X+Y)(w) = X(w)+Y(w)$. Finally, for any constant $c \in \mathbb{R}$, we denote with c the gamble X such that, $\forall w \in \Omega, X(w) = c$.

Given a non-empty set $\mathcal{X}$ of gambles, a *lower prevision* is a function $\underline{P} : \mathcal{X} \to \mathbb{R}$ and an *upper prevision* a function $\overline{P} : \mathcal{X} \to \mathbb{R}$. For any gamble $X \in \mathcal{X}$, the value $\underline{P}(X)$ is interpreted as the maximum price at which the subject would accept to buy the lottery ticket X, and $\overline{P}(X)$ as the minimum price at which he would accept to sell the ticket X.

Walley proposes two fundamental principles to characterise the rationality of a gambler, which together define the notion of coherence. The first principle states that the subject should avoid sure loss, i.e. never be sure to lose money by engaging in any collection of lotteries. We can formalise this principle using the notion of *marginal gamble*, defined for any gamble X as $G(X) = X - \underline{P}(X)$. $G(X)$ corresponds to a transaction in which the subject pays $\underline{P}(X)$ to obtain X. Since $\underline{P}(X)$ is the maximal buying price for X, this transaction is still acceptable. Formally, we say that

Definition 5.6. *$\underline{P}$ **avoids sure loss** if and only if*

$$\max\left(\sum_{j=1}^{k} G(X_j)\right) \geq 0$$

whenever the gambles $X_1, ..., X_k$ are in $\mathcal{X}$ and $k \geq 1$.

The second principle states that the subject's prices should be consistent as a whole. If, for some $x \in \mathbb{R}$, the gamble $X - x$ is greater or equal to a positive combination of acceptable gambles, then $X - x$ should also be acceptable, that is to say $\underline{P}(X) \geq x$. This should hold for any such x, which implies that $\underline{P}(X)$ should be greater or equal to the quantity $\underline{E}(X)$ called *natural extension* of X.

Definition 5.7. *For any gamble $X : \Omega \rightarrow \mathbb{R}$, the* ***natural extension*** *of X is*

$$\underline{E}(X) = \sup\{x \in \mathbb{R} \mid \exists X_1, \ldots, X_k \in \mathcal{X}, \exists \lambda_1, \ldots, \lambda_k \geq 0 : X - x \geq \sum_{j=1}^{k} \lambda_j G(X_j)\}$$

It is proved in [Walley, 1991] that every lower prevision that avoids sure loss admits a natural extension. $\underline{E}(X)$ can be computed efficiently by linear programming [Dantzig *et al.*, 1955]. $\underline{E}(X)$ depends on the marginal gambles $G(X_j) = X_j - \underline{P}(X_j)$ and thus it depends on the lower previsions $\underline{P}(X_j)$ of all the gambles $X_j \in \mathcal{X}$. By taking $k = 1$, $X_1 = X$, $\lambda_1 = 1$ and $x = \underline{P}(X)$ in the definition of $\underline{E}(X)$, one can see that $\underline{E}(X) \geq \underline{P}(X)$ for every gamble $X \in \mathcal{X}$. One may conclude from this that the lower prevision of a rational gambler should not only avoid sure loss, but also coincide with its natural extension.

Definition 5.8. *$\underline{P}$ is* ***coherent*** *if and only if it avoids sure loss and coincides with its natural extension on $\mathcal{X}$, i.e., $\forall X \in \mathcal{X}, \underline{E}(X) = \underline{P}(X)$.*

Definition 5.9. *An upper previsions $\overline{P}$ is* ***coherent*** *when it is the conjugate of some coherent lower prevision $\underline{P}$, i.e., $\forall X \in \mathcal{X}, \overline{P}(X) = -\underline{E}(-X)$.*

It is particularly remarkable (although far from being obvious) that

Theorem 5.4. *When $\mathcal{X} \subseteq 2^{\Omega}$, lower and upper previsions are lower and upper probabilities. Conversely, every pair of lower and upper probabilities are coherent lower and upper previsions.*

This very fundamental result was proved by [Walley, 1991]. The result allows to interpret imprecise probabilities as coherent previsions.

To summarise this background section, we may simply say that belief and plausibility functions are imprecise probabilities and can be interpreted as rational betting prices on events. This interpretation is also valid for Bayesian probability, since this is a special case of belief function.

5.2 Robust belief and plausibility functions

The objective of this section is to show how to construct coherent previsions, when the only information available to the subject is of statistical nature. These previsions are defined in the form of statistically *robust belief functions.* We assume given a sample of $N \geq 30$ observations

$$w_{i_1}, w_{i_2}, \ldots, w_{i_N} \in \Omega$$

We assume that all the conditions required for the construction of Wilson score intervals are satisfied (most importantly, that each scenario has occurred at least once and that more than one scenario has occurred, see at the end of Section 5.1.2). As in Section 5.1.2, the estimated relative frequency $\hat{p}(w)$ of each scenario w defines the statistical prior $\hat{p} : \Omega \to \mathbb{R}$ where

$$\hat{p}(w) = \frac{1}{N} \sum_{k=1}^{N} I_{\{w\}}(w_{i_k})$$

The relative frequencies $\hat{p}(w)$ are unbiased estimators of the true probabilities $p(w)$. The Wilson score intervals are given by $\mathcal{I}(w) = [\underline{p}(w), \overline{p}(w)]$ as in Section 5.1.2 and are $(1 - x)100\%$ confidence intervals. The bounds of confidence intervals do not define imprecise probabilities [Walley, 1991]. We consequently search for specific confidence intervals having the property of being coherent.

To restrict our search amongst the class of all imprecise probabilities, we may impose a number of requirements. One idea suggested by robust Bayes analysts [Blum and Rosenblatt, 1967; Huber, 1981; Berger and Berliner, 1986] is to impose that the credal set associated with the imprecise probability consists of a set of distributions $\mathcal{D}$ close to a given Bayesian prior probability distribution π. Concretely, the credal sets considered in robust Bayes analysis are those of

the form

$$\mathcal{D} = \{P \mid P = (1-\epsilon)\pi + \epsilon q,\ q \in \mathcal{Q}\}$$

where $\epsilon \in [0,1]$ is a fixed parameter, and $\mathcal{Q}$ a subset of the set $\mathcal{P}$ of all probability distributions on Ω. Such credal sets are called *ϵ-contamination models* (sometimes also gross-error models or linear-vacuous mixtures), where the prior π is contaminated by the family $\mathcal{Q}$ with degree ϵ. Basically, such families are used for modeling random processes which follow the distribution π with frequency $1-\epsilon$ and some unknown probability in $\mathcal{Q}$ the rest of the time [Walley, 1991]. We will consider that we have no information concerning $\mathcal{Q} \subseteq \mathcal{P}$ and thus in the remainder make the conservative choice of the maximal family $\mathcal{Q} = \mathcal{P}$. The ϵ-contaminations obtained for such a credal set will be by convenience referred to as *maximal* ϵ-contaminations. A natural candidate here for the prior π is the statistical prior $\hat{p}$, as the relative frequencies are unbiased estimators for the true scenario probabilities. This now restricts our search to credal sets of the form

$$\mathcal{D} = \{P \mid P = (1-\epsilon)\hat{p} + \epsilon q, q \in \mathcal{P}\}$$

In order to choose an ϵ, we impose the constraint that the interval $\mathcal{J}(w) = [\underline{P}(w), \overline{P}(w)]$ where

$$\underline{P}(w) = \inf\{P(w) | P = (1-\epsilon)\hat{p} + \epsilon q, q \in \mathcal{P}\}$$

$$\overline{P}(w) = \sup\{P(w) | P = (1-\epsilon)\hat{p} + \epsilon q, q \in \mathcal{P}\}$$

is an approximate $(1-x)100\%$ confidence interval. Since, for every scenario w, the Wilson score interval $\mathcal{I}(w) = \left[\underline{p}(w), \overline{p}(w)\right]$ is a $(1-x)100\%$ approximate confidence interval, $\mathcal{I}(w) \subseteq \mathcal{J}(w)$ as long as the following two inequalities are

satisfied

$$\underline{p}(w) \geq (1-\epsilon)\hat{p}(w), \quad \overline{p}(w) \leq (1-\epsilon)\hat{p}(w) + \epsilon$$

If these conditions are fulfilled, then $\mathcal{J}(w)$ must a fortiori be a $(1-x)100\%$ approximate confidence interval. The minimal value of ϵ for which $\mathcal{I}(w) \subseteq \mathcal{J}(w)$ holds for every $w \in \Omega$ is

Definition 5.10.

$$\epsilon = \max_{w \in \Omega} \max \left(\frac{\hat{p}(w) - \underline{p}(w)}{\hat{p}(w)}, \frac{\overline{p}(w) - \hat{p}(w)}{1 - \hat{p}(w)} \right)$$

Property 5.3. *The value given in definition 5.10 is the minimum value of ϵ for which $\mathcal{I}(w) \subseteq \mathcal{J}(w)$ holds for every $w \in \Omega$.*

Proof. The inclusions hold if and only if, for every $w \in \Omega$

$$\begin{cases} \underline{p}(w) \geq (1-\epsilon)\hat{p}(w) \\ \overline{p}(w) \leq (1-\epsilon)\hat{p}(w) + \epsilon \end{cases} \Leftrightarrow \begin{cases} \epsilon \geq 1 - \frac{\underline{p}(w)}{\hat{p}(w)} = \frac{\hat{p}(w) - \underline{p}(w)}{\hat{p}(w)} \\ \frac{\overline{p}(w) - \hat{p}(w)}{1 - \hat{p}(w)} \leq \epsilon \end{cases}$$

□

We can show that

Property 5.4. *Let $\underline{P}$ and $\overline{P}$ be the imprecise probabilities corresponding (via definition 5.2 of Section 5.1.3) to the credal set given by the maximal ϵ-contamination of $\hat{p}$ with level of contamination ϵ as indicated in definition 5.10. For every $E \subseteq \Omega$,*

$$\underline{P}(E) = (1-\epsilon)\hat{p}(E) + \epsilon \min(E), \quad \overline{P}(E) = (1-\epsilon)\hat{p}(E) + \epsilon \max(E)$$

Proof. This derives directly from the fact that $\inf\{q(E) \mid q \in \mathcal{P}\} = \min(E)$ and $\sup\{q(E) \mid q \in \mathcal{P}\} = \max(E)$. □

One can easily check that

Property 5.5. *$\underline{P}$ and $\overline{P}$ as defined in property 5.4 are belief and plausibility functions respectively.*

Proof. Consider the function $m : 2^{\Omega} \rightarrow [0,1]$ such that

$$m(X) = (1-\epsilon)\hat{p}(X)\delta(|X|,1) + \epsilon\delta(|X|,|\Omega|)$$

where δ is the Kronecker symbol defined for every integers i and j by

$$\delta(i,j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

The function m is positive and is such that $m(\emptyset) = 0$ and

$$\sum_{X \subseteq \Omega} m(X) = (1-\epsilon)\sum_{w \in \Omega} \hat{p}(w) + \epsilon = 1$$

So m is an evidence mass function [1]. Moreover, we have for every $E \subseteq \Omega$

$$\begin{aligned} \sum_{X \subseteq E} m(X) &= \sum_{w \in E} m(\{w\}) + \epsilon \min(E) = (1-\epsilon)\sum_{w \in E} \hat{p}(w) + \epsilon \min(E) \\ &= (1-\epsilon)\hat{p}(E) + \epsilon \min(E) = \underline{P}(E) \quad \text{(by property 5.4)} \end{aligned}$$

By definition 5.4 in Section 5.1.4, $\underline{P}$ is consequently the belief function *Bel* associated with the mass function m. By definition 5.5, the plausibility function

[1] Note that m admits as focal sets only singletons and Ω.

associated with m is given for every $E \subseteq \Omega$ by

$$\begin{aligned} Pl(E) &= 1 - Bel(E^c) = 1 - \underline{P}(E^c) \\ &= 1 - ((1-\epsilon)\hat{p}(E^c) + \epsilon \min(E^c)) \quad \text{(by property 5.4)} \\ &= 1 - \hat{p}(E^c) + \epsilon\hat{p}(E^c) - \epsilon \min(E^c) \\ &= 1 - (1 - \hat{p}(E)) + \epsilon(1 - \hat{p}(E)) - \epsilon \min(E^c) \quad \text{(since } \hat{p} \text{ is Bayesian)} \\ &= (1-\epsilon)\hat{p}(E) + \epsilon - \epsilon \min(E^c) \end{aligned}$$

We have

$$\min(E^c) = \begin{cases} 1 & \text{if } E^c = \Omega \\ 0 & \text{otherwise} \end{cases} = \begin{cases} 1 & \text{if } E = \emptyset \\ 0 & \text{otherwise} \end{cases} = 1 - \max(E)$$

so

$$\begin{aligned} Pl(E) &= (1-\epsilon)\hat{p}(E) + \epsilon - \epsilon(1 - \max(E)) = (1-\epsilon)\hat{p}(E) + \epsilon \max(E) \\ &= \overline{P}(E) \quad \text{(by property 5.4)} \end{aligned}$$

which means that $\overline{P}$ is the plausibility function associated with m. □

As a result of this property, assuming given a statistical sample of $N \geq 30$ observations and a level of statistical confidence $(1-x)100\%$, we have proposed formulae for the construction of imprecise probabilities ($Bel = \underline{P}$ and $Pl = \overline{P}$ as in property 5.4) whose credal set is a maximal ϵ-contamination of the unbiased statistical prior $\hat{p}$. We have defined $\underline{P}$ and $\overline{P}$ in terms of ϵ and chosen ϵ so that $\mathcal{J}(w) = [\underline{P}(\{w\}), \overline{P}(\{w\})]$ is an approximate $(1-x)100\%$ confidence interval of the true relative frequencies $p(w)$ (for every $w \in \Omega$). We present next an alternative formulation of Bel and Pl in terms of the following notion.

Definition 5.11. *The **epistemic risk-aversion parameter** $K \in \mathbb{R}^+$ is defined as the constant such that*

$$\epsilon = \frac{K}{K+1}$$

We have for the choice of ϵ made earlier in definition 5.10

Property 5.6.

$$K = \frac{\epsilon}{1-\epsilon}$$

The imprecise probabilities derived earlier can be expressed equivalently in terms of K rather than ϵ as follows:

Property 5.7. $\forall E \subseteq \Omega$

$$Bel(E) = \frac{1}{K+1}\left(\hat{p}(E) + K\min(E)\right), \quad Pl(E) = \frac{1}{K+1}\left(\hat{p}(E) + K\max(E)\right)$$

We refer to Bel and Pl represented this way as *robust belief and plausibility functions.*

As an example, let us consider again the umbrella problem. Suppose that over the last 30 days, there has been 12 cloudy days and 18 sunny days. Thus, scenario $w_1 = \{\text{clouds}\}$ has occurred 12 times and scenario $w_2 = \{\neg\text{clouds}\}$ has occurred 18 times. The statistical prior $\hat{p} : \Omega \to [0,1]$ is then defined by

$$\hat{p}(w_1) = \frac{12}{30} = 0.4, \quad \hat{p}(w_2) = \frac{18}{30} = 0.6$$

Assume we want to build the robust belief and plausibility function with a statistical confidence level of 95%, i.e. we choose $x = 5\%$. The $1 - x/2$ percentile of the standard normal distribution is $z = 1.96$. Since the number of observations

is $N = 30$, we have the following bounds for the Wilson score intervals

$$\underline{p}(w_1) = \frac{\hat{p}(w_1) + \frac{1}{2N}z^2 - z\sqrt{\frac{\hat{p}(w_1)(1-\hat{p}(w_1))}{N} + \frac{z^2}{4N^2}}}{1 + \frac{1}{N}z^2} = 0.246$$

$$\overline{p}(w_1) = \frac{\hat{p}(w_1) + \frac{1}{2N}z^2 + z\sqrt{\frac{\hat{p}(w_1)(1-\hat{p}(w_1))}{N} + \frac{z^2}{4N^2}}}{1 + \frac{1}{N}z^2} = 0.577$$

$$\underline{p}(w_2) = \frac{\hat{p}(w_2) + \frac{1}{2N}z^2 - z\sqrt{\frac{\hat{p}(w_2)(1-\hat{p}(w_2))}{N} + \frac{z^2}{4N^2}}}{1 + \frac{1}{N}z^2} = 0.423$$

$$\overline{p}(w_2) = \frac{\hat{p}(w_2) + \frac{1}{2N}z^2 + z\sqrt{\frac{\hat{p}(w_2)(1-\hat{p}(w_2))}{N} + \frac{z^2}{4N^2}}}{1 + \frac{1}{N}z^2} = 0.754$$

So we obtain a level of contamination ϵ

$$\begin{aligned}
\epsilon &= \max_{w\in\Omega} \max\left(\frac{\hat{p}(w) - \underline{p}(w)}{\hat{p}(w)}, \frac{\overline{p}(w) - \hat{p}(w)}{1 - \hat{p}(w)}\right) \\
&= \max\left(\frac{\hat{p}(w_1) - \underline{p}(w_1)}{\hat{p}(w_1)}, \frac{\overline{p}(w_1) - \hat{p}(w_1)}{1 - \hat{p}(w_1)}, \frac{\hat{p}(w_2) - \underline{p}(w_2)}{\hat{p}(w_2)}, \frac{\overline{p}(w_2) - \hat{p}(w_2)}{1 - \hat{p}(w_2)}\right) \\
&= \max(0.385, 0.295, 0.295, 0.385) = 0.385
\end{aligned}$$

and the epistemic risk-aversion parameter

$$K = \frac{\epsilon}{1 - \epsilon} = 0.626$$

If we denote by E_{clouds} the event $\{w_1\}$ and E_{sunny} the event $\{w_2\}$, the robust belief and plausibility functions are such that

- $Bel(E_{clouds}) = \frac{1}{K+1}(\hat{p}(E_{clouds}) + K\min(E_{clouds})) = 0.246$
- $Pl(E_{clouds}) = \frac{1}{K+1}(\hat{p}(E_{clouds}) + K\max(E_{clouds})) = 0.631$
- $Bel(E_{sunny}) = \frac{1}{K+1}(\hat{p}(E_{sunny}) + K\min(E_{sunny})) = 0.369$
- $Pl(E_{sunny}) = \frac{1}{K+1}(\hat{p}(E_{sunny}) + K\max(E_{sunny})) = 0.754$

Robust belief and plausibility functions are imprecise probabilities. We can quantify this imprecision as follows.

Definition 5.12. *The **imprecision** of the probability of an event E is*

$$\Delta(E) = \overline{P}(E) - \underline{P}(E)$$

Trivially:

Property 5.8. *We have $\Delta(\emptyset) = 0$, $\Delta(\Omega) = 1$ and for any event E distinct from $\emptyset$ and Ω,*

$$\Delta(E) = \epsilon = \frac{K}{K+1}$$

In order to further reduce his *indeterminacy* (the imprecision ϵ), the subject needs additional information. In the remainder of this chapter, we consider the case in which this information takes the form of expert opinions.

Note that we have proposed in this section a formula for estimating the risk aversion parameter K when no expert opinion is available and the decision maker wants to guarantee with $(1-x)100\%$ confidence that the true probability of the scenarios w will fall within the lower and upper probabilities $\underline{P}(\{w\})$ and $\overline{P}(\{w\})$. Sometimes, the subject may not know what statistical level of confidence to choose (namely which x to choose as in the Wilson scores theorem 5.3) and may thus wish to take $K = \epsilon/(1-\epsilon)$ but estimate ϵ without relying on the Wilson score intervals and the choice of x. In the remainder of this chapter, we will assume given a generic ϵ and continue to consider that $K = \epsilon/(1-\epsilon)$. We will then show how in two different applications, viz. trust computing (see Section 5.7) and stock trading (see Section 6.8 in Chapter 6), an adequate value of ϵ can be heuristically obtained.

5.3 Modeling expert opinions by abstract argumentation

By an *expert*, we mean a person, device or system capable of elaborating *arguments* understood as justified claims [Krause *et al.*, 1995; Fox, 2003]. Expert opinions are modeled as sets of arguments. All arguments and their conflicts are here modeled using abstract argumentation [Dung, 1995]. As in Chapter 4, we restrict our attention to two types of arguments only. These are *forecast* and *mitigation arguments*. A forecast argument is a justified claim according to which some event $E \subseteq \Omega$ will occur, or, equivalently, since events are seen as disjunctions of scenarios, a claim according to which the true scenario belongs to E. Mitigation arguments are used as in Chapter 4 to cast doubt on the validity of justifications. Mitigation arguments can thus attack both forecast and mitigation arguments.

Formally, given a universe Ω, we model the opinions of experts and their conflicts in terms of

- a set of arguments Arg composed of a set A of forecast arguments and a set M of mitigation arguments (the sets A and M need not be disjoint),
- a binary relation of attack between arguments $att \subseteq M \times Arg$, all attacks coming from the mitigation arguments, and
- a support function $X_\bullet : A \rightarrow 2^\Omega - \{\emptyset\}$ which associates to each forecast argument $a \in A$ an event X_a supported by a.

For instance, consider the universe $\Omega = \{w_1, w_2\}$ and the scenarios $w_1 = \{clouds\}$ and $w_2 = \{\neg clouds\}$ obtained for the umbrella problem (see decision table in Figure 3.7, Section 3.6). Let $A = \{a_1, a_2\}$ be the set of forecast arguments, where

- a_1 = 'there will be no clouds, because the air today is dry',
- a_2 = 'there will be clouds, because the wind is bringing humid air',

and $M = \{m_1\}$ be the set of mitigation arguments, where

- m_1 = 'the humidity of the air in the wind is not so high, according to the measurements made by the Meteosat satellite'.

The set of arguments is $Arg = \{a_1, a_2, m_1\}$ and the attack relationship between the arguments in this set is simply $att = \{(m_1, a_2)\}$. Finally, the support function is such that

- $X_\bullet(a_1) = X_{a_1} = \{w_2\}$ and
- $X_\bullet(a_2) = X_{a_2} = \{w_1\}$.

Since Arg is a set and not a list of arguments, arguments repeated by several experts are not taken into account more than once. Indeed, we want to avoid redundant information. Each new argument should add some information to the analysis of the situation. We refer to this important feature of arguments as *complementarity*.

Concerning the attacks, note that only mitigation arguments can attack other arguments. However, mitigation arguments can have a double status and serve as forecast arguments. Therefore, it seems that one would only need one kind of argument here, namely mitigation arguments. The reason why we distinguish forecast arguments from the others is that only forecast arguments need to be aggregated in order to estimate probabilities.

The support function $X_\bullet$ only allows forecasting arguments to support the occurrence of an event and not to be against it. This restriction is illusory, as any argument against an event E can be equivalently substituted by an argument supporting the complement $E^c = \Omega - E$ of that event. Note also that the event

X_a supported by a forecast argument a cannot be empty, since the event $\emptyset$ cannot occur.

Finally, notice that conflicts and disagreements between expert opinions are two distinct notions. Conflicts correspond in our model to attacks and serve to cast doubt on the justifications of arguments. In the terminology of argumentation, our attacks are used to *undermine* arguments. Disagreements on the opposite relate to the conclusions of arguments. As an example, suppose Expert 1 gives an argument supporting event E and Expert 2 gives an argument supporting event E'. We say that Expert 1 and 2

- totally disagree with each other when $E' \subseteq E^c$ (or equivalently $E \subseteq E'^c$),
- partially agree when $E \cap E' \neq \emptyset$,
- totally agree when $E = E'$.

If $E \subset E'$, then Expert 2 would partially agree with Expert 1 and we may say that the argument of Expert 1 is overly precise for Expert 2. Conversely, Expert 2 would say that the argument of Expert 1 is too vague. In the terminology of argumentation, a disagreement between arguments supporting respectively E and E^c would be called a *rebuttal*. Conflicts and disagreements are totally disentangled, so that two experts may

- conflict but not disagree (they come to the same conclusion for conflicting reasons),
- not conflict but disagree (they derive contradictory conclusions from the same empirical theory),
- both conflict and disagree (they have conflicting reasons and draw contradictory conclusions) and
- neither conflict nor disagree.

Our model allows to capture both conflicts (unreliability) and disagreements (contradictions) between opinions, as well as their vagueness and imprecision.

5.4 Combining statistics and expert opinions

We now consider the problem of constructing imprecise probabilities given both statistical data and expert opinions. Formally, we assume given

- a finite universe $\Omega = \{w_1, ..., w_n\}$ of n disjoint and exhaustive scenarios,
- a statistical sample of N observations $w_{i_1}, \ldots, w_{i_N} \in \Omega$,
- a generic constant $\epsilon > 0$ such that $\epsilon \leq 1$ and the corresponding epistemic risk-aversion parameter $K = \epsilon/(1-\epsilon)$,
- an abstract argumentation framework $F = (Arg, att)$ where Arg is a set of expert arguments partitioned into forecast arguments A and mitigation arguments M and $att \subseteq M \times Arg$ an attack relation between arguments,
- a support function $X_\bullet : A \to 2^\Omega - \{\emptyset\}$ associating to every forecast argument a an event X_a supported by a, and
- a measure of argument strength $s_F : Arg \to [0,1]$.

We will deploy the notion of strength given in Chapter 4. From the universe Ω and statistical sample $w_{i_1}, \ldots, w_{i_N}$, we derive the unbiased statistical prior $\hat{p} : \Omega \to [0,1]$ as in Section 5.1.1.

Expert opinions bring additional information to the subject which is needed to reduce his indeterminacy $\Delta(E) = \epsilon$ (as discussed at the end of Section 5.2). It is thus important to quantify this information. This can arguably be done in several ways. A subject may perfectly believe that expert opinions are highly valuable or that they have no value at all (with respect to the statistical data

available). Both attitudes are perfectly acceptable. A subject may also choose to assign equal value to all forecast arguments, or wish to discriminate between them, for instance by looking at their individual strength, etc. We will

- adopt the convention that the amount of statistical information counts as 1 unit,
- make the assumptions that the information contained in a forecast argument is proportional to its strength and is equal to $V_A * s_F(a)$ units, where $V_A \in \mathbb{R}^+$ is a new subjective parameter interpreted as the *informational value of forecast arguments per unit of strength*, and
- assume that, by complementarity, the informational value of forecast arguments is additive.

Under these assumptions, we define

Definition 5.13. *The **amount of information of the forecast arguments** is*

$$I_A = V_A \sum_{a \in A} s_F(a)$$

Definition 5.14. *The **total amount of information** available is*

$$I = I_A + 1$$

Like the parameter ϵ, we will present several ways for helping the subject to choose V_A in practical applications, or, in other words, to give expert opinions the right level of importance. Basically, V_A can be set by

- estimating the reduction of the frequency of error ϵ due to the arguments,
- imposing a constraint on the amplitude of the bias introduced by the arguments on the statistical prior $\hat{p}$, or

- plotting the variations of the long-term average performance of the decision maker (using a realistic benchmark) with respect to V_A and choosing for V_A the value that yields the optimal performance on the benchmark.

The expressions proposed in section 5.2 to construct robust belief and plausibility functions from statistical data can be generalised to the case where expert opinions are also available. The generalisation of this result yields to a another formula that is what really embodies the contribution of this chapter. We define the *dialectical lower and upper probabilities* $\underline{P}$ and $\overline{P}$ as

Definition 5.15. *The **lower and upper dialectical probabilities** $\underline{P}$ and $\overline{P}$ are defined for every event $E \subseteq \Omega$ by*

$$\underline{P}(E) = \frac{1}{K+I}\left[\hat{p}(E) + V_A \sum_{a \in A} s_F(a)\hat{p}(E|X_a) + K \min(E)\right]$$

$$\overline{P}(E) = \frac{1}{K+I}\left[\hat{p}(E) + V_A \sum_{a \in A} s_F(a)\hat{p}(E|X_a) + K \max(E)\right]$$

Each term of the form $\hat{p}(E|X_a)$ corresponds to the prior conditional probability of E given the event supported by a. Recall that this conditional probability is given by

$$\hat{p}(E|X_a) = \frac{\hat{p}(E \cap X_a)}{\hat{p}(X_a)}$$

Each conditional probability $\hat{p}(E|X_a)$ is weighted by the informational value $V_A s_F(a)$ of the forecast argument a. Therefore, each forecast argument adds a bias to the statistical prior that is proportional to its informational value. It is easy to see that

Property 5.9. *When $V_A = 0$ or when $A = \emptyset$, the dialectical lower and upper probabilities are equal to the robust belief and plausibility functions respectively.*

Dialectical lower and upper probabilities thus generalise robust belief and plausibility functions.

As an illustration, let us consider again the umbrella problem. The arguments $Arg = \{a_1, a_2, m_1\}$ and attacks $att = \{(m_1, a_2)\}$ are as in Section 5.3. The event supported by a_1 is $X_{a_1} = \{w_2\} = E_{sunny}$ and the event supported by a_2 is $X_{a_2} = \{w_1\} = E_{clouds}$. The strength of a_1 is $s_F(a_1) = 1$ since a_1 is not attacked and the strength of a_2 is $s_F(a_2) = 0.25$ since a_2 is attacked by the argument m_1 and a_2 does not counter-attack m_1. Choosing for instance $V_A = 2$, the total amount of information of the forecast arguments is

$$I_A = V_A \sum_{a \in A} s_F(a) = 2.5$$

The total amount of information available is $I = I_A + 1 = 3.5$. The epistemic risk-aversion parameter chosen is $K = 0.626$ as in Section 5.2, thus we still assume a 95% statistical confidence level. Observe that we simply have

- $\hat{p}(E_{clouds}|X_{a_1}) = 0$
- $\hat{p}(E_{clouds}|X_{a_2}) = 1$
- $\hat{p}(E_{sunny}|X_{a_1}) = 1$
- $\hat{p}(E_{sunny}|X_{a_2}) = 0$

The dialectical lower and upper probabilities are such that

- $\underline{P}(E_{clouds}) = \frac{1}{K+I}\left[\hat{p}(E) + V_A s_F(a_2) + 0\right] = 0.218$
- $\overline{P}(E_{clouds}) = \frac{1}{K+I}\left[\hat{p}(E) + V_A s_F(a_2) + K\right] = 0.370$
- $\underline{P}(E_{sunny}) = \frac{1}{K+I}\left[\hat{p}(E) + V_A s_F(a_1) + 0\right] = 0.630$
- $\overline{P}(E_{sunny}) = \frac{1}{K+I}\left[\hat{p}(E) + V_A s_F(a_1) + K\right] = 0.782$

5.5 Coherence of dialectical probabilities

In this section, we show that dialectical probabilities (given in definition 5.15) are coherent previsions, as defined in Section 5.1.5. In order to prove this result, we establish that

1. upper dialectical probabilities are the dual of the lower ones,
2. lower dialectical probabilities avoid sure loss, and
3. lower dialectical probabilities coincide with their natural extension.

The first point is fairly easy to establish.

Property 5.10. *For every event* $E \subseteq \Omega$, $\overline{P}(E) = 1 - \underline{P}(E^c)$.

Proof. $\forall E \subseteq \Omega$, it holds that

- $\hat{p}(E) = 1 - \hat{p}(E^c)$
- $\hat{p}(E|X_a) = \frac{\hat{p}(E \cap X_a)}{\hat{p}(X_a)} = \frac{\hat{p}(X_a - E^c \cap X_a)}{\hat{p}(X_a)} = 1 - \hat{p}(E^c|X_a)$
- $\max(E) = 1 - \min(E^c)$

Therefore, by substitution

$$\begin{aligned}
\overline{P}(E) &= \frac{1}{K+I}\left[\hat{p}(E) + V_A \sum_{a \in A} s_F(a)\hat{p}(E|X_a) + K\max(E)\right] \\
&= \frac{1}{K+I}\left[(1 - \hat{p}(E^c)) + V_A \sum_{a \in A} s_F(a)(1 - \hat{p}(E^c|X_a)) + K(1 - \min(E^c))\right] \\
&= \frac{1}{K+I}\left[1 + V_A \sum_{a \in A} s_F(a) + K\right] \\
&\quad - \frac{1}{K+I}\left[\hat{p}(E^c) + V_A \sum_{a \in A} s_F(a)\hat{p}(E^c|X_a) + K\min(E^c)\right] \\
&= 1 - \underline{P}(E^c)
\end{aligned}$$

□

Before we prove the second point, let us first introduce the following notion:

Definition 5.16. *The **biased prior** $\hat{p}_A : \Omega \to [0,1]$ is given $\forall w \in \Omega$ by*

$$\hat{p}_A(w) = \frac{1}{I}\left[\hat{p}(w) + V_A \sum_{a \in A} s_F(a)\hat{p}(w|X_a)\right]$$

We may now rewrite the dialectical probabilities in terms of $\hat{p}_A$, a follows:

Property 5.11. *The lower and upper dialectical probabilities are equal to*

$$\underline{P}(E) = \frac{1}{K+I}(I\hat{p}_A(E) + K\min(E)) \quad \overline{P}(E) = \frac{1}{K+I}(I\hat{p}_A(E) + K\max(E))$$

We will use this property to prove that lower dialectical probabilities avoid sure loss.

Property 5.12. *For any collection of events $X_1, \ldots, X_k \subseteq \Omega$:*

$$\max\left(\sum_{j=1}^{k} G(X_j)\right) = \max_{w \in \Omega} \sum_{j=1}^{k} (X_j(w) - \underline{P}(X_j)) \geq 0$$

where $G(X)$ is the notation used for the marginal gamble $X - \underline{P}(X)$ (defined in Section 5.1.5).

Proof. Since $G(\Omega) = \Omega - \underline{P}(\Omega) = 1 - 1 = 0$, we may assume without loss of generality that $\forall j \in \{1, ..., k\}$, $X_j \subset \Omega$. For every event $X_j \subset \Omega$, we have $\min(X_j) = 0$ so the expression of the lower probability for X_j simplifies to

$$\underline{P}(X_j) = \frac{1}{K+I}\left[\hat{p}(X_j) + V_A \sum_{a \in A} s_F(a)\hat{p}(X_j|X_a)\right]$$

Since, by property 5.11, $\forall X_j$

$$\underline{P}(X_j) = \frac{1}{K+I}\left(I\hat{p}_A(X_j) + K\min(X_j)\right)$$

we may write that

$$\sum_{j=1}^{k} G(X_j) = \sum_{j=1}^{k} \left(X_j - \frac{I\hat{p}_A(X_j)}{K+I} \right)$$

Let us then consider the gambles $g : \Omega \to \mathbb{R}$ and $f : \Omega \to \mathbb{R}$ defined $\forall w \in \Omega$ as

$$g(w) = \sum_{j=1}^{k} \left(X_j(w) - \frac{I\hat{p}_A(X_j)}{K+I} \right)$$
$$f(w) = \sum_{j=1}^{k} (X_j(w) - \hat{p}_A(X_j))$$

Obviously $g \geq f$. To prove that $\underline{P}$ avoids sure loss, we must prove that $\max_{w\in\Omega} g(w) \geq 0$. Since $g \geq f$, it is sufficient to prove that $\max_{w\in\Omega} f(w) \geq 0$. Observe that by linearity of the expectation operation:

$$E_{\hat{p}_A}(f) = \sum_{j=1}^{k} (E_{\hat{p}_A}(X_j) - E_{\hat{p}_A}(\hat{p}_A(X_j)))$$

Since the expectation of a constant gamble is equal to that constant:

$$E_{\hat{p}_A}(f) = \sum_{j=1}^{k} (E_{\hat{p}_A}(X_j) - \hat{p}_A(X_j))$$

Finally, since the expectation of an event is equal to its probability:

$$E_{\hat{p}_A}(f) = \sum_{j=1}^{k} (\hat{p}_A(X_j) - \hat{p}_A(X_j)) = 0$$

We have found a non-null positive combination of values of f that is null. This would not be possible if all the values of f were strictly negative. Thus, one of them is positive or null, and so is a fortiori the maximum of f. This concludes the proof. □

Since dialectical lower probabilities avoid sure loss, they admit a natural extension $\underline{E}$. Before proving that dialectical probabilities coincide with their natural extension (on events), we are going to derive closed-form formulae for $\underline{E}(X)$ and $\overline{E}(X)$ valid for any gamble X.

Property 5.13. *For any gamble $X : \Omega \to \mathbb{R}$,*

$$\underline{E}(X) = \frac{1}{K+I}\left(IE_{\hat{p}_A}(X) + K\min(X)\right), \ \overline{E}(X) = \frac{1}{K+I}\left(IE_{\hat{p}_A}(X) + K\max(X)\right)$$

Proof. For any $X : \Omega \to \mathbb{R}$, $\underline{E}(X)$ is defined as the solution of the linear optimisation problem which consists in maximising the variable x [2] subject to the positivity constraints $\lambda_j \geq 0$ and $\forall w \in \Omega$:

$$X(w) - x \geq \sum_{X_j \subseteq \Omega} \lambda_j (X_j(w) - \underline{P}(X_j))$$

(see definition 5.7 in Section 5.1.5). The inclusion $X_j \subseteq \Omega$ can be replaced by a strict inclusion since the marginal gamble $G(\Omega) = \Omega - \underline{P}(\Omega)$ is null. Gambles X are now seen as vectors of dimension $n = |\Omega|$ with components $X_i = X(w_i)$, $\langle X, Y \rangle$ denotes the canonical scalar product $X'Y$, where X' is the transpose of X. With these notations, the lower probability $\underline{P}(X_j)$ for every event $X_j \subset \Omega$ can be rewritten

$$\underline{P}(X_j) = \langle \frac{I\hat{p}_A}{K+I}, X_j \rangle$$

Let us introduce the linear form $\sigma : \mathbb{R}^n \to \mathbb{R}$ defined $\forall X \in \mathbb{R}^n$ as

$$\sigma(X) = \langle \frac{I\hat{p}_A}{K+I}, X \rangle$$

By linearity and using the σ notation, the earlier inequality constraint is equiv-

[2] x is a generic variable here and has nothing to do with the x used in theorem 5.3 and thereafter to define confidence intervals.

alent to

$$X - x \geq \sum_{X_j \subset \Omega} \lambda_j X_j - \sigma\left(\sum_{X_j \subset \Omega} \lambda_j X_j\right)$$

The set of all positive combinations of $X_j \subset \Omega$ is equal to $(\mathbb{R}^+)^n$, so that our problem is equivalent to maximising x such that $X - x \geq Y - \sigma(Y)$ for some $Y \in (\mathbb{R}^+)^n$. Denoting by $(e_1, \ldots, e_n)$ the canonical basis of $\mathbb{R}^n$, whereby each e_i is the vector whose i-th component is equal to 1 and all other components are null, and $\sigma_j = \sigma(e_j)$, the solution x^* of the problem is given by

$$x^* = \max_{y_i \geq 0} \min_{i \in \{1,\ldots,n\}} \left(x_i - (y_i - \sum_{j=1}^{n} \sigma_j y_j)\right) = \max_{y_i \geq 0} \left(\sum_{j=1}^{n} \sigma_j y_j + \min_{i \in \{1,\ldots,n\}} (x_i - y_i)\right)$$

Clearly, the value of x^* is unchanged if we add the further constraint that the values of $x_i - y_i$ should all be equal to the same value m. The value of m would then be such that $m = x_i - y_i$ for all i and since y_i are positive, m must satisfy $m \leq \min(X)$. Substituting y_i by $x_i - m$, we obtain

$$\begin{aligned} x^* &= \max_{m \leq \min(X)} \left(\sum_{j=1}^{n} \sigma_j (x_j - m) + m\right) = \sum_{j=1}^{n} \sigma_j x_j + \max_{m \leq \min(X)} m\left(1 - \sum_{j=1}^{n} \sigma_j\right) \\ &= \sum_{j=1}^{n} \sigma_j x_j + \min(X)\left(1 - \sum_{j=1}^{n} \sigma_j\right) = \frac{I}{K+I} E_{\hat{p}_A}(X) + \frac{K}{K+I} \min(X) \end{aligned}$$

Thus, we have proved that

$$\underline{E}(X) = \frac{1}{K+I} \left(I E_{\hat{p}_A}(X) + K \min(X)\right)$$

$\overline{E}(X)$ is equal to the conjugate $-\underline{E}(-X)$, so

$$\overline{E}(X) = \frac{1}{K+I} \left(I E_{\hat{p}_A}(X) + K \max(X)\right)$$

□

We can now prove that lower dialectical properties coincide with their natural extension:

Property 5.14. $\forall X \subseteq \Omega$, $\underline{E}(X) = \underline{P}(X)$.

Proof. It is easy to see that, for any $X \subseteq \Omega$,

$$E_{\hat{p}_A}(X) = \sum_{w \in \Omega} \hat{p}_A(w) X(w) = \sum_{w \in \Omega} \hat{p}_A(w) = \hat{p}_A(X)$$

The property thus follows directly from property 5.13 and definition 5.15. □

It trivially follows from properties 5.10, 5.12 and 5.14 that

Theorem 5.5. *Dialectical probabilities are coherent.*

It is particularly remarkable that coherence holds independently of the disagreements and conflicts existing between the expert's opinions. Since they are coherent, dialectical probabilities inherit all the normative properties of coherent previsions. For completeness, we recall from [Walley, 1991] the main list of properties verified by $\underline{P}$ and $\overline{P}$ for any events $A, B, A_1, \ldots, A_s, B_1, \ldots, B_m$:

1. $0 \leq \underline{P}(A) \leq \overline{P}(A) \leq 1$
2. $\underline{P}(\emptyset) = \overline{P}(\emptyset) = 0$, $\underline{P}(\Omega) = \overline{P}(\Omega) = 1$
3. if $A \supseteq B$, then $\underline{P}(A) \geq \underline{P}(B)$ and $\overline{P}(A) \geq \overline{P}(B)$
4. $\underline{P}(A \cup B) \leq \underline{P}(A) + \overline{P}(B)$, $\overline{P}(A \cup B) \leq \overline{P}(A) + \overline{P}(B)$ (see also prop. 8)
5. if $A \cap B = \emptyset$, then $\underline{P}(A \cup B) \geq \underline{P}(A) + \underline{P}(B)$ and $\overline{P}(A \cup B) \geq \underline{P}(A) + \overline{P}(B)$
6. $\underline{P}(A) + \underline{P}(B) \leq 1 + \underline{P}(A \cap B)$, and if $A \cup B = \Omega$, then $\overline{P}(A) + \overline{P}(B) \geq 1 + \overline{P}(A \cap B)$

7. $|\underline{P}(A) - \underline{P}(B)| \leq \overline{P}(A \triangle B)$, $|\overline{P}(A) - \overline{P}(B)| \leq \overline{P}(A \triangle B)$

8. $\underline{P}(A) + \underline{P}(B) \leq \underline{P}(A \cup B) + \overline{P}(A \cap B) \leq \overline{P}(A) + \overline{P}(B)$,
   $\underline{P}(A) + \underline{P}(B) \leq \overline{P}(A \cup B) + \underline{P}(A \cap B) \leq \overline{P}(A) + \overline{P}(B)$,
   $\underline{P}(A \cup B) + \underline{P}(A \cap B) \leq \underline{P}(A) + \overline{P}(B) \leq \overline{P}(A \cup B) + \overline{P}(A \cap B)$ – the inequalities in property 4 can thus be made tighter as follows: $\underline{P}(A \cup B) \leq \underline{P}(A) + \overline{P}(B) - \underline{P}(A \cap B)$ and $\overline{P}(A \cup B) \leq \overline{P}(A) + \overline{P}(B) - \underline{P}(A \cap B)$

9. if $\overline{P}(A_s \triangle A) \to 0$ as $s \to \infty$, then $\underline{P}(A_s) \to \underline{P}(A)$ and $\overline{P}(A_s) \to \overline{P}(A)$

10. $\overline{P}(\bigcup_{j=1}^{s} A_j) \leq \sum_{j=1}^{s} \overline{P}(A_j)$ and if A_1,..., A_s are disjoint, then $\underline{P}(\bigcup_{j=1}^{s} A_j) \geq \sum_{j=1}^{s} \underline{P}(A_j)$ and $\underline{P}(\bigcup_{j=1}^{\infty} A_j) \geq \sum_{j=1}^{\infty} \underline{P}(A_j)$

11. $(\forall \lambda_1, ..., \lambda_s \geq 0$ and $\mu \in \mathbb{R})$
    if $B \geq \sum_{j=1}^{s} \lambda_j A_j + \mu$, then $\underline{P}(B) \geq \sum_{j=1}^{s} \lambda_j \underline{P}(A_j) + \mu$
    if $B \leq \sum_{j=1}^{s} \lambda_j A_j + \mu$, then $\overline{P}(B) \leq \sum_{j=1}^{s} \lambda_j \overline{P}(A_j) + \mu$

12. if $\lambda_1, ..., \lambda_s, \beta_1, ..., \beta_m \geq 0$ and $\sum_{j=1}^{s} \lambda_j A_j \leq \sum_{i=1}^{m} \beta_i B_i$, then $\sum_{j=1}^{s} \lambda_j \underline{P}(A_j) \leq \beta_1 \underline{P}(B_1) + \sum_{i=2}^{m} \beta_i \overline{P}(B_i)$ and $\lambda_1 \overline{P}(A_1) + \sum_{j=2}^{s} \lambda_j \underline{P}(A_j) \leq \sum_{i=1}^{m} \beta_i \overline{P}(B_i)$

5.6 Properties of dialectical probabilities

The aim of this section is to gain insight on dialectical probabilities and show some of their information-theoretic properties.

Dialectical lower and upper probabilities remain very similar to robust belief and plausibility functions. To see this, remark that the biased prior $\hat{p}_A$ introduced in the previous section (see definition 5.16) defines a Bayesian probability distribution $\hat{p}_A : \Omega \to [0, 1]$. Then we may see dialectical lower and upper probabilities as the imprecise probabilities associated with the credal set

$$\mathcal{D}_A = \{P \,|\, P = (1 - \epsilon_A)\, \hat{p}_A + \epsilon_A q, q \in \mathcal{P}\}$$

which is a maximal ϵ_A-contamination of $\hat{p}_A$ with level of contamination

$$\epsilon_A = \frac{K}{K+I}$$

Therefore, $\hat{p}_A$ still plays the role of prior. Because $\hat{p}_A$ is different from $\hat{p}$, we may say that $\hat{p}_A$ has a statistical bias. This bias is caused by the forecast arguments. This bias shall hopefully make $\hat{p}_A$ a better prior than the statistical prior $\hat{p}$.

Remark that the dialectical probabilities can be rewritten as follows:

Property 5.15. $\forall E \subseteq \Omega$,

$$\underline{P}(E) = (1 - \epsilon_A)\hat{p}_A(E) + \epsilon_A \min(E)$$

$$\overline{P}(E) = (1 - \epsilon_A)\hat{p}_A(E) + \epsilon_A \max(E)$$

Proof. This directly derives from property 5.4 and the fact that dialectical lower and upper probabilities are the imprecise probabilities associated with the maximal ϵ_A-contamination of $\hat{p}_A$. □

Clearly, since the imprecision of the probability of events is defined as $\Delta(E) = \overline{P}(E) - \underline{P}(E)$ (see definition 5.12 in Section 5.2)

Property 5.16. $\Delta(\emptyset) = 0$, $\Delta(\Omega) = 0$ *and for every* E *such that* $\emptyset \subset E \subset \Omega$,

$$\Delta(E) = \epsilon_A$$

Proof. This derives directly from property 5.15 and the fact that $\min(\emptyset) = \max(\emptyset) = 0$, $\min(\Omega) = \max(\Omega) = 1$ and for every event E distinct from $\emptyset$ and Ω, $\min(E) = 0$ and $\max(E) = 1$. □

Thus, ϵ_A measures the indeterminacy of the subject.

Property 5.17. *For a fixed value of the epistemic risk-aversion parameter K and a fixed argumentation framework F, it holds that*

1. $\epsilon_A \leq \epsilon$, *and*

2. ϵ_A *is a decreasing function of* V_A.

Proof. Since $I = 1 + I_A$ by definition 5.14 and $I_A = V_A \sum_{a \in A} s_F(a) \geq 0$, we have $I = I_A + 1 \geq 1$ and $\epsilon_A = K/(K+I) \leq K/(K+1) = \epsilon$. I_A and I increase with V_A, so ϵ_A decreases with V_A. □

So, as we expected, the informational value of arguments reduces the indeterminacy of the subject (expert opinions allow to make probabilistic estimations more precise).

Theorem 5.6. *Dialectical probabilities are the belief and plausibility functions whose evidence mass function $m_A : 2^\Omega \to [0,1]$ is defined for every event X as*

$$m_A(X) = (1 - \epsilon_A)\hat{p}_A(X)\delta(|X|, 1) + \epsilon_A \delta(|X|, |\Omega|)$$

Proof. This directly derives from the proof of property 5.5 and the fact that dialectical lower and upper probabilities are the imprecise probabilities associated with the maximal ϵ_A-contamination of $\hat{p}_A$. □

Belief and plausibility functions have been extensively studied within generalised information theory [Klir, 2006] where measures of indeterminacy and uncertainty have been developed. For any belief and plausibility functions with associated evidence mass function m, the indeterminacy of the subject is given by the *generalised Hartley measure*

$$GH = \sum_{B \subseteq \Omega} m(B) \log_2 |B|$$

The previous theorem allows us to compute the information-theoretic indeterminacy of the subject in the case of the dialectical probability.

Property 5.18. *The information-theoretic indeterminacy is*

$$GH = \epsilon_A \log_2 |\Omega|$$

Proof. For every $B \subseteq \Omega$ which is not a singleton or Ω, we have $m_A(B) = 0$. For every singleton $B = \{w\}$, we have $|B| = 1$ and $\log_2 |B| = 0$, so $GH = m_A(\Omega) \log_2 |\Omega| = \epsilon_A \log_2 |\Omega|$. □

The information-theoretic indeterminacy is thus proportional to the indeterminacy ϵ_A and the proportionality factor which is equal to $\log_2 |\Omega|$ increases with the total number of scenarios.

Moreover, when $I \ll K$, the indeterminacy $\epsilon_A = K/(K + I)$ is close to 1 and $GH \approx \log_2 |\Omega|$, which remarkably corresponds to the classical expression of the Hartley measure [Hartley, 1928]. This reflects the state of a subject that is almost completely ignorant about the true probability distribution of the scenarios. When more information is available, the imprecision and indeterminacy decrease inversely proportionally to I. Thus, the quantity I truly plays the role of *total amount of information* contained in the statistical sample and the expert opinions from the subject's viewpoint. In particular, K is the value of the amount of information I for which the imprecision is reduced from 1 to 1/2. Namely, the epistemic risk-aversion parameter K represents the "half-life" of the subject's indeterminacy. Let $\epsilon_A(I = s)$ stand for the value of ϵ_A when I is set to s. Then:

Property 5.19. $\epsilon_A(I = K) = 1/2$.

Proof. $\epsilon_A(I = K) = K/(K + K) = 1/2$. □

More generally, the subject's indeterminacy ϵ_A depends on the total amount of information I in the following way.

Property 5.20. *It holds that*

- *for any integer* $n > 0$

$$\epsilon_A(I = (n-1)K) = \frac{1}{n}$$

- $\lim_{s \to 1} \epsilon_A(I = s) = \epsilon$
- $\epsilon_A(I = s) \sim_{s \to +\infty} K/s$

Proof. $\epsilon_A(I = (n-1)K) = K/(K + (n-1)K) = K/(nK) = 1/n$. Since the function $s \mapsto K/(K+s)$ is contiuous in the neighbourhood of $s = 1$, we have $\lim_{s \to 1} \epsilon_A(I = s) = K/(K+1) = \epsilon$. $\epsilon_A(I = s) = K/(K+s) \sim_{s \to +\infty} K/s$. □

Thus, the indeterminacy is maximum and equal to ϵ when there are no arguments and varies inversely proportionally with the total amount of information available. K is the quantum of information necessary to reduce the indeterminacy from $1/n$ to $1/(n+1)$, for any integer $n > 0$.

We know that indeterminacy ϵ_A monotonically decreases with the amount of information I_A provided by the experts. Since I_A is by definition additive over the forecast arguments, we can qualitatively predict the value of the indeterminacy depending on the strength and acceptability of these arguments.

Property 5.21.

$$\frac{K}{K + V_A|A| + 1} \leq \epsilon_A \leq \frac{K}{K+1}$$

Proof. For every forecast argument $a \in A$ it holds that $0 \leq s_F(a) \leq 1$, so we have $0 \leq I_A \leq V_A|A|$. Hence the result. □

When K and V_A are fixed, this allows to determine a lower bound on the number of forecast arguments to elicit from experts in order to have a sufficiently low probabilistic imprecision. For instance, if one has $K = 1$ and $V_A = 1$, one always needs at least three forecast arguments to obtain $\Delta(E) = 20\%$ precision and at least eight forecast arguments to obtain $\Delta(E) = 10\%$ precision.

The indeterminacy is minimal when all forecast arguments have a strength of one, i.e. when none of them is mitigated (the experts have no doubt concerning their forecasts). The indeterminacy is maximal when either

- $V_A = 0$ (the experts are not trusted by the subject),
- $A = \emptyset$ (the experts do not have any opinion), or
- every forecast argument attacks itself [3] (all given forecasts contradicts themselves).

We can use our knowledge of the impact of the acceptability status of arguments on their strength (see Chapter 4) to obtain a tighter upper bound for indeterminacy.

Property 5.22. *The indeterminacy is bounded up by*

$$\epsilon_A \leq K/\left(K + V_A \sum_{a \in A} \underline{s}_F(a) + 1\right)$$

where for every forecast argument a,

$$\underline{s}_F(a) = \begin{cases} \max\left(\frac{1}{2(n(x)+1)}, \frac{1}{1+\sqrt{k}}, \frac{k+3}{4(k+1)}\right) & \text{if } a \text{ is contained in an admissible extension} \\ \frac{1}{2(n(x)+1)} & \text{if } a \text{ is only contained in a conflict-free extension} \end{cases}$$

Proof. This directly derives from properties 4.4 and 4.5 in Chapter 4. □

[3] Forecast arguments may be mitigation arguments too and thus attack themselves.

When K, V_A and F are fixed, this allows to determine a guaranteed level of precision without having to compute the precise strength of each forecast argument. Instead, one simply needs to identify the conflict-free and admissible arguments. For instance, if one has $K = 1$, $V_A = 1$ and the framework F has an out-degree $k = 1$ and the maximum number of attacks against an argument is $n(x) = 2$, then each conflict-free forecast argument has a strength of at least $s_F(a) = 1/6$ and each admissible argument has a strength of at least $s_F(a) = 1/2$. Thus, with three admissible forecast arguments the level of precision that is guaranteed is $\Delta(E) = 29\%$ and the level of precision guaranteed with three conflict-free arguments is only $\Delta(E) = 40\%$.

Other bounds on the indeterminacy can be obtained when the framework is a disjoint union of frameworks $F = F_1 + \ldots + F_k$. Recall that the geometric mean of $x_1, \ldots, x_k \in \mathbb{R}$ is defined as

$$H(x_1, \ldots, x_k) = \frac{k}{\frac{1}{x_1} + \ldots + \frac{1}{x_k}}$$

We have

Property 5.23. *If F is the disjoint union of $F_1, \ldots, F_k$, then*

$$\frac{1}{k} H(\epsilon_{A_1}, \ldots, \epsilon_{A_k}) \leq \epsilon_A \leq \min(\epsilon_{A_1}, \ldots, \epsilon_{A_k})$$

where ϵ_{A_i} denotes the indeterminacy of the subject when only the opinions from the framework F_i are considered.

Proof. If the given abstract argumentation framework $F = F_1 + \ldots + F_k$ where $F_1 = (Arg_1, att_1), \ldots, F_k = (Arg_k, att_k)$ are mutually disjoint frameworks (see the study of dynamic properties of the strength measure in Section 4.3.4), then by property 4.9, for every $a \in F_i$, we have $s_F(a) = s_{F_i}(a)$. Denoting $A_1, \ldots, A_k$

the sets of forecast arguments in $F_1, \ldots, F_k$ respectively, we have

$$\epsilon_A = K/\left(K + V_A \sum_{i=1}^{k} \sum_{a \in A_i} s_{F_i}(a) + 1\right)$$

ϵ_A can be expressed in terms of the quantities

$$\epsilon_{A_i} = \frac{K}{K + V_A \sum_{a \in A_i} s_{F_i}(a) + 1}$$

by

$$\epsilon_A = \frac{1}{\sum_{i=1}^{k} \frac{1}{\epsilon_{A_i}} - (k-1)(1 + \frac{1}{K})}$$

Hence the inequality

$$\frac{1}{k} H(\epsilon_{A_1}, \ldots, \epsilon_{A_k}) \leq \epsilon_A$$

For every i, it is also clear that $\epsilon_A \leq \epsilon_{A_i}$, thus

$$\epsilon_A \leq \min(\epsilon_{A_1}, \ldots, \epsilon_{A_k})$$

□

This allows to get an idea of the level of precision one can obtain by merging the probability estimations made independently by several experts. For instance, if one has three groups of experts using independent bodies of knowledge to formulate their arguments and the precision of their probability estimates are 20%, 30% and 40%, then by merging their opinions one can guarantee that the new precision will be comprised between 9.2% and 20%.

Based on the generalised Hartley measure a measure of uncertainty is also defined in Dempster-Shafer theory. This is called *strife* or *generalised entropy*

and defined for any belief function with evidence mass function m as

$$ST = GH - Z$$

where

$$Z = \sum_{X \subseteq \Omega} m(X) \log_2 \left(\sum_{Y \subseteq \Omega} m(Y) |X \cap Y| \right)$$

In the case of the dialectical probability,

Property 5.24. *The generalised entropy or strife is*

$$ST = \epsilon_A \log_2 \left(\frac{|\Omega|}{1 + \epsilon_A(|\Omega| - 1)} \right) - \sum_{w \in \Omega} (1 - \epsilon_A)\hat{p}_A(w) \log_2 \left((1 - \epsilon_A)\hat{p}_A(w) + \epsilon_A \right)$$

Proof. The evidence mass function of dialectical lower and upper probabilities (see theorem 5.6) is such that $m_A(B) \neq 0$ only if $|B| = 1$ or $B = \Omega$. Thus, it holds that

$$Z = \sum_{w \in \Omega} m_A(\{w\}) \log_2 \left(m_A(\{w\}) + m_A(\Omega) \right)$$

$$+ m_A(\Omega) \log_2 \left(\sum_{w \in \Omega} m_A(\{w\}) + m_A(\Omega)|\Omega| \right)$$

Since m_A is normalised, we have in fact

$$Z = \sum_{w \in \Omega} m_A(\{w\}) \log_2 \left(m_A(\{w\}) + m_A(\Omega) \right) + m_A(\Omega) \log_2 \left(1 + m_A(\Omega)(|\Omega| - 1) \right)$$

Since, by theorem 5.6, $m_A(\{w\}) = (1 - \epsilon_A)\hat{p}_A(w)$ and $m_A(\Omega) = \epsilon_A$, we obtain

$$Z = \sum_{w \in \Omega} (1 - \epsilon_A)\hat{p}_A(w) \log_2 \left((1 - \epsilon_A)\hat{p}_A(w) + \epsilon_A \right) + \epsilon_A \log_2 \left(1 + \epsilon_A(|\Omega| - 1) \right)$$

Since $GH = \epsilon_A \log_2 |\Omega|$ by property 5.18 and by definition $ST = GH - Z$, we

get

$$ST = \epsilon_A \log_2 \left(\frac{|\Omega|}{1 + \epsilon_A(|\Omega| - 1)} \right) - \sum_{w \in \Omega} (1 - \epsilon_A)\hat{p}_A(w) \log_2 \left((1 - \epsilon_A)\hat{p}_A(w) + \epsilon_A \right)$$

□

In particular, if $I \gg K$, then $\epsilon_A \ll 1$ and we have

$$ST \approx - \sum_{w \in \Omega} \hat{p}_A(w) \log_2 \hat{p}_A(w)$$

which remarkably corresponds to the classical expression of entropy [Shannon, 1948]. We will resist the temptation to investigate the significance of this result. [Klir, 2006] explains indeed that the strife ST is not a good measure of uncertainty, as it is not sub-additive, and suggests to use instead the *maximum Shannon entropy*, defined for any imprecise probability with credal set $\mathcal{D}$, as

$$\overline{S} = \max_{p \in \mathcal{D}} - \sum_{w \in \Omega} p(w) \log_2 p(w)$$

Expert opinions may reduce the indeterminacy ϵ_A by bringing in new information (increasing the total amount of information I), but this does not necessarily decrease the subject's uncertainty (measured by either ST or $\overline{S}$). Information is what decreases the indeterminacy ϵ_A and is necessary (although not sufficient) to reduce the uncertainty, since information may sometimes increase the uncertainty by making the biased prior $\hat{p}_A$ closer to the uniform probability distribution $u : \Omega \to [0, 1]$, defined $\forall w \in \Omega$, as $u(w) = 1/|\Omega|$ and whose entropy $S(u) = \log_2 |\Omega|$ is maximal.

5.7 Application to trust computing

Belief functions have been used as a mathematical model of trust in distributed reputation management [Yu and Singh, 2002]. Since dialectical probabilities are belief functions, we may use them for computing trust. Belief functions for trust are usually derived from statistical data only. We show that, for some adequate choice of the epistemic risk-aversion parameter K, the robust belief function of Section 5.2 coincides with the belief function normally employed in Yu & Singh's model. We are also naturally interested in comparing these belief functions to dialectical probabilities and see whether arguments add value to the purely statistical method.

5.7.1 Background on the evidence-based trust model

The model developed by [Yu and Singh, 2002] is called the evidence-based trust model. This model uses belief functions as a mathematical model of trust. This model is based upon a simple universe $\Omega = \{T, \neg T\}$, where T (respectively $\neg T$) means that the subject or agent a_i (called the *evaluator*) can consider a given agent a_j (called the *target*) to be trustworthy (respectively untrustworthy). The belief of a_i concerning a_j is modeled by a belief function $Bel : 2^{\Omega} \rightarrow [0, 1]$. The value $Bel(\{T\}) = m(\{T\})$ is interpreted as the trust of a_i in a_j and $Bel(\{\neg T\}) = m(\{\neg T\})$ as the distrust. The cautiousness or risk aversion of the evaluator is modeled by a parameter denoted ρ comprised between 0 and 1. The evaluator decides to trust the target if and only if trust exceeds distrust by the threshold value ρ, that is to say, if and only if

$$m(\{T\}) - m(\{\neg T\}) \geq \rho$$

The target may for instance, endorse the role of a service provider (as is the

case in grid computing). For each interaction, it is assumed that the quality of the service provided by the target can be quantified by the evaluator. We denote by q the service quality below which the service is deemed *too poor* and by Q the service quality above which the service is deemed *satisfying*. Naturally, q must be strictly smaller than Q. When the quality is strictly greater than q but strictly smaller than Q, the service quality is said to be *inappreciable*.

In the evidence-based trust model, the evaluator knows

- N^- the number of times the service quality was too poor
- N^+ the number of times the service quality was satisfying
- $N_?$ the number of times the service quality was inappreciable
- $N = N^- + N^+ + N_?$ the total number of past interactions with the target

Finally, the belief function *Bel* used by Yu & Singh to compute trust is determined by the evidence mass function $m : 2^\Omega \to [0,1]$ where

$$m(\emptyset) = 0, \quad m(\{T\}) = \frac{N^+}{N}, \quad m(\{\neg T\}) = \frac{N^-}{N}, \quad m(\Omega) = \frac{N_?}{N}$$

5.7.2 Robust belief functions for trust

Each time the service quality is too poor counts as one occurrence of the scenario $\neg T$, and each time the service quality is satisfying counts as one occurrence of the scenario T. We therefore consider the statistical prior $\hat{p} : \Omega \to [0,1]$ defined on $\Omega = \{T, \neg T\}$ as

$$\hat{p}(T) = \frac{N^+}{N^+ + N^-}, \quad \hat{p}(\neg T) = \frac{N^-}{N^+ + N^-}$$

The definition of the robust belief function relies upon that statistical prior.

Theorem 5.7. *If $K = N_?/(N - N_?)$, then the robust belief function coincides with Yu & Singh's belief function.*

Proof. The evidence mass function of the robust belief function is

$$m(\emptyset) = 0, \quad m(\{T\}) = (1-\epsilon)\hat{p}(T), \quad m(\{\neg T\}) = (1-\epsilon)\hat{p}(\neg T), \quad m(\Omega) = \epsilon$$

Yu & Singh's belief function has an evidence mass function defined as

$$m'(\emptyset) = 0, \quad m'(\{T\}) = \frac{N^+}{N}, \quad m'(\{\neg T\}) = \frac{N^-}{N}, \quad m'(\Omega) = \frac{N_?}{N}$$

When $K = N_?/(N - N_?)$, the two evidence mass functions are identical:

$$\begin{aligned}
m(\emptyset) &= 0 = m'(\emptyset) \\
m(\{T\} &= (1-\epsilon)\hat{p}(T) = (1 - \frac{N_?}{N})\frac{N^+}{N^+ + N^-} = \frac{N^+}{N} = m'(\{T\}) \\
m(\{\neg T\}) &= (1-\epsilon)\hat{p}(\neg T) = (1 - \frac{N_?}{N})\frac{N^-}{N^+ + N^-} = \frac{N^-}{N} = m'(\{\neg T\}) \\
m(\Omega) &= \epsilon = \frac{K}{K+1} = \frac{\frac{N_?}{N-N_?}}{\frac{N_?}{N-N_?}+1} = \frac{N_?}{N? + N - N_?} = \frac{N_?}{N} = m'(\Omega)
\end{aligned}$$

□

Robust belief functions consequently generalise Yu & Singh's belief functions to any value possible value of K. In the remainder, we will assume that K is however fixed to the specific value given in theorem 5.7 in order to allow experimental comparisons between $m = m'$ and m_A.

5.7.3 Arguments for reasoning about trust

Statistics provide information about the past behaviour of the target, whilst arguments can provide information concerning its future behaviour. In grid computing, the arguments of interest may pertain to various *dimensions* of

trust such as the availability, security, privacy and reliability of services. For each such dimension, the idea is to build arguments from the content of the contract linking the target to the evaluator. We consider a simple abstract argumentation framework F per dimension, having as possible arguments

- $\neg T$: don't trust the target because there is no guarantee specific to that trust dimension given in the form of a contract clause
- T: trust the target because there exists a guarantee specific to that trust dimension given in the form of a contract clause
- V: don't rely on contract clauses concerning this trust dimension because the target has usually violated them in the past

For each trust dimension d, either there exists a contract clause concerning d and T is put into F, or there does not exist such a clause in the contract and we put instead $\neg T$ in F. If there is a contract clause, and, in the past, we have observed that when there was such a clause, d did not have an acceptable level in a majority of cases at execution time, then we also put in F the argument V as well as an attack from V against T. Consequently, each individual framework contains only one or two arguments.

5.7.4 Dialectical probabilities for trust

We thus define one dialectical probability per trust dimension. We want to allow the evaluator to decide separately if the target can be trusted with respect to each dimension and consider that the evaluator eventually decides to trust the target if and only if it decides to trust the target with respect to all four dimensions (availability, security, privacy and reliability).

In order to build a dialectical probability, we need to choose a value for the

parameter V_A (see Section 5.4). We choose

$$V_A = \frac{\hat{p}(T) - \hat{p}_{\neg c}(T)}{\hat{p}_{\neg c}(T)}$$

whereby $\hat{p}_{\neg c}(T)$ corresponds to the probability of the target being trustworthy given that there is no contract clause, defined as

$$\hat{p}_{\neg c}(T) = \frac{N^+_{\neg c}}{N^+_{\neg c} + N^-_{\neg c}}$$

where we denote by

- $N^+_{\neg c}$, the total number of times the service quality was satisfying and there was no contract clause, and
- $N^-_{\neg c}$, the total number of times the service quality was too poor and there was no contract clause.

For this choice of V_A, the biased prior $\hat{p}_A$ (see definition 5.16 in Section 5.5) satisfies the following intuitively appealing constraint.

Property 5.25. *When there is no contract clause the biased prior $\hat{p}_A(T)$ equals the conditional probability $\hat{p}_{\neg c}(T)$.*

Proof. If there is no contract clause, then by definition F contains only the argument $\neg T$. This argument is not attacked and has a strength of one. The value of $\hat{p}_A(T)$ is in that case

$$\hat{p}_A(T) = \frac{1}{1 + V_A}\hat{p}(T)$$

since $\hat{p}(T|\neg T) = 0$. The solution V_A of the equation $\hat{p}_A(T) = \hat{p}_{\neg c}(T)$ is then after simple calculus

$$V_A = \frac{\hat{p}(T) - \hat{p}_{\neg c}(T)}{\hat{p}_{\neg c}(T)}$$

Thus, for that choice of V_A, we have indeed $\hat{p}_A(T) = \hat{p}_{\neg c}(T)$. □

We can thus use the dialectical probability (for a given dimension and argumentation framework). This is computed via the evidence mass function m_A given in theorem 5.6, Section 5.4. This gives

$$m_A(\emptyset) = 0, m_A(\{T\}) = (1-\epsilon_A)\hat{p}_A(T), m_A(\{\neg T\}) = (1-\epsilon_A)\hat{p}_A(\neg T), m_A(\Omega) = \epsilon_A.$$

As for Yu & Singh's mass function, the evaluator decides to trust with respect to a dimension if and only if the following inequality holds for m_A for that dimension:

$$m_A(\{T\}) - m_A(\{\neg T\}) \geq \rho$$

Next, we compare experimentally the performance of evaluators using the original belief function and our dialectical probability.

5.7.5 Performance of dialectical probabilities

To test our model, we fix an evaluator and target and generate a large number of interactions with random outcomes. For each interaction, having informed the target about the specificity of the situation we first

- ask the target to make a decision concerning the target (trust or not trust)
- then, we determine the precise outcome of the interaction and finally
- assign a score to the target, depending on the correctness of its decision.

We assume that the evaluator decides to trust if and only if the condition

$$Bel(\{T\}) - Bel(\{\neg T\}) \geq \rho$$

is satisfied where ρ has a constant value over time. We will naturally be interested in the influence of ρ on the performance of the technique. For a fixed value of ρ and after a large number of interactions, we want to compare the percentage of correct decisions made using Yu & Singh's belief function (which corresponds to the robust belief function, see Section 5.7.2) versus the performance obtained using the dialectical probabilities (Section 5.7.4).

Each interaction is random, which means that its characteristics are determined randomly, but also, that its outcome is generated randomly from its characteristics. The basic characteristics of an interaction are the presence or not of a contract clause guaranteeing a satisfying level of service quality for each trust dimension. The presence of each type of clause is modeled by a Boolean random variable following a Bernoulli distribution with mean 0.8. This means that the probability for an interaction to have a contract clause (on reliability for instance) is 80%. All four random variables are identically and independently distributed.

The randomness of the outcome of a situation with known basic characteristics is solely due to the unpredictability of the target agent. Each outcome is modeled as a quadruple of Beta-distributed variables noted A (availability), S (security), P (privacy) and R (reliability). By definition of a Beta distribution, these four variable take their values inside the interval $[0, 1]$. The variances of the Beta distributions are fixed to $\sigma^2 = 3\%$ but the means can take two possible values: $\mu = 0.80$ or $\mu = 0.40$. We use

- the higher mean of $\mu = 0.80$ (for instance for the distribution of the variable A) when there exists a contract clause guaranteeing the service availability and the target has decided to respect the clause,
- the lower mean of $\mu = 0.40$ otherwise, when there is no contract clause or when there is one but the target has decided not to care about it.

A Beta distribution with high μ models the target's intention to produce a service of good quality, and a distribution with low μ a lack of effort on the side of the target. The intention to make this effort is mainly conditioned by the presence of a contract clause. This idea here is that our statistical setting should capture the fact that service providers tend to try and honour the contracts they sign. Our model of trust also needs to be slightly more realistic than this and take into account the existence of fraud. Therefore, we introduce an additional parameter θ representing the frequency with which the target tries to cheat (ignore the clauses of the contract). Thus, when the fraud frequency is $\theta = 0$, the target is perfectly honest and never ignores contract clauses ($\mu = 0.80$ is systematically used when there is a contract and $\mu = 0.40$ otherwise). When on the opposite the fraud frequency is $\theta = 1$, the target is never honest ($\mu = 0.40$ all the time).

In each interaction, once the parameters of the Beta distributions have been fixed, the variables A, S, P and R can be measured. The level of service availability A (similarly for security S, privacy P and reliability R) is deemed

- sufficient if and only if $A \geq Q = 0.6$,
- insufficient if and only if $A \leq q = 0.4$,
- inappreciable if and only if $q = 0.4 < A < Q = 0.6$.

Overall, the outcome is satisfying if and only if all four levels of availability, security, privacy and reliability are sufficient. The outcome is disappointing otherwise (at least one of the levels of availability, security, privacy or reliability is not sufficient). The trust decision is deemed correct if and only if

- the evaluator decides to trust and the outcome is satisfying, or
- the evaluator decides not to trust and the outcome is disappointing.

We use a sample of size 1000 interactions and start registering the evaluator's score at interaction number 20 only, so as to allow the target to compute $\hat{p}$.

Our first experiment assumes $\theta = 0$ (no fraud) and compares the percentage of correct decisions obtained with the belief functions depending on the threshold parameter ρ (level of cautiousness of the evaluator). The two curves obtained are depicted in Figure 5.1.

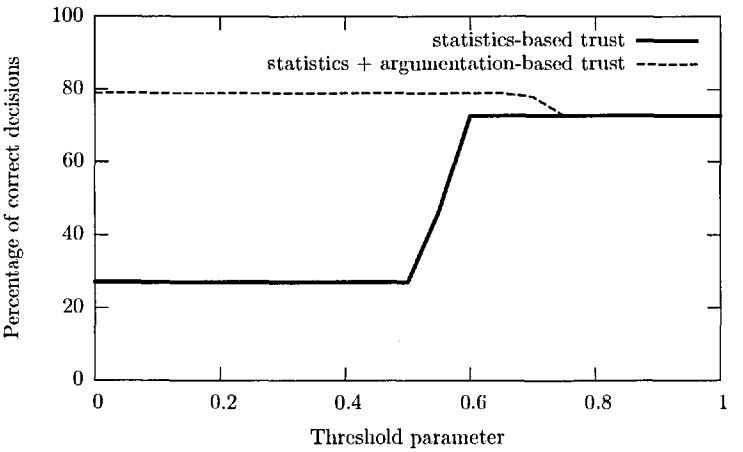


Figure 5.1: Performance depending on ρ and assuming $\theta = 0$ (no fraud).

The performance curve obtained for Yu & Singh's belief function shows a plateau at 27.04% performance for $\rho \leq 0.5$, a short performance affine transition for $\rho \in [0.5, 0.6]$ and another higher plateau at 72.86% for $\rho \geq 0.6$. The dialectical probability leads to a long performance plateau at 78.98% for $\rho \leq 0.65$, a short performance transition for $\rho \in [0.65, 0.75]$ and a slightly lower plateau also at 72.86% for $\rho \geq 0.75$. For $\theta = 0$, the performance is thus increased by

- 51.94% points for $\rho \in [0, 0.5]$
- more than 6.12% points for $\rho \in [0.5, 0.65]$
- 0% point or less than 6.12% points for $\rho \in [0.65, 1]$

The reason why no improvement is possible for large values of ρ is that such values make the evaluator overly cautious and lead him/her to systematically not trust (under both belief functions).

If the target delivers a service with identically and independently distributed service quality over time, the evaluator will constantly make the same decision. If the decision is always to trust, then the percentage of correct decisions will be equal to $F = N^+/N$. If the decision is always to distrust, then the percentage of correct decisions will be equal to $1-F = (N^- + N_?)/N$. Thus, the theoretical performance of Yu & Singh's belief function is at most equal to $\max(F, 1-F)$. Here we have $F = 27.04\%$, $1-F = 72.85\%$ and consequently MaxPerf $= 72.85\%$. The performance experimentally obtained with Yu & Singh's belief function is indeed limited by MaxPerf, but the dialectical probability allows to break that theoretical barrier. In fact it offers a level of performance that is either equal or above it for any value of ρ.

Let us examine now how these results change when fraud is taken into account. The performance depending on ρ and θ in shown in the table of Figure 5.2 for Yu & Singh's belief function and in the table of Figure 5.3 for the dialectical probability. To compare the two belief functions, we can use these tables to derive a worst-case performance measure with respect to the fraud frequency θ. This measure still depends on ρ. The worst-case performance of each belief function is shown in Figure 5.4. The worst case performance under the dialectical probability is clearly higher (irrespective of the choice of ρ), with at best an improvement of 53.47% points when $\rho \approx 0$. The improvement is null for high values of ρ (as we should have expected). The average worst-case performance improvement calculated for ρ ranging over $[0, 1]$ is of 28.21% points.

ρ	$\theta = 0$	$\theta = 0.25$	$\theta = 0.5$	$\theta = 0.75$	$\theta = 1$
0	27.04%	20.71%	16.73%	91.84%	99.69%
0.05	27.04%	20.71%	21.94%	91.84%	99.69%
0.1	27.04%	20.71%	76.12%	91.84%	99.69%
0.15	27.04%	20.71%	85.31%	91.84%	99.69%
0.2	27.04%	20.71%	85.31%	91.84%	99.69%
0.25	27.04%	21.43%	85.31%	91.84%	99.69%
0.3	27.04%	28.67%	85.31%	91.84%	99.69%
0.35	27.04%	72.85%	85.31%	91.84%	99.69%
0.4	27.04%	78.47%	85.31%	91.84%	99.69%
0.45	27.04%	79.08%	85.31%	91.84%	99.69%
0.5	27.04%	79.18%	85.31%	91.84%	99.69%
0.55	46.43%	79.18%	85.31%	91.84%	99.69%
0.6	72.86%	79.18%	85.31%	91.84%	99.69%
0.65	72.86%	79.18%	85.31%	91.84%	99.69%
0.7	72.86%	79.18%	85.31%	91.84%	99.69%
0.75	72.86%	79.18%	85.31%	91.84%	99.69%
0.8	72.86%	79.18%	85.31%	91.84%	99.69%
0.85	72.86%	79.18%	85.31%	91.84%	99.69%
0.9	72.86%	79.18%	85.31%	91.84%	99.69%
0.95	72.86%	79.18%	85.31%	91.84%	99.69%
1	72.86%	79.18%	85.31%	91.84%	99.69%

Figure 5.2: Influence of θ on the performance of Yu & Singh's belief function.

ρ	$\theta = 0$	$\theta = 0.25$	$\theta = 0.5$	$\theta = 0.75$	$\theta = 1$
0	78.97%	75.51%	70.20%	84.49%	99.69%
0.05	78.97%	75.51%	70.20%	91.12%	99.69%
0.1	78.97%	75.51%	70.20%	91.84%	99.69%
0.15	78.97%	75.51%	70.20%	91.84%	99.69%
0.2	78.97%	75.51%	70.20%	91.84%	99.69%
0.25	78.97%	75.51%	70.20%	91.84%	99.69%
0.3	78.97%	75.51%	81.53%	91.84%	99.69%
0.35	78.97%	75.51%	85%	91.84%	99.69%
0.4	78.97%	75.51%	85.31%	91.84%	99.69%
0.45	78.97%	75.51%	85.31%	91.84%	99.69%
0.5	78.97%	75.51%	85.31%	91.84%	99.69%
0.55	78.97%	78.97%	85.31%	91.84%	99.69%
0.6	78.97%	79.18%	85.31%	91.84%	99.69%
0.65	78.97%	79.18%	85.31%	91.84%	99.69%
0.7	77.95%	79.18%	85.31%	91.84%	99.69%
0.75	72.86%	79.18%	85.31%	91.84%	99.69%
0.8	72.86%	79.18%	85.31%	91.84%	99.69%
0.85	72.86%	79.18%	85.31%	91.84%	99.69%
0.9	72.86%	79.18%	85.31%	91.84%	99.69%
0.95	72.86%	79.18%	85.31%	91.84%	99.69%
1	72.86%	79.18%	85.31%	91.84%	99.69%

Figure 5.3: Influence of θ on the performance of dialectical probabilities.

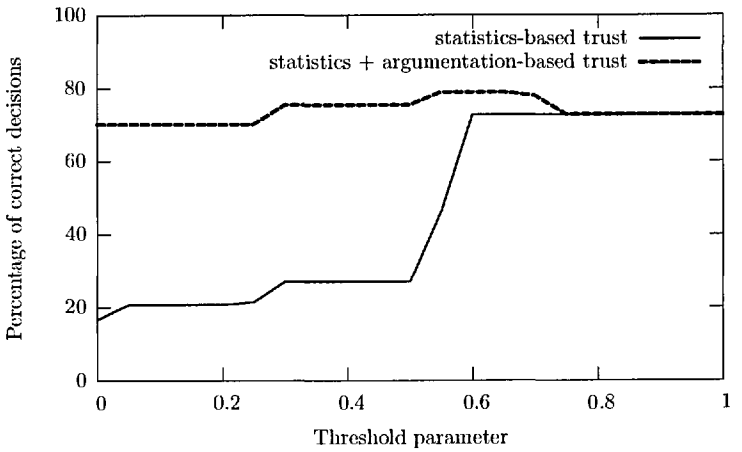


Figure 5.4: Worst-case performance with respect to θ depending on ρ.

5.8 Related work

This section discusses the relationship existing between our approach and methods used in subjective probability as well as work on uncertainty done in the areas of logic, artificial intelligence, machine learning, multi-agent systems and decision theory.

In subjective probability, the intuitive comparative school (Koopman and Good) exploits the concept of ordering relation – not more probable than – to derive an empirical probability measure. A second school is based on Ramsey and Savage utility decision making approach, whereby the choices of a decision maker (modeled as an expected utility maximiser) are used to reveal a set of probabilities. Subjective probabilities, when elicited by such means, are referred to as psychological probabilities to recognise the lack of total rationality and coherence expected from human responses [Chesley, 1975]. Psychological probabilities are descriptive of human beliefs and may not conform to the basic axioms of probability theory. Instead, this chapter has provided an approach that always guarantees the normative goodness of probabilities in terms of co-

herence.

A general procedure for eliciting imprecise probabilities from a subject is proposed in [Walley, 1991]. The idea is to obtain a number of qualitative or quantitative probability judgements and model them as marginally acceptable gambles. If this set of marginal gambles avoids sure loss, then one can use the natural extension to assess the probability of any event. This method is not easily applicable. The first problem is that the method fails whenever conflicting judgements are employed. The second problem is that checking that the property of avoiding sure loss holds is computationally expensive and having to compute linear extension by linear optimisation may be impractical. Dialectical probabilities avoid sure loss and are coherent even when arguments conflict. Moreover, we dispose of analytic expressions to compute them directly so we notably avoid the use of optimisation (dialectical probabilities are already optimal, since we have proved that they coincide with their natural extension).

Belief functions are coherent [Walley, 2000] and have been used for the design of many expert systems [Biswas *et al.*, 1988; Hsia and Shenoy, 1989; Kak *et al.*, 1990; Krause and Clark, 1993; Saffiotti and Umkehrer, 1991]. Dempster-Shafer theory of evidence is not concerned with the problem of how to construct evidence mass functions. In this chapter, our contribution has been to introduce methods for constructing belief functions based on statistical information and expert opinions.

In logic, argumentation theory is fundamentally concerned with the characterisation and computation of rationally acceptable sets of arguments [Dung, 1995] or assumptions [Bondarenko *et al.*, 1997; Dung *et al.*, 2006]. The primal goal of this theory is not to assess uncertainty but rather to study what a rational subject may believe. Argumentation can be seen essentially as a qualitative and logical approach to uncertainty. Probability is therefore not considered as

essential to argumentation theory. However, a few researchers in logic and artificial intelligence have advanced the idea of estimating both objective [Poole, 1993] and subjective probabilities [Krause *et al.*, 1995; Ambler, 1996] by means of argument aggregation. The basic idea behind Poole's probabilistic calculus is to estimate the probability of a claim being true by summing the probability of statistically independent, logically derived arguments that support that same claim. In other words, Poole's work can be seen as argumentation aggregation for statistical inference, i.e. the inference of probabilistic information from an initial set of probabilistic judgements. On the opposite, we focus on problems in which no probabilistic information is available to the subject. The authors of [Krause *et al.*, 1995; Ambler, 1996] embrace a more general and abstract view on argument aggregation than Poole and interestingly propose the use of algebraic calculus to assess (symbolically or numerically) the strength of arguments depending on their internal structure. The method used in this thesis also involves a notion of argument strength, which may be assessed from the arguments internal structure, but which may as well be derived from an analysis of the dialectical relationships existing between the arguments.

Our experiments on trust have shown that simple arguments could increase the performance of purely statistical methods. Recently, argumentation has been used to improve the performance of machine learning techniques [Mozina *et al.*, 2007]. The idea is to use arguments given by experts to explain logically the correct classification of examples of individuals in a class. These arguments are then used to guide the construction of a theory (set of if-then rules linking the attributes of individuals to classes). Arguments thus restrict the complexity of the search for a theory but more importantly, also lead to theories that perform statistically better outside the learning set.

Argumentation has been advocated by [Nielsen and Parsons, 2007] as a gen-

eral way to guide probabilistic reasoning in multi-agent systems. In their setting – in which an agent's knowledge is modelled by a Bayesian network – abstract argumentation-driven debates allow the agents of the system to reason about the probabilistic dependencies that exist between a set of propositions of interest and to eventually reach an optimal consensus, i.e. agree on which propositions should objectively be regarded as true or false.

Finally, textbooks on decision theory say little about the actual construction of probabilities [Fishburn, 1964; French, 1987]. In fact, standard decision theory does not try to answer how probabilities should be constructed but considers them as given a priori. The authors of [Tan and Pearl, 1994], amongst others, indicate that the specification of complete sets of probabilities and utilities makes decision theory impractical in complex tasks involving common sense knowledge. This important issue has favoured the emergence of qualitative approaches to decision making [Tan and Pearl, 1994; Boutilier, 1994; Dubois and Prade, 1995; Bonet and Geffner, 1996; Amgoud and Prade, 2004]. These works rely on mathematical models of uncertainty that not only differ from probability but, as for decision theory, do not seek to explain how to measure uncertainty. Instead, we have provided a practical method for estimating probabilities given statistical data and expert opinions.

5.9 Summary and conclusion

We have proposed a method for assessing imprecisely and subjectively the probabilities of future events based on a statistical sample and the opinions of experts in the form of arguments. We have first shown how to construct robust belief and plausibility functions based on the statistical sample solely and with any arbitrarily fixed level of statistical confidence. Then, we have shown that these belief functions correspond to the lower and upper envelopes of maximal

ϵ-contaminations of an unbiased statistical prior with a level of contamination ϵ that increases with the chosen level of confidence.

In order to take into account the experts' opinions, we have first introduced an abstract argumentation model able to capture the claims made by the experts, the disagreements between the experts' forecasts and conflicts between their arguments. In order to discriminate between forecast arguments, we have used the dialectical measure of argument strength from Chapter 4.

We have then exposed a general method for exploiting (conflicting) expert opinions in combination with statistical information and given a concrete formula for the construction of imprecise probabilities. The resulting probabilities have been proved to remain in the form of belief and plausibility functions and to be equivalent to maximal ϵ-contamination models. Thus, the proposed formula is a contribution to Dempster and Shafer's theory in that it addresses the issue of constructing evidence mass functions, which is not addressed by this theory. A thorough analysis of dialectical probabilities in the context of robust statistics and generalised uncertainty and information theory has followed, with the presentation of subjective measures of indeterminacy (the generalised Hartley measure) and uncertainty (the generalised entropy or strife).

Dialectical probabilities have been shown to be coherent previsions, so the formula proposed brings to the field of imprecise probability a way to achieve coherence even when the subject's opinion is based on conflicting judgements, as opposed to Walley's scheme which requires the probabilistic judgements to be mutually compatible. This provides us with a strong guarantee on the rationality of this subjective model of probability. Coherence offers us a rigorous mathematical framework for solving problems of statistical inference, decision analysis, risk analysis and decision making, which is the object of study of the next chapter. Dialectical probabilities have been applied with success to trust

computing. In the next chapter, we are going to use them to solve decision making problems in the field of finance.

Chapter 6

Statistical inference, risk analysis & decision making with dialectical probabilities

This chapter is dedicated to statistical inference, risk analysis and decision making when the decision maker's uncertainty is modeled by dialectical probabilities. Dialectical probabilities have been defined in Chapter 5. These probabilities are imprecise and non-Bayesian. In decision under risk, a Bayesian probability distribution is needed to choose the best decision: the expected utility criterion (see Chapter 2) needs to be adapted and generalised to the case of imprecise probabilities. The theory of coherent previsions [Walley, 1991] fortunately offers such a generalisation. As a general theory of imprecise probabilities, it also provides general methods for statistical inference, decision and risk analysis.

The first contribution of this chapter is to instantiate all these methods in the case of dialectical probabilities. We do so, so as to allow the direct exploitation of

dialectical probabilities for practical problems. The second contribution is that we relate the main criterion for decision making with imprecise probabilities to the classical criteria for decision making under strict uncertainty or risk (defined in Chapter 2). Finally, we provide a new application of argumentation and imprecise probabilities to finance and demonstrate experimentally the potential offered by dialectical probabilities for decision making in this domain.

The chapter is organised as follows. Section 6.1 is background and introduces the methods used under imprecise probabilities (in general) for solving problems of statistical inference, risk analysis and decision making. We instantiate these method in the particular case of dialectical probabilities in Section 6.2 for statistical inference, Section 6.3 for risk analysis, Section 6.4 and Section 6.5 for decision analysis and decision making. In Section 6.6, we study the (decision-theoretic) axioms fulfilled by the main decision criterion used in imprecise probability theory. In Section 6.7, we study the links between the two decision criteria used in imprecise probability theory (when making decisions with generic imprecise probabilities) and the classical decision criteria for decision making under strict uncertainty and risk. Section 6.8 presents an application of argumentation and dialectical probabilities to financial stock trading. We discuss related work in Section 6.9 and conclude in Section 6.10.

6.1 Background: using imprecise probabilities

All the definitions and results exposed in the section hold for imprecise probabilities in general and are not specific to dialectical probabilities. In Section 5.1.5, we have defined the natural extension $\underline{E}$ of a coherent lower prevision or lower probability $\underline{P} : \Omega \to \mathbb{R}$, where $\Omega = \{w_1, \ldots, w_n\}$ denotes as usual a finite set of n scenarios. Recall that the natural extension $\underline{E}$ is defined for any gamble $X : \Omega \to \mathbb{R}$.

6.1.1 Statistical inference

The natural extension can be used for *statistical inference*, i.e. for comparing the likelihood of events and for computing conditional probabilities. Indeed, for any two events $A, B \subseteq \Omega$, we can say according to [Walley, 1991] that

- A is more likely than B if $\underline{E}(A - B) \geq 0$,
- A is at least twice as likely as B if $\underline{E}(A - 2B) \geq 0$, and more generally, when $x \geq 1$, that
- A is at least x times as likely as B if $\underline{E}(A - xB) \geq 0$.

The conditional lower probability $\underline{P}(A|B)$ of A given B is given by the solution $z \in \mathbb{R}$ of the equation

$$\underline{E}(B\,[A - z]) = 0$$

called the *generalised Bayes rule* by [Walley, 1991], who showed that when the imprecise probability is Bayesian (precise), the generalised Bayes rule becomes identical to the standard Bayes rules of conditioning.

6.1.2 Decision analysis

The concept of natural extension can also be used for *decision analysis*, i.e. for comparing decisions and for decision making. Assume that D is a set of possible decisions and $U : D \times \Omega \to \mathbb{R}$ a value function, so that $U(d, w)$ measures the desirability of the consequence of d under scenario $w \in \Omega$ for the subject. Note here that U could for instance be determined using the method introduced in Chapter 4. The gamble X_d defined $\forall w \in \Omega$ as $X_d(w) = U(d, w)$ models the decision d. According to [Walley, 1991], we may say that d is preferred to d' whenever $\underline{E}(X_d - X_{d'}) \geq 0$. This means that the transaction where the subject gets X_d in exchange of $X_{d'}$ is acceptable. One may also say that d is twice as

much preferred as d' whenever $\underline{E}(X_d - 2X_{d'}) \geq 0$, and more generally that d is preferred λ times as much as d' whenever $\underline{E}(X_d - \lambda X_{d'}) \geq 0$. Intuitively, $\underline{E}(X_d - X_{d'})$ provides a robust under-estimation of the expected value of d in comparison to d'. When $\underline{P}$ is a Bayesian probability distribution, the expression of $\underline{E}(X_d)$ collapses to the expression of the expected utility:

$$\underline{E}(X_d) = E_P(X_d) = \sum_{i=1}^{n} P(w_i)U(d, w_i)$$

Thus, the concept of natural extension generalises the expected utility criterion of decision theory to the case of imprecise probabilities.

For decision making, one needs a strict preference relation to rank decisions.

Definition 6.1. *Let $\mathcal{X} = \{X_d \,|\, d \in D\}$. For any $X_d, X_{d'} \in \mathcal{X}$, one writes $X_d \succ X_{d'}$ if and only if*

- $\underline{E}(X_d - X_{d'}) > 0$, *or*
- $X_d \geq X_{d'}$ *and* $X_d \neq X_{d'}$.

Recall that for any gambles X and Y, the inequality $X \geq Y$ means that, $\forall w \in \Omega$, $X(w) \geq Y(w)$. Also, $X \neq Y$ means that, $\exists w \in \Omega$, $X(w) \neq Y(w)$.

Definition 6.2. *Let $X \in \mathcal{X}$. X is **inadmissible** when there exists $Y \in \mathcal{X}$ such that $Y \geq X$ and $Y \neq X$. Otherwise, X is said to be **admissible**.*

Clearly, the best decisions d should at least correspond to admissible gambles X_d. [Walley, 1991] proposes two criteria for choosing the best decision under imprecise probabilities: the *minimax* criterion and the *maximality* criterion. The maximality criterion is defined in terms of the upper natural extension $\overline{E}$, which is defined for any gamble X as $\overline{E}(X) = -\underline{E}(-X)$.

Definition 6.3. *The **minimax criterion** consists in choosing a decision d such that X_d is a **minimax gamble**, i.e.*

- X_d *is admissible, and*
- $\underline{E}(X_d)$ *is maximal.*

Definition 6.4. *The* ***maximality criterion*** *consists in choosing a decision d such that X_d is a* ***maximal gamble****, i.e.*

- X_d *is admissible, and*
- $\min_{d' \in D} \overline{E}(X_d - X_{d'})$ *is maximal.*

It is proved in [Walley, 1991] that the maximal gambles are exactly the gambles that are undominated under $\succ$. It is also proved that every minimax gamble is also a maximal gamble. However, it has been observed by [Berger, 1985] that minimax gambles often do not lead to particularly good decisions. Hence, decisions are most often made by applying the maximality criterion.

6.1.3 Financial risk analysis

Finally, upper previsions can be used in *risk analysis* to measure financial risk. Consider a gamble X such that for every scenario w, $X(w)$ represents the return of an investor after a period of time T if w occurs on a portfolio (i.e. a combination of financial instruments such as a bonds, stocks and options). Let R_F be the risk-free rate over T (typically of the order of 5% per year). Then, according to [Pelessoni and Vicig, 2001], the risk of X can be measured by

$$\rho(X) = \frac{1}{R_F} \overline{E}(-X)$$

[Pelessoni and Vicig, 2001] have proved that whenever $\underline{P}$ is a lower probability or a coherent lower prevision, the function ρ defines a coherent risk measure in the sense of [Artzner *et al.*, 1999], i.e. it satisfies for every gambles X and Y the following four axioms:

- $\forall c \in \mathbb{R}, \rho(X + cR_F) = \rho(X) - c$ (translation invariance)
- $\forall \lambda \in \mathbb{R}^+, \rho(\lambda X) = \lambda \rho(X)$ (positive homogeneity)
- if $X \leq Y$ then $\rho(Y) \leq \rho(X)$ (monotonicity)
- $\rho(X + Y) \leq \rho(X) + \rho(Y)$ (subadditivity)

It has also been proved that, conversely, for every coherent risk measure ρ, there exists a coherent lower prevision $\underline{P}$ such that for every gamble X,

$$\rho(X) = \frac{1}{R_F}\overline{E}(-X)$$

Moreover, note that, if ρ is a coherent risk measure, then minimising the risk $\rho(X)$ of a portfolio is equivalent to maximising $\underline{E}(X)$.

6.2 Statistical inference

In this and the next two sections, we instantiate the methods previously discussed in Section 6.1 to the special case of dialectical probabilities. In all these sections, we assume given a dialectical probability as defined in Sections 5.4 and 5.5, and still denote by K the epistemic risk-aversion parameter, I the total amount of information and $\hat{p}_A : \Omega \to [0, 1]$ the biased prior.

Statistical inference is important, as it allows to compare the likelihood of events and to assess the conditional probabilities of events. As presented in Section 6.1.1, in imprecise probability theory, an event A is said to be at least x times as likely as B whenever $\underline{E}(A - xB) \geq 0$. In the case of dialectical probabilities, we have

Theorem 6.1. *For any events $E, B \subseteq \Omega$ such that $B \neq \emptyset$, E can be said to be*

x times as likely as B, where

$$x = \begin{cases} \frac{I\hat{p}_A(E)+K}{I\hat{p}_A(B)+K} & \text{if } B \subseteq E \\ \frac{I\hat{p}_A(E)}{I\hat{p}_A(B)+K} & \text{otherwise} \end{cases}$$

Proof. Replacing $\underline{E}$ by its analytical expression (given in property 5.13), the condition $\underline{E}(E - xB) \geq 0$ becomes

$$IE_{\hat{p}_A}(E - xB) + K \min(E - xB) \geq 0$$

Since $B \neq \emptyset$, it is easy to see that the solutions x of this inequality must be such that $x \geq 0$. Let us consider the gamble $Y = E - xB$. We have, $\forall w \in \Omega$, $Y(w) = 1$ if $w \in (E - B)$, $Y(w) = 1 - x$ if $w \in E \cap B$ and $Y(w) = -x$ if $w \in (B - E)$. Since $x \geq 0$, we have

$$\min(Y) = \begin{cases} -x \text{ if } B - E \neq \emptyset \\ 1 - x \text{ if } B - E = \emptyset \end{cases}$$

If $B - E = \emptyset$ (or equivalently $B \subseteq E$), the condition becomes by linearity of the expectation operator $E_{\hat{p}_A}$ and its coincidence with $\hat{p}_A$ on events, $I(\hat{p}_A(E) - x\hat{p}_A(B)) + K(1 - x) \geq 0$, whose maximal solution is simply $x = (I\hat{p}_A(E) + K)/(I\hat{p}_A(B)+K)$. Otherwise, the condition becomes $I(\hat{p}_A(E) - x\hat{p}_A(B)) - Kx \geq 0$, whose maximal solution is simply $x = I\hat{p}_A(E)/(I\hat{p}_A(B) + K)$. □

Note that, as dialectical probabilities become more precise ($I \gg K$), the value of x converges towards a limit that is the intuitively expected ratio between the prior probabilities of the two events: $\hat{p}_A(E)/\hat{p}_A(B)$. This result confirms the suitability of dialectical probabilities for statistical reasoning.

Let us now consider conditional lower and upper probabilities induced by a

dialectical lower probability $\underline{P}$. We have

Theorem 6.2. *For any events $E, B \subseteq \Omega$ such that $\underline{P}(B) > 0$, the conditional lower probability $\underline{P}(E|B)$ is given by*

$$\underline{P}(E|B) = \begin{cases} 1 & \text{if } B \subseteq E \\ \frac{I\hat{p}_A(B \cap E)}{I\hat{p}_A(B)+K} & \text{otherwise} \end{cases}$$

Proof. $\underline{P}(E|B)$ is defined as the solution $z \in \mathbb{R}$ of the generalised Bayes rule $\underline{E}(B\,[E - z]) = 0$, which here is equivalent to

$$IE_{\hat{p}_A}(B\,[E - z]) + K \min(B\,[E - z]) = 0$$

It is easy to see that the only possible solutions of this equation are such that $z \in [0, 1]$. Let us then consider the gamble $Y = B\,[E - z]$. We have, $\forall w \in \Omega$, $Y(w) = 0$ if $w \notin B$, $Y(w) = 1 - z$ if $w \in B \cap E$ and $Y(w) = -z$ if $w \in (B - E)$. Since $z \in [0, 1]$,

$$\min(Y) = \begin{cases} -z \text{ if } B - E \neq \emptyset \\ 0 \text{ if } B - E = \emptyset \end{cases} = \begin{cases} -z \text{ if } B \not\subseteq E \\ 0 \text{ if } B \subseteq E \end{cases}$$

If $B \subseteq E$, the generalised Bayes rule is simplified to $\hat{p}_A(B) - z\hat{p}_A(B) = 0$ with solution $z = 1$. When $B \not\subseteq E$, the rule becomes $I(\hat{p}_A(B \cap E) - z\hat{p}_A(B)) - Kz = 0$ whose solution is $z = I\hat{p}_A(B \cap E)/(I\hat{p}_A(B) + K)$. □

Theorem 6.3. *For any events $E, B \subseteq \Omega$ such that $\underline{P}(B) > 0$, the conditional upper probability $\overline{P}(E|B)$ is given by*

$$\overline{P}(E|B) = \begin{cases} 0 & \text{if } B \cap E = \emptyset \\ \frac{I\hat{p}_A(B \cap E)+K}{I\hat{p}_A(B)+K} & \text{otherwise} \end{cases}$$

Proof. $\overline{P}(E|B)$ is defined as the solution z of the equation $\overline{E}(B\,[E-z]) = 0$, which is equivalent to

$$IE_{\hat{p}_A}(B\,[E-z]) + K\max(B\,[E-z]) = 0$$

It is easy to see that the only possible solutions of this equation are such that $z \in [0,1]$. Since $z \in [0,1]$,

$$\max(Y) = \begin{cases} 1-z & \text{if } B \cap E \neq \emptyset \\ 0 & \text{if } B \cap E = \emptyset \end{cases}$$

If $B \cap E = \emptyset$, the equation is simplified to $-z\hat{p}_A(B) = 0$ with solution $z = 0$. If $B \cap E \neq \emptyset$, the equation becomes $I(\hat{p}_A(B \cap E) - z\hat{p}_A(B)) + K(1-z) = 0$ whose solution is $z = (I\hat{p}_A(B \cap E) + K)/(I\hat{p}_A(B) + K)$. □

Given the expressions of $\overline{P}(E|B)$ and $\underline{P}(E|B)$ in these last two theorems, it is possible to observe that $\overline{P}(E|B)$ and $\underline{P}(E|B)$ are linked by the relation $\overline{P}(E|B) = 1 - \underline{P}(E^c|B)$ and thus are always the dual of each other. When dialectical probabilities become more precise ($I \gg K$), the lower and upper conditional probabilities converge to the conditional probability

$$\hat{p}_A(E|B) = \frac{\hat{p}_A(E \cap B)}{\hat{p}_A(B)}$$

This result again gives an intuitive confirmation of the suitability of dialectical probabilities for statistical reasoning.

6.3 Financial risk analysis

We assume here that $X : \Omega \to \mathbb{R}$ is a gamble such that, for every scenario w, $X(w)$ represents the return over a period T of some financial investment X if w occurs. As in Section 6.1.3, we denote by R_F the risk-free rate. Given a dialectical lower probability $\underline{P}$ defined on Ω, we have

Theorem 6.4. *The risk on the investment with uncertain returns $X : \Omega \to \mathbb{R}$ is given by*

$$\rho(X) = \frac{-1}{R_F}\left[\frac{I}{I+K}E_{\hat{p}_A}(X) + \frac{K}{I+K}\min(X)\right]$$

Proof. According to [Pelessoni and Vicig, 2001], the risk is $\rho(X) = \frac{1}{R_F}\overline{E}(-X)$. In the property 5.13 of Section 5.5, we have proved that, for any gamble Y,

$$\overline{E}(Y) = \frac{1}{K+I}(IE_{\hat{p}_A}(Y) + K\max(Y))$$

In particular, for $Y = -X$, we obtain

$$\overline{E}(-X) = \frac{1}{K+I}(IE_{\hat{p}_A}(-X) + K\max(-X))$$

which by linearity of the expectation and duality of max and min operators can be rewritten as

$$\overline{E}(-X) = -\left[\frac{I}{K+I}E_{\hat{p}_A}(X) + \frac{K}{K+I}\min(X)\right]$$

hence the expression

$$\rho(X) = \frac{-1}{R_F}\left[\frac{I}{I+K}E_{\hat{p}_A}(X) + \frac{K}{I+K}\min(X)\right]$$

□

Here, the term $\min(X)$ represents the worst-case scenario return. The term in brackets corresponds to a linear combination of the expected return $E_{\hat{p}_A}(X)$ with the worst-case scenario return. The weight put on the worst-case scenario return is as small as the amount of information I is great in comparison to K. The risk measure ρ takes positive values when this linear combination is negative, and takes negative values otherwise. Thus, negative risk means that the investor can be confident that his investment will generate a positive return. In particular, if all his wealth is invested at the risk-free rate, then the risk is $\rho(R_F) = -1$. Finally, and as already said in Section 6.1.3, minimising the risk $\rho(X)$ of investment X is equivalent to maximising $\underline{E}(X)$, which represents a pessimistic estimation of the expected return on X. Thus, dialectical probabilities determine coherent financial risk measures that have an intuitive interpretation.

6.4 Decision analysis

Decision analysis amounts to comparing decisions and decision making amounts to selecting optimal decisions. Let us assume given a set of possible decisions D and a value function $U : D \times \Omega \to \mathbb{R}$. As discussed in Section 6.1.2, each decision d can be associated with the gamble X_d defined $\forall w \in \Omega$ as $X_d(w) = U(d, w)$. Assuming given as in the previous sections a dialectical probability, we have

Theorem 6.5. *Let d and d' be two decisions in D. d is preferred at least λ times as much as d' whenever*

$$IE_{\hat{p}_A}(X_d - \lambda X_{d'}) + K \min(X_d - \lambda X_{d'}) \geq 0$$

Proof. As discussed in Section 6.1.2, d is preferred at least λ times as much as

d' whenever

$$\underline{E}(X_d - \lambda X_{d'}) \geq 0$$

which, by property 5.13, is equivalent to

$$\frac{I}{K+I} E_{\hat{p}_A}(X_d - \lambda X_{d'}) + \frac{K}{K+I} \min(X_d - \lambda X_{d'}) \geq 0$$

and, since $K + I \geq I \geq 1 > 0$, this is also equivalent to

$$I E_{\hat{p}_A}(X_d - \lambda X_{d'}) + K \min(X_d - \lambda X_{d'}) \geq 0$$

□

Let us see now how to select the best decisions when using dialectical probabilities. We have seen, in Section 6.1.2, that two criteria could be used, namely the maximin criterion and the maximality criterion. No matter the criterion adopted, the best decisions must correspond to admissible gambles. Let us first remark that, in the case of dialectical probabilities, admissible gambles are undominated in the sense of *imprecise stochastic dominance* [Jaffray, 1999].

Definition 6.5. *Given a credal set $\mathcal{D}$ and two decisions d and d', we say that d* ***stochastically dominates*** *d' if and only if $\forall c \in \mathbb{R}$ and $\forall P \in \mathcal{D}$*

$$P(\{w | X_d(w) \leq c\}) \leq P(\{w | X_{d'} \leq c\})$$

and the inequality is strict for some $c_0 \in \mathbb{R}$ and $P_0 \in \mathcal{D}$.

In the case of dialectical probabilities, we have the following theorem.

Theorem 6.6. *Every inadmissible gamble $X_d \in \mathcal{X}$ is stochastically dominated.*

Proof. It is sufficient to show that if $X_{d'} \geq X_d$ and $X_{d'} \neq X_d$, then d' stochastically dominates d. Since we assume a dialectical probability, the credal set

considered is (see beginning of Section 5.6) $\mathcal{D}_A$, the maximal ϵ_A-contamination of the biased prior $\hat{p}_A$ with contamination level $\epsilon_A = K/(K+I)$. If $X_{d'} \geq X_d$ and $X_{d'} \neq X_d$, then $\forall c \in \mathbb{R}$

$$\{w|X_{d'}(w) \leq c\} \subseteq \{w|X_d \leq c\} \quad \text{and} \quad P(\{w|X_{d'}(w) \leq c\}) \leq P(\{w|X_d \leq c\})$$

for any Bayesian probability P. Since $X_{d'} \neq X_d$, there exists $c_0 \in \mathbb{R}$ for which the set inclusion is strict: $\{w|X_{d'}(w) \leq c_0\} \subset \{w|X_d \leq c_0\}$. Let $P_0 = (1 - \epsilon_A)\hat{p}_A + \epsilon_A q_0$, where q_0 is the uniform probability distribution $q_0 : \Omega \to \{1/|\Omega|\}$. P_0 obviously belongs to $\mathcal{D}_A$. $\forall w \in \Omega : P_0(w) > 0$ therefore P_0 is strictly monotonic and $P_0(\{w|X_{d'}(w) \leq c_0\}) < P_0(\{w|X_d \leq c_0\})$. □

We now propose a general algorithm for computing maximal gambles (see Algorithm 4 here-after). Maximal gambles correspond to the best decisions under the maximality criterion, which constitutes the main criterion in the theory of imprecise probabilities, as discussed in Section 6.1.2. The algorithm we propose is general, in the sense that it applies to any imprecise probability and not only to dialectical probabilities. In [Walley, 1991], no algorithm is described to allow the computation of maximal gambles so this algorithm can be seen as a contribution to the theory of imprecise probabilities. Our algorithm is based on an optimisation method known as *semi-infinite programming* [Blankenship and Falk, 1976].

Theorem 6.7. *Algorithm 4 returns a maximal gamble $X^* \in \mathcal{X}$.*

Proof. It is proved in [Walley, 1991] that if $\mathcal{X}$ is a compact set under the supremum-norm topology and $\underline{P}$ is coherent, then there exists at least one maximal gamble in $\mathcal{X}$. In order to find a maximal gamble $X \in \mathcal{X}$, one may

Algorithm 4 Finds a maximal gamble $X^* \in \mathcal{X}$ (optimal decision) given a gamble $X \in \mathcal{X}$ as input.

1: **Initialise**: set $k = 0$ and $\mathcal{Y}^0 = \{X\}$.
2: **Compute**: $x_k = (X_k, m_k)$ which maximises $m_k \in \mathbb{R}$ subject to $X_k \in \mathcal{X}$, $P \in \mathcal{D}$ and $\forall Y \in \mathcal{Y}^k$, $m_k \leq E_P(X_k - Y)$.
3: **Compute**: $Y_{k+1} \in \mathcal{X}$ which minimises $\overline{E}(X_k - Y_{k+1})$.
4: **Check**: if $m_k \leq \overline{E}(X_k - Y_{k+1})$ then go to step 5, else update $\mathcal{Y}^{k+1} = \mathcal{Y}^k \cup \{Y_{k+1}\}$ and return to step 2 with $k = k + 1$.
5: **Initialise**: $X = x_k$.
6: **Compute**: $X^* \in \mathcal{X}$ which maximises $\sum_{i=1}^{|\Omega|} X_i^*$ subject to $X^* \geq X$.

proceed as follows. First, find a solution X of the maxmin problem

$$\max_{X \in \mathcal{X}} \min_{Y \in \mathcal{X}} \overline{E}(X - Y)$$

and then transform X into an admissible gamble $X^* \geq X$. When $\mathcal{X}$ is compact under the supremum norm, the existence of a maximal gamble implies that the solution X of the maxmin problem will be such that $\forall Y \in \mathcal{X}$, $\overline{E}(X - Y) \geq 0$. By monotonicity of $\overline{E}$, it is clear that any $X^* \geq X$ would also be such that $\forall Y \in \mathcal{X}$, $\overline{E}(X^* - Y) \geq 0$. Hence, X^* is guaranteed to be a maximal gamble.

The initial maxmin problem is equivalent to solving

$$\max_{(X,m) \in \mathcal{X} \times \mathbb{R}} m$$

subject to the constraint that $\forall Y \in \mathcal{X}$, $m \leq \overline{E}(X - Y)$. In this problem, the variable m plays the role of the quantity $\min_{Y \in \mathcal{X}} \overline{E}(X - Y)$ which depends on X and that we want to maximise. Introducing the function f defined for every $x = (X, m) \in \mathcal{X} \times \mathbb{R}$ as $f(x) = -m$ and the function g defined for every (x, y) where $x = (X, m) \in \mathcal{X} \times \mathbb{R}$ and $y = Y \in \mathcal{X}$ as $g(x, y) = m - \overline{E}(X - Y)$, the initial maxmin problem can be rewritten under the form of the following semi-infinite program

$$\min_x f(x)$$

subject to the constraint that $\forall y$, $g(x,y) \leq 0$. Since the domains of the variables x and y are non-empty and compact and the functions f and g are continuous, the maxmin problem can be solved using the algorithm of [Blankenship and Falk, 1976]

1. Initialise: set $k = 0$ and choose $\mathcal{Y}^0 \subseteq \mathcal{X}$.
2. Compute: x_k which minimises $f(x)$ subject to $\forall y \in \mathcal{Y}^k$, $g(x,y) \leq 0$.
3. Compute: y_{k+1} which maximises $g(x^k, y)$.
4. Check: if $g(x^k, y^{k+1}) \leq 0$, then go to step 5, else update $\mathcal{Y}^{k+1} = \mathcal{Y}^k \cup \{y_{k+1}\}$ and return to step 2 with $k = k + 1$.
5. Return: x_k.

and resulting in steps 1 to 5 of Algorithm 4. The final transformation of the solution of the maxmin problem into an admissible gamble corresponds to step 6 of the algorithm. □

This algorithm will be used later on in the application Section 6.8 for computing optimal trading decisions.

6.5 Forcing decisions

Typically, dialectical probabilities are imprecise, so that $\succ$ (see definition 6.1) is not a complete ordering of $\mathcal{X}$ and there exists more than one maximal gamble. As a result, the decision maker may still be undecided. The subject's indeterminacy (inability to estimate probabilities precisely) is the main cause of his undecidedness. Therefore, in this section, we review a number of existing methods available for reducing the indeterminacy to zero, or, in other words, for transforming imprecise probabilities into Bayesian probability distributions.

All these methods only apply to imprecise probabilities that correspond to belief functions. Fortunately, dialectical probabilities are belief functions (see Section 5.6). We are thus going to apply all the existing methods for transforming a belief function into a Bayesian probability distribution to dialectical probabilities.

First, recall from Section 5.6 theorem 5.6 that a dialectical probability is a belief function whose evidence mass function $m_A : 2^\Omega \rightarrow [0,1]$ is defined for every event X by

$$m_A(X) = (1 - \epsilon_A)\hat{p}_A(X)\delta(|X|, 1) + \epsilon_A \delta(|X|, |\Omega|)$$

where δ is the Kronecker symbol (see the proof of property 5.5 for a definition of this).

Definition 6.6. *The **pignistic probability** is the Bayesian probability distribution BetP obtained by application of the pignistic transform [Smets, 1990] and is defined for any belief function with evidence mass function m by*

$$BetP(w) = \sum_{X \subseteq \Omega, w \in X} \frac{m(X)}{|X|}$$

In the case of the dialectical probability, we straightforwardly have by definition of m_A that

Property 6.1. *The pignistic probability is given $\forall w \in \Omega$ by*

$$BetP(w) = (1 - \epsilon_A)\hat{p}_A(w) + \frac{\epsilon_A}{|\Omega|}$$

Notice that $BetP = (1 - \epsilon_A)\hat{p}_A + \epsilon_A u$, where $u : \Omega \rightarrow \{1/|\Omega|\}$ is the uniform probability distribution. Thus, it appears that for the dialectical probability, the pignistic probability is a linear combination of the biased prior $\hat{p}_A$ and the uniform distribution u. Recall that u is known to be [Klir, 2006] the Bayesian

probability distribution P that maximises the Shannon entropy

$$S(P) = \sum_{i=1}^{n} -P(w_i)\log_2(P(w_i))$$

Definition 6.7. *The* ***cautious probability*** *is the Bayesian probability distribution $CautP$ obtained by application of the cautious probabilistic transform [Daniel, 2006] and is defined for any belief function with evidence mass function m by*

$$CautP(w) = \frac{\sum_{X\subseteq\Omega, w\in X} m(X)}{\sum_{v\in\Omega}\sum_{X\subseteq\Omega, v\in X} m(X)}$$

In the case of the dialectical probability, we have

Property 6.2. *The cautious probability is given $\forall w \in \Omega$ by*

$$CautP(w) = \frac{(1-\epsilon_A)\hat{p}_A(w) + \epsilon_A}{1+\epsilon_A(|\Omega|-1)}$$

Proof. Observe that

- first $\sum_{X\subseteq\Omega, w\in X} m_A(X) = (1-\epsilon_A)\hat{p}_A(w) + \epsilon_A$, and finally that
- $\sum_{v\in\Omega}\sum_{X\subseteq\Omega, v\in X} m_A(X) = (1-\epsilon_A)\sum_{v\in\Omega}\hat{p}_A(v) + \epsilon_A|\Omega| = 1+\epsilon_A(|\Omega|-1)$.

□

It is easy to check that $CautP = (1-\epsilon')\hat{p}_A + \epsilon' u$, where

$$\epsilon' = \frac{\epsilon_A|\Omega|}{1+\epsilon_A(|\Omega|-1)}$$

and u is the uniform probability distribution. The cautious probability is thus also a linear combination of $\hat{p}_A$ and u, although in that case, the weight associated to u is ϵ' instead of ϵ_A. The ratio between these weights is

$$\frac{\epsilon'}{\epsilon_A} = \frac{|\Omega|}{1+\epsilon_A(|\Omega|-1)} \geq \frac{|\Omega|}{|\Omega|} = 1$$

This means that ϵ' is an amplification of ϵ_A. Understandably, $CautP$ is a more cautious version than the pignistic probability.

Definition 6.8. *The **proportional probability** is the Bayesian probability distribution PropP obtained by application of the proportional probabilistic transformation [Daniel, 2006] and is defined for any belief function with evidence mass function m by*

$$PropP(w) = m(\{w\}) \sum_{X \subseteq \Omega, w \in X} \frac{m(X)}{\sum_{v \in X} m(\{v\})}$$

In the case of the dialectical probability, we have

Property 6.3. *The proportional probability is given $\forall w \in \Omega$ by*

$$PropP(w) = \hat{p}_A(w)$$

Proof. By definition of m_A, we have $PropP(w) =$

$$(1-\epsilon_A)\hat{p}_A(w)\left(\sum_{X \subset \Omega, w \in X} \frac{(1-\epsilon_A)\hat{p}_A(X)\delta(|X|,1)}{\sum_{v \in X}(1-\epsilon_A)\hat{p}_A(v)} + \frac{\epsilon_A}{\sum_{v \in \Omega}(1-\epsilon_A)\hat{p}_A(v)}\right) =$$

$$(1-\epsilon_A)\hat{p}_A(w)\left(\frac{(1-\epsilon_A)\hat{p}_A(w)}{(1-\epsilon_A)\hat{p}_A(w)} + \frac{\epsilon_A}{1-\epsilon_A}\right) =$$

$$(1-\epsilon_A)\hat{p}_A(w)\left(1 + \frac{\epsilon_A}{1-\epsilon_A}\right) = \hat{p}_A(w)$$

□

The proportional probability is simply equal to the biased prior $\hat{p}_A$. Here, the prior is not combined with u, which makes the proportional probability the least cautious of all probabilities seen so far.

Definition 6.9. *The **plausible probability** is the Bayesian probability distri-*

bution Pl_P obtained by application of the plausibility transformation method [Cobb and Shenoy, 2006] and is defined for any plausibility function Pl by

$$Pl_P(w) = \frac{Pl(\{w\})}{\sum_{v\in\Omega} Pl(\{v\})}$$

In the case of the dialectical probability, we have

Property 6.4. *The plausible probability is given* $\forall w \in \Omega$ *by*

$$Pl_P(w) = \frac{(1-\epsilon_A)\hat{p}_A(w) + \epsilon_A}{1 + \epsilon_A(|\Omega| - 1)}$$

Proof. We have for a dialectical probability

$$Pl(\{w\}) = 1 - Bel(\{w\}^c) = 1 - (1-\epsilon_A)\hat{p}_A(\{w\}^c) = 1 - (1-\epsilon_A)(1-\hat{p}_A(w))$$

so $Pl_P(w) = (1-\epsilon_A)\hat{p}_A(w) + \epsilon_A$. Clearly,

$$\sum_{v\in\Omega} Pl_P(\{v\}) = (1-\epsilon_A) + \epsilon_A|\Omega| = 1 + \epsilon_A(|\Omega| - 1)$$

□

The plausible probability is remarkably identical to the cautious probability $CautP$ and consequently also represents a very cautious Bayesian substitute for dialectical probabilities.

6.6 Axioms fulfilled by the maximality criterion

Recall that the main criterion for choosing the best decisions under any imprecise probability, and in particular for dialectical probabilities, is the maximality criterion. By definition, this criterion consists in choosing an admissible gamble

X_d maximising $\min_{d'} \overline{E}(X_d - X_{d'})$ (see Section 6.1.2). We would like to know which rationality axioms from decision theory amongst those of Section 2.1.2 are satisfied by this criterion. Here, we do not assume that the imprecise probability is a dialectical probability, but consider the general case. This decision-theoretic analysis has not been given by [Walley, 1991]. Note that, for this analysis to be possible, we need to assume that the set of possible decisions is finite.

Theorem 6.8. *The maximality criterion fulfils axioms 2.1 (complete ranking), 2.2 (independence of labelling), 2.3 (independence of value scale), 2.4 (strong domination), 2.6 (independence of addition of a constant to a column) and 2.7 (independence of column duplication) under strict uncertainty (vacuous previsions) but does not fulfil axiom 2.5 (independence of irrelevant alternatives).*

Proof. Axiom 2.1. By finiteness of D, the value $\min_{d'} \overline{E}(X_d - Y_{d'})$ exists for every $d \in D$ and any valuation in $\mathbb{R}$ of decisions induces a complete ranking of the decisions.

Axiom 2.2. The ordering of the set of scenarios and set of decisions has no influence on $\min_{d'} \overline{E}(X_d - Y_{d'})$ and subsequently on the ordering obtained.

Axiom 2.3. Let $f(x) = ax + b$ be an affine transformation of the utility scale with $a, b \in \mathbb{R}$ and $a > 0$. For every Bayesian probability P and gamble X, we have $E_P(f(X)) = \sum_{w\in\Omega} P(w)f(X(w)) = \sum_{w\in\Omega} P(w)(aX(w) + b) = aE_P(X) + b = f(E_P(X))$. For any set $\mathcal{D}$ of Bayesian probabilities (credal set), we have $\overline{E}(f(X)) = \sup_{P\in\mathcal{D}} E_P(f(X)) = \sup_{P\in\mathcal{D}} f(E_P(X)) = f(\overline{E}(X))$ since f is strictly monotone increasing. Similarly, we have $\min_{d'} \overline{E}(f(X_d) - f(X_{d'})) = f(\min_{d'} \overline{E}(X_d - X_{d'}))$, so the strict preference of d_1 over d_2 is unchanged by application of the affine transformation f to the utility scale.

Axiom 2.4. If $X_{d_1}(w) > X_{d_2}(w)$ for every scenario $w \in \Omega$, then there exists a constant $\mu > 0$ such that $X_{d_1} \geq X_{d_2} + \mu$. For every $d' \in D$, we have $\overline{E}(X_{d_1} - Y_{d'}) \geq \overline{E}(X_{d_2} - Y_{d'} + \mu) = \overline{E}(X_{d_2} - Y_{d'}) + \mu$. Clearly, we also have

$\min_{d'} \overline{E}(X_{d_1} - Y_{d'}) \geq \min_{d'} \overline{E}(X_{d_2} - Y_{d'}) + \mu$ and, by consequence, d_1 is strictly preferred to d_2.

Axiom 2.5. Take for instance the vacuous prevision $\overline{E} = \max$. Assume $\Omega = \{w_1, w_2\}$ and consider the gambles X_{d_1}, X_{d_2}, X_{d_3}, and X_{d_4} defined by $X_{d_1} = 0$, $X_{d_2} = 1$, $X_{d_3}(w_1) = 0$, $X_{d_3}(w_2) = 1/2$, $X_{d_4}(w_1) = 1$, $X_{d_4}(w_2) = 3/2$. If $D = \{d_1, d_2, d_3\}$, then we have the ranking $d_1 < d_3 < d_2$:

- $\min_{d' \in D} \overline{E}(X_{d_1} - X_{d'}) = \min(0, -1, 0) = -1$
- $\min_{d' \in D} \overline{E}(X_{d_2} - X_{d'}) = \min(1, 0, 1) = 0$
- $\min_{d' \in D} \overline{E}(X_{d_3} - X_{d'}) = \min(1/2, -1/2, 0) = -1/2$

but if we add d_4 and $D = \{d_1, d_2, d_3, d_4\}$ we obtain instead $d_1 = d_3 < d_2$:

- $\min_{d' \in D} \overline{E}(X_{d_1} - X_{d'}) = \min(0, -1, 0, -1) = -1$
- $\min_{d' \in D} \overline{E}(X_{d_2} - X_{d'}) = \min(1, 0, 1, 0) = 0$
- $\min_{d' \in D} \overline{E}(X_{d_3} - X_{d'}) = \min(1/2, -1/2, 0, -1) = -1$

so the strict preference of d_3 over d_1 depends on d_4.

Axiom 2.6. Consider a utility function $U : D \times \Omega \to \mathbb{R}$ and U' such that for every d and $w \neq w_j$ we have $U'(d, w) = U(d, w)$ and $U'(d, w_j) = U(d, w_j) + c$, where c is a constant (this is equivalent to adding the constant c to the utility of outcomes in column j in the decision table). Define $X_d(w) = U(d, w)$ and $X'_d = U'(d, w)$ for every $d \in D$ and $w \in \Omega$. Remark that, for every $w \in \Omega$, we have the equality $X_d(w) - X_{d'}(w) = X'_d(w) - X'_{d'}(w)$ and therefore, the strict preference of d_1 over d_2 under U is equivalent to strict preference of d_1 over d_2 under U'. The addition of a constant to the utility of every outcome under a chosen scenario does not affect the strict preference of a decision over another one.

Axiom 2.7. The question of independence of the preferences with respect to column duplication only makes sense in the context of decision making under strict uncertainty, as the duplication of a scenario in Ω would otherwise violate the property of disjointedness of the universe Ω. Thus, assuming that $\overline{E} = \max$ (vacuous upper prevision), it is clear that the values $\min_{d'} \overline{E}(X_d - X_{d'})$ are unchanged (the max of a list of numbers does not change if one of the numbers is duplicated) and the same subsequently holds for the preferences over the decisions. □

So the maximality criterion for imprecise probabilities fulfils the same set of axioms as Savage's minimax regret criterion (see Section 2.1.2). We will examine more closely in the next section the relationship between these two criteria.

6.7 Link with decision criteria under strict uncertainty and risk

In decision making under strict uncertainty, no probabilistic information is assumed to be available. Reasoning under strict uncertainty is (intuitively) equivalent to reasoning under maximally imprecise probabilities. These probabilities, called the *vacuous lower and upper probabilities* are defined respectively for every event E as $\min(E)$ and $\max(E)$. Clearly, vacuous probabilities convey no probabilistic information. It is thus a natural question to check whether the minimax and maximality criteria under vacuous probabilities correspond to any of the criteria for decision making under strict uncertainty of Section 2.1.2 and, if not, which are the properties that distinguish them. Importantly, one should bear in mind that dialectical probabilities become vacuous as the confidence level $(1 - x)100\%$ imposed by the decision maker becomes close to 100% or when the ratio K/I is very large (see Section 5.6).

Property 6.5. *Under the vacuous previsions, all decisions that are optimal according to the minimax criterion are also optimal according to Wald's maximin return criterion.*

Proof. A decisions d is optimal according to the minimax criterion if and only if the gamble X_d is admissible and maximises $\underline{E}(X_d)$. When the prevision is vacuous, $\underline{E}$ is simply equal to min. So, d is optimal under a vacuous prevision if and only if X_d is admissible and d has a maximal security level $s_d = \min_{w\in\Omega} U(d,w)$. □

The converse however does not hold in general because of the admissibility condition. The minimax criterion is thus stronger than Wald's criterion.

Although Savage's regret criterion and the maximality criterion satisfy the same decision-theoretic axioms (result of the previous section), the two criteria differ. Indeed,

Property 6.6. *Under the vacuous previsions,*

- *the maximality criterion is equivalent to choosing a decision d such that X_d is admissible and solves*

$$\max_d \min_{d'} \max_w (U(d,w) - U(d',w))$$

- *Savage's regret criterion is equivalent to choosing a decision d that solves*

$$\max_d \min_{d'} \min_w (U(d,w) - U(d',w))$$

Proof. A decision d is optimal according to the maximality criterion if and only if the gamble X_d is admissible and maximises $\min_{d'\in D} \overline{E}(X_d - X_{d'})$. Under the vacuous previsions, $\overline{E}$ is simply equal to max. So, a decision d is optimal

according to the maximality criterion if and only if X_d is admissible and d maximises

$$\min_{d' \in D} \max_{w \in \Omega} (U(d, w) - U(d', w))$$

In other words, the maximality criterion consists in choosing a decision d such that X_d is admissible and d solves

$$\max_{d} \min_{d'} \max_{w} (U(d, w) - U(d', w))$$

Savage's regret criterion consists in choosing a decision d that minimises

$$\rho_d = \max_{w} \max_{d'} (U(d', w) - U(d, w))$$

Savage's criterion is thus equivalent to maximising $-\rho(d)$, and thus can be formulated equivalently as choosing a decision d that solves

$$\max_{d} \min_{w} \min_{d'} (U(d, w) - U(d', w))$$

or since the min operators always commute

$$\max_{d} \min_{d'} \min_{w} (U(d, w) - U(d', w))$$

□

Thus, the two criteria differ in their inner parts: Savage's regret involves a min and the maximality criterion a max. The regret is based on worst-case scenario comparisons to other decisions, whilst the maximality criterion is based on best-case scenario comparisons. Thus, intuitively, the maximality criterion is less "pessimistic" than Savage's regret criterion.

Laplace's principle of insufficient reason cannot be directly compared to the

minimax and maximality criteria. Nevertheless, we can still state that

Property 6.7. *Under the vacuous previsions, the Bayes gambles X_d determined by the pignistic, cautious or plausible and proportional probabilities are in one-to-one correspondence with the decisions d that are optimal according to Laplace's principle of insufficient reason.*

Proof. It can easily be shown that the pignistic, cautious and proportional probabilities that derive from the vacuous probability simply correspond to the uniform distribution. The corresponding Bayes gambles X_d are thus by definition the gambles with maximal expected utility under the uniform distribution. Hence, they are trivially in one-to-one correspondence with the decisions d that are optimal according to Laplace's principle of insufficient reason. □

Let us now study the case of decision making under risk, where a precise or Bayesian probability distribution P is available. We can show that both the minimax criterion and the maximality criterion are compatible with the expected utility criterion.

Property 6.8. *In decision under risk, the decisions that are optimal according to the minimax criterion are also optimal under the expected utility criterion.*

Proof. The credal set is the singleton $\mathcal{D} = \{P\}$, so the natural extension $\underline{E}$ is simply equal to the expect utility E_P. The converse does not hold in general because minimax gambles must fulfil the admissibility condition. □

Property 6.9. *In decision under risk, the decisions that are optimal according to the maximality criterion are also optimal under the expected utility criterion.*

Proof. Since in that case $\overline{E} = E_P$ and E_P is linear, we have

$$\overline{E}(X_d - X_{d'}) = E_P(X_d - X_{d'}) = E_P(X_d) - E_P(X_{d'})$$

and we may write

$$\min_{d'} \overline{E}(X_d - X_{d'}) = E_P(X_d) - \max_{d'} E_P(X_{d'})$$

The ranking obtained using the maximality criterion is perfectly equivalent to the one obtained using the expected utility criterion. Again, the converse does not hold (decisions with maximal expected utility are not always optimal under the maximality criterion) because the admissibility condition may not hold. □

Both the minimax and maximality criteria generalise the expected utility criterion to imprecise probabilities. The minimax criterion degenerates to the pessimistic criterion of Wald and the maximality criterion degenerates to a more optimistic version than Savage's regret criterion.

6.8 Application to stock trading

The purpose of this section is to provide an experimental validation of the substantive goodness of dialectical probabilities and the suitability of the maximality criterion. We also wish to illustrate the use of abstract argumentation in finance and assess its added value.

6.8.1 Description of the problem

One of the most basic activities for investors in financial markets consists in buying and selling stocks, an activity referred to as stock trading. This application is restricted to single stock trading, although investors normally compose portfolios with tens or hundreds of stocks, using optimisation techniques such as mean-variance portfolio selection [Markowitz, 1952].

We will consider a fixed stock (American Express) and the repeated decision problem of buying or selling the stock on a daily basis. We assume that the

present and past values of the stock's prices can be observed by the investor. We assume that the investor has an initial budget of 1, that the investor is not allowed to short-sell the stock [1] nor use leverage (borrow money) to magnify his profits. Scenarios can be generated from a sample of past daily returns over an horizon of one year by clustering (we use the deterministic annealing algorithm of [Rose, 1998]). We use as benchmark B the period of trading days starting on January, 2nd 2003 and ending on September, 2nd 2008 (see Figure 6.1) and we will use the Sharpe ratio [Sharpe, 1966]

$$S = \frac{\mu - R_F}{\sigma}$$

where

- μ is the annualised mean of the daily investment returns,
- σ is the annualised volatility of the returns and
- $R_F = 5\%$ is the yearly risk-free rate

as a measure of performance for investment strategies.

We thus consider a single stock and denote the series of its adjusted daily closing prices $x_0, x_1, \ldots, x_i, \ldots$. The initial price of the stock is x_0 and for every integer i, x_i denotes the price of the stock per unit after i periods. All periods have a same duration of $T = 1$ trading day. We will denote by t_0 the date of the first trading day and for every integer i the date after iT trading days have elapsed, so on a calendar with trading days only the time $t_i = t_0 + iT$. At time t_0, the investor is assumed to start with a budget of 1 and no stock. This sum or *cash* is placed in the bank at a constant risk-free rate of R_F expressed in

[1] Short-selling a stock means borrowing the stock from someone, selling it immediately after and re-buying it later in time at a cheaper price, so as to return it to its owner. Short-selling allows to make money on a stock whose price is falling. It is mathematically equivalent to investing on a negative amount of the stock.

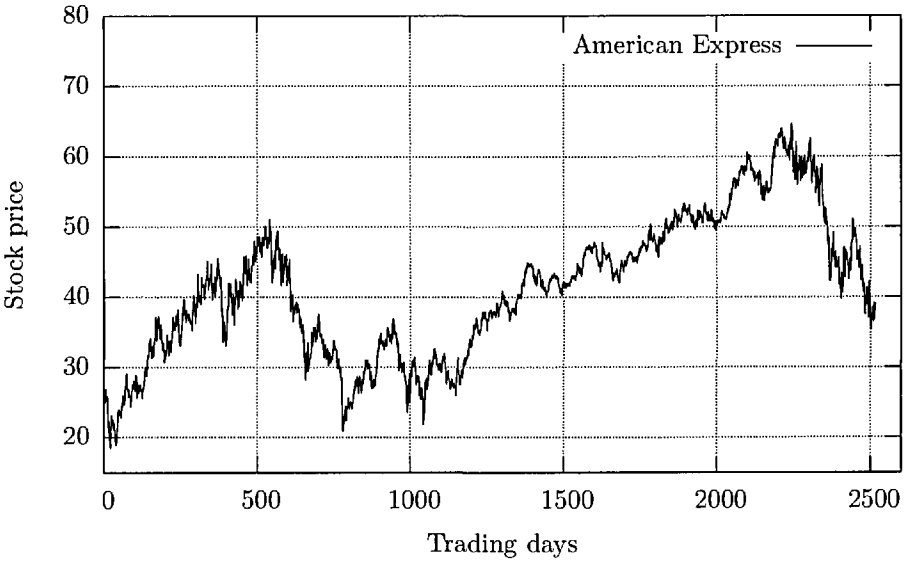


Figure 6.1: American express stock from August, 11 1998 to August, 8 2008.

percents per period T. At any time t_i, where $i \geq 0$, the investor may use part of his cash to buy an amount of $q_i \in \mathbb{R}$ units of the stock. If $q_i < 0$, we will say that the investor sells stocks, and if $q_i = 0$, that he does nothing. No buying or selling is permitted intra-period. Buying and selling stocks incurs transaction costs usually proportional to the amount of the transaction. Thus, in order to buy (or sell) q_i units of the stocks, the investor must pay $|q_i|x_iTC/100$ in cash, where TC is the transaction cost expressed in percents. We denote by C_i the amount of cash at time t_i and Q_i the quantity of stock owned by the investor in units before any decision has been made at that time. The total gross wealth of the investor is then given by $W_i = C_i + Q_i x_i$.

$$\begin{cases} C_0 = 1, Q_0 = 0, W_0 = 1 \\ C_{i+1} = C_i - |q_i| \, x_i \, (TC/100) - q_i x_i \\ Q_{i+1} = Q_i + q_i \\ W_{i+1} = C_{i+1} + Q_{i+1} x_{i+1} \end{cases}$$

The performance of the investor depends on the evolution of his gross wealth W_i. A useful metric in quantitative finance [Lhabitant, 2001] to measure the performance of an investor is the Sharpe ratio [Sharpe, 1966], given earlier. A high Sharpe ratio value means that the average yearly growth rate of the wealth is large compared to the irregularity of its variations. As said earlier, the Sharpe ratio is given by the formula

$$S = \frac{\mu - R_F}{\sigma}$$

The two parameters μ and σ of this formula are precisely computed in the following way. First, compute the sequence of *log-returns* for every $i > 0$ as $r_i = \ln(1 + \frac{W_i - W_{i-1}}{W_i - 1})$. Then, compute the log-return $\hat{r}$ as the mean (arithmetic average) of the sequence (r_i). Compute the log-volatility v as the standard deviation of the sequence (r_i). Let $x = 252/T$ be the annualization factor. Compute the yearly log-return r_y as $\hat{r}x$ and the yearly log-standard deviation v_y as $v\sqrt{x}$. Finally, compute

- the aggregated yearly return $\mu = e^{r_y} - 1$, and
- the yearly volatility $\sigma = e^{v_y} - 1$.

The Sharpe ratio is not much significant if calculated on an horizon of less than 5 years and reporting periods T of more than a quarter of a year. A negative Sharpe ratio means that the investor under-performs the risk-free rate, and a positive one that the investor on the contrary outperforms it. Historically, reference market indexes (which are dynamically updated portfolios) such as

S&P 500 and FTSE 100 realise Sharpe ratios going up to almost 0.5 and (at the time of writing this thesis) the world's top hedge funds have Sharpe ratios of value slightly above 1.

6.8.2 Building arguments from technical analysis

In finance, the forecast of stock prices is based on fundamental analysis and/or technical analysis. Fundamental analysis is defined (InvestorWords.com) as a method of security valuation which involves examining the company's financials and operations, especially sales, earnings, growth potential, assets, debt, management, products, and competition. Fundamental analysis takes into consideration only those variables that are directly related to the company itself, rather than the overall state of the market or technical analysis data. Technical analysis is a method of evaluating securities by relying on the assumption that market data, such as charts of price, volume, and open interest, can help predict future (usually short-term) market trends. Unlike fundamental analysis, the intrinsic value of the security is not considered. Technical analysts believe that they can accurately predict the future price of a stock by looking at its historical prices and other trading variables. Technical analysis assumes that market psychology influences trading in a way that enables predicting when a stock will rise or fall. For that reason, many technical analysts are also market timers, who believe that technical analysis can be applied just as easily to the market as a whole as to an individual stock.

We adopt here technical analysis for its simplicity of use. For any given stock and time, we seek to identify the current stock price trend and potential risk of that trend being interrupted and then reversed (the phenomenon is called a *divergence*). Trends and divergences allow to build a number of useful forecast arguments. To build such arguments, one looks at existing technical indicators

(basically functions of the past prices of the stock and volumes exchanged). We use here MACD (moving average convergence/divergence) and RSI (relative strength index) as indicators. Mitigation arguments can be built by looking at confirmation signals, which are indicators used in technical analysis to back up or invalidate the forecast of technical analysts. We use here moving average and centerline crossovers to build mitigation arguments. Denoting by x_i the current price of the stock and $x_{i+1}(w)$ its price after T if scenario w occurs, we introduce the following arguments:

- an *uptrend argument* noted $T_\uparrow$
  - if the stock is oversold (the value of the RSI has dropped below 30 and still remains below 50)
  - with supported event $X_{T_\uparrow} = \{w \in \Omega \,|\, x_{i+1}(w) \geq x_i\}$
- a *downtrend argument* noted $T_\downarrow$
  - if the stock is overbought (the value of the RSI has jumped above 70 and still remains above 50)
  - with supported event $X_{T_\downarrow} = \{w \in \Omega \,|\, x_{i+1}(w) \leq x_i\}$
- a *positive divergence argument* noted $D_\uparrow$
  - if the stock and the values of MACD have been going down and MACD only has started to go up
  - with supported event $X_{D_\uparrow} = \{w \in \Omega \,|\, x_{i+1}(w) > x_i\}$
  - attacking the downtrend argument $T_\downarrow$
- a *negative divergence argument* noted $D_\downarrow$
  - if the stock and the values of MACD have been going up and MACD only has started to go down

  - with supported event $X_{D_\downarrow} = \{w \in \Omega \,|\, x_{i+1}(w) < x_i\}$
  - attacking the uptrend argument $T_\uparrow$
- a *lack of confirmation argument* noted $\neg D_\uparrow$ (respectively $\neg D_\downarrow$)
  - if there is a positive (respectively negative) divergence argument
  - attacking the divergence argument $D_\uparrow$ (respectively $D_\downarrow$)
- a *bullish moving average crossover argument* noted MA-$X_\uparrow$
  - if MACD has crossed its 9-day exponential moving average from below (bullish EMA-crossover) and there has been a uninterrupted flow of positive divergence arguments since then
  - attacking the lack of confirmation argument $\neg D_\uparrow$
- a *bullish centerline crossover argument* noted C-$X_\uparrow$
  - if MACD has jumped from negative to positive values (bullish centerline crossover) and there has been a uninterrupted flow of positive divergence arguments since then
  - attacking the lack of confirmation argument $\neg D_\uparrow$
- a *bearish moving average crossover argument* noted MA-$X_\downarrow$
  - if MACD has crossed its 9-day exponential moving average from above (bearish EMA-crossover) and there has been a uninterrupted flow of negative divergence arguments since then
  - attacking the lack of confirmation argument $\neg D_\downarrow$
- a *bearish centerline crossover argument* noted C-$X_\downarrow$
  - if MACD has jumped from positive to negative values (bearish centerline crossover) and there has been a uninterrupted flow of negative divergence arguments since then

– attacking the lack of confirmation argument $\neg D_{\downarrow}$

The possible attacks between arguments can be visualised in Fig. 6.2. Observe that divergence arguments attack trend arguments and that divergence arguments are always attacked by a lack of confirmation arguments (this is a systematic principle of cautiousness). Confirmations can be provided either by centerline or moving average crossover arguments. We believe that such an abstract argumentation framework constitutes a possible and honest reconstitution of the reasoning process used by technical analysts and traders on a daily basis.

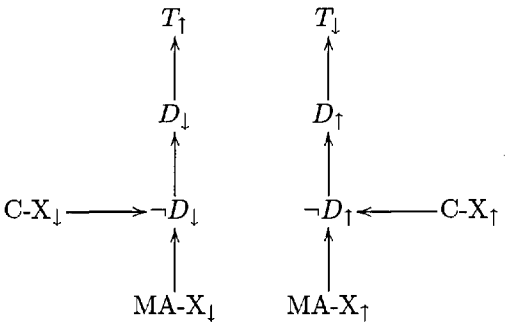


Figure 6.2: Attacks in the stock trading abstract argumentation framework.

At any given time, the argumentation framework may only be in one of the 11 configurations listed in Table 6.1. These configurations differ essentially by the presence of trend and divergence arguments and the number of confirmations of the divergence argument. We are now going to assess, for each one of these configurations, the strength of each argument according to our measure of strength given in Chapter 4. Intuitively, in absence of divergence arguments the trend argument should have maximal strength. A divergence argument should decrease the strength of the trend argument. The impact of the divergence arguments should however diminish with the number of divergence confirmation arguments (moving average and centerline crossover arguments). The strength

values obtained for every argument in the various possible configurations of our financial framework are computed in Table 6.2. As expected, trend arguments have a maximal strength of 1 in the absence of divergence arguments. When a divergence exists, but has no confirmation, the strength of the trend argument collapses to 0.5 and the divergence argument has an even lower strength of 0.25. However, a first confirmation argument restores the strength of the divergence argument to 0.5, and a second to brings it to 0.583. With this first and second confirmation of divergence, the strength of the trend argument further plunges to 0.386 and 0.348. Notice that in the absence of trend argument, the strength of the divergence argument is unchanged. This is quite natural as the trend argument does not constitute an attack against the divergence argument.

Set of arguments Arg	Set of attacks att
$\emptyset$	$\emptyset$
$\{T_\uparrow\}$	$\emptyset$
$\{T_\downarrow\}$	$\emptyset$
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow\}$	$\{(D_\downarrow, T_\uparrow), (\neg D_\downarrow, D_\downarrow)\}$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow\}$	$\{(D_\uparrow, T_\downarrow), (\neg D_\uparrow, D_\uparrow)\}$
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow, \text{MA-X}_\downarrow\}$	$\{(D_\downarrow, T_\uparrow), (\neg D_\downarrow, D_\downarrow), (\text{MA-X}_\downarrow, \neg D_\downarrow)\}$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow, \text{MA-X}_\uparrow\}$	$\{(D_\uparrow, T_\downarrow), (\neg D_\uparrow, D_\uparrow), (\text{MA-X}_\uparrow, \neg D_\uparrow)\}$
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow, \text{C-X}_\downarrow\}$	$\{(D_\downarrow, T_\uparrow), (\neg D_\downarrow, D_\downarrow), (\text{C-X}_\downarrow, \neg D_\downarrow)\}$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow, \text{C-X}_\uparrow\}$	$\{(D_\uparrow, T_\downarrow), (\neg D_\uparrow, D_\uparrow), (\text{C-X}_\uparrow, \neg D_\uparrow)\}$
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow, \text{MA-X}_\downarrow, \text{C-X}_\downarrow\}$	$\{(D_\downarrow, T_\uparrow), (\neg D_\downarrow, D_\downarrow), (\text{MA-X}_\downarrow, \neg D_\downarrow), (\text{C-X}_\downarrow, \neg D_\downarrow)\}$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow, \text{MA-X}_\uparrow, \text{C-X}_\uparrow\}$	$\{(D_\uparrow, T_\downarrow), (\neg D_\uparrow, D_\uparrow), (\text{MA-X}_\uparrow, \neg D_\uparrow), (\text{C-X}_\uparrow, \neg D_\uparrow)\}$

Table 6.1: Possible configurations of the abstract argumentation framework.

6.8.3 Decision making

A simple strategy for a Bayesian investor (using either the pignistic, cautious or proportional probability distributions) would consist at each t_i to buy a quantity q_i of the stock that maximises the expectation of his future wealth $E(W_{i+1}(q_i))$ and, more generally, for an investor using an imprecise probability model, to choose q_i so that his future wealth $X(q_i) = W_{i+1}$ is a maximal gamble. Let us

Arguments Arg	Strength of arguments
$\emptyset$	N.A.
$\{T_\uparrow\}$	1.0
$\{T_\downarrow\}$	1.0
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow\}$	$0.5, 0.25, 1.0$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow\}$	$0.5, 0.25, 1.0$
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow, \text{MA-X}_\downarrow\}$	$0.386, 0.5, 0.25, 1.0$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow, \text{MA-X}_\uparrow\}$	$0.386, 0.5, 0.25, 1.0$
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow, \text{C-X}_\downarrow\}$	$0.386, 0.5, 0.25, 1.0$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow, \text{C-X}_\uparrow\}$	$0.386, 0.5, 0.25, 1.0$
$\{T_\uparrow, D_\downarrow, \neg D_\downarrow, \text{MA-X}_\downarrow, \text{C-X}_\downarrow\}$	$0.348, 0.583, 0.167, 1.0, 1.0$
$\{T_\downarrow, D_\uparrow, \neg D_\uparrow, \text{MA-X}_\uparrow, \text{C-X}_\uparrow\}$	$0.348, 0.583, 0.167, 1.0, 1.0$

Table 6.2: Strength of arguments depending on the framework configuration.

then first identify the set of gambles $\mathcal{X}$. First, budget constraints exist on the decision variable q_i. Since we assume that the investor cannot borrow any cash, the future cash value C_{i+1} must remain always positive, therefore $q_i \in \mathbb{R}$ must satisfy the following constraints

$$\begin{cases} C_i - q_i\, x_i\, (TC/100) - q_i x_i \geq 0 \\ C_i + q_i\, x_i\, (TC/100) - q_i x_i \geq 0 \end{cases}$$

Secondly, the amount of stock to be sold is limited by the amount already owned or equivalently, the quantity Q_{i+1} must always remain positive

$$Q_i + q_i \geq 0$$

The decision variable q_i is consequently bounded up and below. The gamble $X(q_i)$ associated to the decision q_i is defined $\forall w \in \Omega$ as

$$X(q_i)(w) = C_i - |q_i|\, x_i\, (TC/100) - q_i x_i + (Q_i + q_i) x_{i+1}(w)$$

The set $\mathcal{X} = \{X(q_i) \,|\, q_i \in \mathbb{R}, C_{i+1}(q_i) \geq 0, Q_{i+1}(q_i) \geq 0\}$ is non-empty, since it contains notably $X(0)$. $\mathcal{X}$ is clearly compact, as it is the image by continuous

function into the space of gambles of the variable q_i which ranges over closed interval of possible values. A maximal gamble can be found by application of Algorithm 4 from Section 6.4. This algorithm can be slightly specialised to the stock trading problem as shown in Algorithm 5.

Algorithm 5 Finds a maximal gamble X^*.

1: **Initialise**: set $k = 0$ and $\mathcal{Y}^0 = \{X(0)\}$ where the gamble $X(0)$ is defined $\forall w \in \Omega$ as $X(0)(w) = C_i + Q_i x_{i+1}(w)$.
2: **Compute**: $x_k = (X_k, m_k)$ which maximises $m_k \in \mathbb{R}$ subject to $X_k \in \mathcal{X}$ and $\forall Y \in \mathcal{Y}^k$, $m_k \leq \overline{E}(X_k - Y)$ using Algorithm 6.
3: **Compute**: $Y_{k+1} \in \mathcal{X}$ which minimises $\overline{E}(X_k - Y_{k+1})$ using Algorithm 7.
4: **Check**: if $m_k \leq \overline{E}(X_k - Y_{k+1})$ then go to step 5, else update $\mathcal{Y}_{k+1} = \mathcal{Y}_k \cup \{Y_{k+1}\}$ and return to step 2 with $k = k + 1$.
5: **Initialise**: $X = x_k$.
6: **Compute**: $X^* \in \mathcal{X}$ which maximises $\sum_{i=1}^{|\Omega|} X_i^*$ subject to $X^* \geq X$.

Property 6.10. *Algorithm 5 returns a maximal gamble.*

Proof. This directly derives from the correctness of Algorithm 4 and the fact that $X(0) \in \mathcal{X}$. Indeed, $X(0)$ is by definition the gamble representing the future wealth of the investor if he does nothing at time i. $X(0)$ belongs to $\mathcal{X}$ since $q_i = 0 \in \mathbb{R}$ and when $q_i = 0$, we have $C_{i+1} = C_i \geq 0$ (the investor cannot borrow any money) and $Q_{i+1} = Q_i \geq 0$ (the investor cannot short-sell the stock). □

Property 6.11. *Algorithm 6 gives an optimisation program with linear objective function and linear and quadratic constraints that is equivalent to maximising $m_k \in \mathbb{R}$ subject to the constraints $X_k \in \mathcal{X}$ and $\forall Y \in \mathcal{Y}^k$, $m_k \leq \overline{E}(X_k - Y)$ under the dialectical probability.*

Proof. The first two inequality constraints are equivalent to $C_{i+1} \geq 0$. The third inequality constraint is equivalent to $Q_{i+1} \geq 0$. The fourth and fifth inequality constraints are together equivalent to $z_i \geq |q_i|$. Since the objective m to be maximised decreases as z_i becomes larger (sixth and last inequality

Algorithm 6 Maximise $m_k \in \mathbb{R}$ s.t. $X_k \in \mathcal{X}$ and $\forall Y \in \mathcal{Y}^k$, $m_k \leq \overline{E}(X_k - Y)$.

1: **Variables:** $q_i, z_i, t_1, \ldots, t_{|\Omega|}, X_k(w_1), \ldots, X_k(w_{|\Omega|}), P_1, \ldots, P_{|\Omega|}, m \in \mathbb{R}$
2: **Maximise:** m
3: **Subject to the linear and quadratic constraints:**

$$\begin{cases} C_i - q_i\, x_i\, (TC/100) - q_i x_i \geq 0 \\ C_i + q_i\, x_i\, (TC/100) - q_i x_i \geq 0 \\ Q_i + q_i \geq 0 \\ z_i \geq q_i \\ z_i \geq -q_i \\ \forall j \in \{1, \ldots, |\Omega|\} : X_k(w_j) = C_i - z_i\, x_i\, (TC/100) - q_i x_i + (Q_i + q_i) x_{i+1}(w_j) \\ \forall j \in \{1, \ldots, |\Omega|\} : t_j \geq 0 \\ t_1 + \ldots + t_{|\Omega|} = 1 \\ \forall j \in \{1, \ldots, |\Omega|\} : P_j = (1-\epsilon)\hat{p}_A(w_j) + \epsilon t_j \\ \forall Y \in \mathcal{Y}^k : m \leq E_P(X_k - Y) \end{cases}$$

constraints), the variable z_i plays exactly the role of $|q_i|$. Thus, the first six constraints are equivalent to $X_k \in \mathcal{X}$.

Under the dialectical probability, the constraint $P \in \mathcal{D}_A$ is equivalent to the linear constraints $\forall j \in \{1, \ldots, |\Omega|\}$: $P_j = (1-\epsilon)\hat{p}_A(w_j) + \epsilon t_j$, $t_j \geq 0$ and $\sum_{j=1}^{|\Omega|} t_j = 1$. So, the seventh, eighth and ninth constraints are equivalent to $P \in \mathcal{D}_A$. For given gambles X_k and Y, maximising m subject to $m \leq E_P(X_k - Y)$ and $P \in \mathcal{D}_A$ is equivalent to maximising m subject to $m \leq \overline{E}(X_k - Y)$.

The objective m is linear. All constraints are clearly linear except the last inequality constraints $m \leq E_P(X_k - Y)$ which is quadratic, as it involves factors that multiply pairs of variables of the program. □

Property 6.12. *Algorithm 7 gives a linear optimisation program that is equivalent to minimising $\overline{E}(X_k - Y_{k+1})$ subject to $Y_{k+1} \in \mathcal{X}$ under the dialectical probability.*

Proof. This directly derives from the expression of $\overline{E}$ for the dialectical probability given in property 5.13 and the fact that all constraints after the second

Algorithm 7 Minimise $\overline{E}(X_k - Y_{k+1})$ subject to $Y_{k+1} \in \mathcal{X}$.

1: **Variables:** $m, q'_i, z'_i, Y_{k+1}(w_1), \ldots, Y_{k+1}(w_{|\Omega|}) \in \mathbb{R}$
2: **Minimise:** m
3: **Subject to the linear constraints:**

$$\begin{cases} \forall w \in \Omega : m \geq \left[IE_{\hat{p}_A}(X_k - Y_{k+1}) + K(X_k(w) - Y_{k+1}(w))\right]/(K+I) \\ C_i - q'_i x_i \, (TC/100) - q'_i x_i \geq 0 \\ C_i + q'_i x_i \, (TC/100) - q'_i x_i \geq 0 \\ Q_i + q'_i \geq 0 \\ z'_i \geq q'_i \\ z'_i \geq -q'_i \\ \forall j \in \{1, \ldots, |\Omega|\} : Y_{k+1}(w_j) = C_i - z'_i x_i \, (TC/100) - q'_i x_i + (Q_i + q'_i) x_{i+1}(w_j) \end{cases}$$

one are equivalent to $Y_{k+1} \in \mathcal{X}$ (as in the proof of correctness of the previous algorithm). □

6.8.4 Choice of the parameters K and V_A

The statistical confidence parameter x used for calculating the Wilson score intervals, ϵ and eventually K can be chosen arbitrarily. In finance, a high statistical confidence implies more cautiousness and less risk taking. If no risk is taken, the investor's return is the risk-free rate and the Sharpe ratio is null. If the investor takes too much risk, the volatility of his investment's returns explodes and leads inevitably to low Sharpe ratio values as well. Statistics do not offer any method for choosing the ideal value of x in our problem. We have shown in Chapter 5 that robust belief functions are equivalent to maximal contaminations of the statistical prior $\hat{p}$ with level of contamination ϵ. Thus, ϵ represents the frequency with which the random process observed does not follow the distribution $\hat{p}$, but some other unknown probability q. We have also shown that dialectical probabilities collapse to robust beliefs whenever the set of forecast arguments is empty ($A = \emptyset$), see property 5.9. Therefore, ϵ should be chosen so as to match with the frequency $Fr(A \neq \emptyset)$ with which the set

of forecast arguments is not empty. This frequency can be measured from the benchmark B. We thus choose $\epsilon = Fr(A \neq \emptyset)$ and $K = \epsilon/(1-\epsilon)$, or equivalently

$$K = \frac{Fr(A \neq \emptyset)}{1 - Fr(A \neq \emptyset)}$$

We have also seen (see property 5.17 in particular) that when the set of forecast arguments is non-empty ($A \neq \emptyset$), the level of contamination ϵ_A is lower than ϵ. The value ϵ_A is interpreted as the frequency with which the random process does not follow the distribution of the biased prior $\hat{p}_A$. Assuming that the forecast arguments are good enough, we can expect the frequency failure ϵ_A to be, let us say, half of ϵ. Under this assumption, the average informational value I_A of the arguments must be such that $\epsilon/2 = K/(K + I_A + 1)$, or, after simplification, I_A is in average equal to $K + 1$. Denoting now by $\langle s \rangle$ the average value (in time) of the total strength of the forecast arguments (also be estimated from the benchmark B), we choose

$$V_A = \frac{K+1}{\langle s \rangle}$$

6.8.5 Performance results

We may now compare the investment performance obtained with the American Express stock (AXP), using maximal gamble strategies and, as probabilistic model, either

- the statistical prior $\hat{p}$, calculated from an historical sample of daily returns over one year and updated monthly,
- the robust belief functions based on $\hat{p}$, using a rolling window of 10 years to estimate the epistemic risk aversion parameter K, or
- the dialectical probabilities based on $\hat{p}$, using the same value of K and

statistics	stock	stat. prior	robust belief	dialectical prob.
average yearly return	4.08%	0.90%	7.51%	16.17%
worst daily log-return	-14.60%	-14.60%	-11.50%	-11.50%
best daily log-return	12.01%	10.42%	12.01%	12.01%
yearly volatility	35.63%	32.01%	21.13%	24.38%
Sharpe ratio	-0.02	-0.12	0.11	0.41
yearly skewness	-0.43%	-0.85%	1.71%	2.59%

Table 6.3: Performance of the statistical prior, robust beliefs and dialectical probabilities compared to the stock's performance.

same 10 years window to estimate V_A.

The transaction costs TC are taken equal to 0.25% and the yearly risk-free rate assumed is $R_F = 5\%$. The evolutions of the gross wealth, starting from an initial budget of 1 are shown in Figure 6.3. The bottom curve corresponds to trading with the statistical prior, the middle curve to the robust belief and the top curve to the dialectical probability.

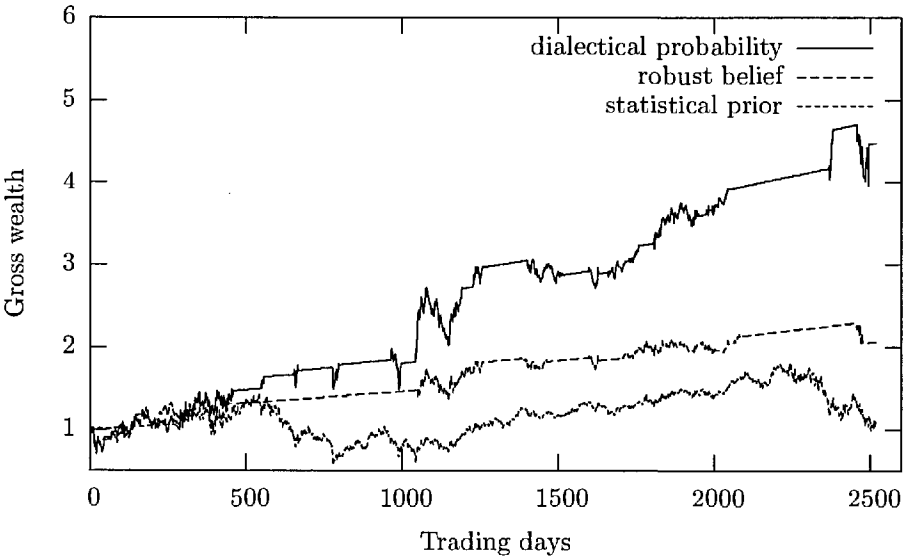


Figure 6.3: Evolution of the gross wealth depending on the probability used.

We now analyse quantitatively and qualitatively the results obtained (see Table 6.3). The difference between the middle and bottom curve shows the added value of robust belief functions and justifies the importance of imprecise probabilities. The average yearly returns with the statistical prior are of 0.90%, and with the robust beliefs are of 7.51%. The worst daily log-return is improved from -14.60% to -11.50%, the best daily log-return from 10.42% to 12.01%. The volatility (risk) is also reduced from 32.01% to 21.13%. The overall performance (Sharpe ratio) is increased from the negative value (below the risk-free rate) -0.12 to 0.11, which means we outperform the risk-free rate. The robust beliefs make the log-returns distribution positively skewed, which means that the downside risk (part of the risk that really matters) is reduced. We observe here that robust belief functions allow to build much safer and more efficient investment strategies than statistical priors, by greatly reducing the volatility and increasing the return from a level below the risk-free rate with the statistical prior to a level slightly above.

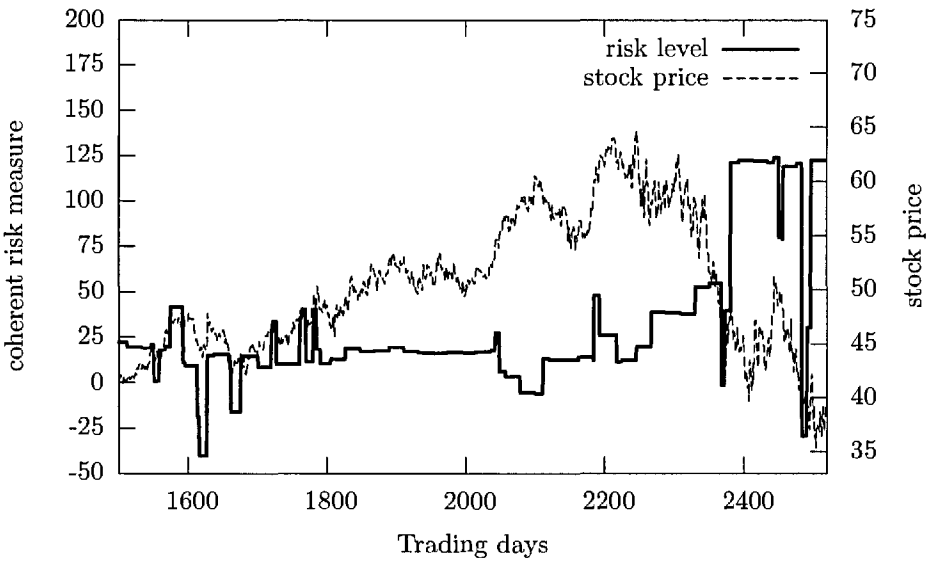


Figure 6.4: Risk of holding the stock per unit of money invested.

The difference between the top and middle curve shows the added value of argumentation (and more generally the potential of using subjective information) over robust belief functions. The average yearly returns with the dialectical prior are further increased from 7.51% to 16.17%. The worst and best daily returns remain the same. The volatility is slightly increased from 21.13% to 24.38%, although the skewness is in fact improved. The overall performance gain is quite important, with a Sharpe ratio jumping from 0.11 to 0.41. This performance is high despite the credit crunch period (last 250 trading days) where the coherent risk measure does not fail to detect an abnormally high level of risk (observe the jump made on Figure 6.4 after roughly 2400 trading days).

In conclusion, it appears that dialectical probabilities constitute a good basis for the construction of efficient and dynamic investment strategies, with quasi equal levels of risk and much higher returns. This advantages is provided by carefully chosen, weighted and combined arguments concerning price trends or divergences generated fully automatically and drawn from expert knowledge in technical analysis.

6.9 Related work

Imprecise probability is a rather recent topic and its use for statistical inference, decision and risk analysis has been investigated so far essentially by only a few researchers in statistics, artificial intelligence, finance and operations research.

In statistics, imprecise probabilities are used to obtain robust intervals as well as maximum likelihood estimations for probabilities, conditional probabilities, expectations, variances, entropy or mutual information [Berger and Berliner, 1986; Hutter, 2003]. Notably, the ϵ-contamination models used by robust Bayes analysts [Huber, 1981; Berger and Berliner, 1986], but also the imprecise Dirichlet model introduced by [Walley, 1991] relax the classical Bayesian assumption

that a prior state of ignorance can be represented by a single prior probability distribution. Like these authors, we have provided in Chapter 5 robust intervals for probabilities, conditional probabilities, expectations and entropy. A robust interval for variance is given in the appendix (the proof of this technical result obtained by Lagrangian duality is omitted) for dialectical probabilities. Interestingly, the imprecise Dirichlet model uses an hyper-parameter noted s that allows the subject to control the speed at which probabilistic imprecision decreases with the amount of information. The epistemic risk-aversion parameter K introduced in our approach plays essentially the same role.

In the area of artificial intelligence, credal sets have been found appropriate for representing uncertainty in situations where the information is incomplete [Zaffalon, 2002a], imprecise [Walley and Fine, 1979; Zaffalon, 2002b; Kozine and Utkin, 2005; Utkin, 2006], vague [Zadeh, 1978; de Cooman, 1997; Walley and de Cooman, 2001] or contradictory [Walley, 1996]. Imprecise probabilities are thus appreciated as a model of partial ignorance, incompleteness, indeterminacy and indecision [Zaffalon, 2005]. We have discussed these aspects in the case of dialectical probabilities in Chapters 5 and 6 and, in this chapter, we have shown in detail the connection between decision making under total probabilistic ignorance (vacuous previsions) and strict uncertainty. We have also observed empirically that our trading agent was less frequently undecided when taking into account expert opinions (the wealth grows more often at the risk-free rate under robust belief functions than it does under dialectical probabilities).

In finance, authors such as [Knight, 1921] and more recently [Kane, 1999] have argued against the existence of a common observable stationary probability measure for security prices. Indeed, the periodic returns of a firm's stock do not really fulfil the definition of a random experiments in statistics (i.e. an experiment that can be repeated many times in exactly the same conditions) as e.g.

firms constantly introduce new products and services. Also, [Cragg and Malkiel, 1968] observed that different security analysis firms' estimates differed greatly. The very existence of active and volatile securities markets challenges the assumption of unanimous agreement as to what parties feel as security is worth. Kane ascertained that "imprecision gives active and volatile markets a reason to exist". These observations justify the importance of imprecise probabilities in finance. Notice that for instance, coherent financial risk measures [Artzner *et al.*, 1999] are equivalent to coherent upper previsions [Pelessoni and Vicig, 2001; Vicig, 2008]. In this chapter, we have proposed such a measure of financial risk based on dialectical probabilities and demonstrated experimentally its ability to detect risky situations (credit crunch).

In operations research, [Ben-Tal and Nemirovski, 2000] showed that even small perturbations of uncertain quantities could result in highly infeasible solutions of linear optimisation problems. For these problems, a robust formulation was proposed in terms of a matrix of uncertain coefficients ranging within an uncertainty set. [Bertsimas and Brown, 2008], who observed that the literature on robust optimisation was "essentially silent on the question of constructing these uncertainty sets" proposed a systematic method for constructing uncertainty sets from coherent risk measures [Artzner *et al.*, 1999]. They produced then a number of uncertainty sets for robust linear optimisation from several important coherent risk measures, such as e.g. CVaR and one-sided moment measures. Naturally, the same technique may be used to produce uncertainty sets in robust linear optimisation from our coherent risk measure ρ based on dialectical probabilities.

[Treynor and Ferguson, 1985] have explained theoretically how investors possessing private and non-price information could statistically achieve unusual profits by using technical analysis. In particular, [Osler, 2003] provided ex-

perimental evidence and an explanation for the predictive success of technical analysis in the foreign exchange market. Our experiments in this chapter have confirmed that unusual profits could be achieved by using technical analysis. [Lo *et al.*, 2000] introduced kernel regression as a mathematical method for removing the noise component of stock price time series, thus simplifying the problem of automatic pattern detection in technical analysis. In our experiments, pattern detection was based on argumentation and thus was done quite differently. [Gately, 1996] and [Refenes *et al.*, 1997] have developed automatic trading systems using technical analysis based on neural networks and [Simutis, 2000] and [Dourra and Siy, 2002] used fuzzy logic. In this chapter, we have explained how to design an automatic trading system using technical analysis based instead on imprecise probabilities.

6.10 Summary and conclusion

In this chapter, we have seen that the maximality criterion generalised the expected utility criterion to the case of imprecise probabilities and have studied the decision-theoretic properties fulfilled by this criterion. We have observed that the maximality criterion had the same properties as Savage's regret criterion, but that the maximality criterion was less pessimistic than the regret criterion. We have also introduced a general algorithm based on semi-infinite programming for finding maximal gambles under any imprecise probability.

We have applied existing results from the theory of imprecise probabilities and thus showed how to solve problems of statistical inference, financial risk analysis and decision analysis when using dialectical probabilities in particular. The efficiency of dialectical probabilities combined with the maximality criterion have been experimentally validated in a financial setting. We have demonstrated experimentally the potential of expert arguments drawn from technical analysis

and their ability to improve the performance of investment strategies based on purely statistical estimations of probability.

Chapter 7

Conclusion

In this thesis, we have proposed a methodology for constructing decision making models in decision theory. In Chapter 2, we have introduced decision theory and described the three possible components of a decision model, namely, a decision table, a value function and a Bayesian probability distribution. Then, the construction of each one of these components has been considered as a problem on its own and has been examined in detail in a separate chapter of the thesis.

Chapter 3 has been dedicated to the construction of decision tables. By definition, a decision table is the graphical representation of a mapping $T : D \times \Omega \rightarrow X$, where D is a set of possible decisions, Ω an exhaustive set of mutually exclusive scenarios and X a set of outcomes. A decision table thus represents all the possible outcomes of decisions in a given situation. First, basic influence diagrams have been introduced as a tool for describing logical and uncertain knowledge about a given decision domain. We have used assumption-based argumentation to analyse and process basic influence diagrams. We have proved that the new liberal stable semantics could be used to process basic

influence diagrams in a rational way and derive decision tables. We have studied the suitability of hypergraph algorithms to compute the liberal stable semantics and have applied this technique to build a decision table for the game of poker.

Chapter 4 has been dedicated to the construction of value functions. By definition, a value function is a function $V : X \to \mathbb{R}$ that quantifies the desirability of each outcome in X. We have used abstract argumentation to model the uncertain knowledge of a decision maker regarding the desirability of his objectives. In order to weigh the uncertainties, we have developed a quantitative measure of argument strength. This measure allows to estimate the strength of arguments in dialectic settings, whereby each argument is a justified claim on its own and the arguments may undermine or "attack" each other. The strength of an argument is defined as the value (the guaranteed expected gain) of a proponent of the argument confronting an opponent of the argument via a repeated game of argumentation strategy. We have used a simple mechanism of preference aggregation from cardinal welfare theory to calculate a value function V, by combining the decision maker's arguments concerning the desirability of his objectives and the precise strength of these arguments. We have studied the main properties of the value functions obtained and have applied the proposed technique for obtaining value functions to build value functions for poker players with various psychological and emotional patterns.

Chapter 5 has been dedicated to the construction of imprecise probabilities. Imprecise probabilities generalise the Bayesian model of probability. Probabilities serve to quantify the chance of occurrence of each possible scenario in Ω. We have developed a framework for constructing imprecise probabilities called dialectical probabilities taking the form of belief functions, given as input some statistical data plus a number of expert opinions concerning the expected occurrence of scenarios in Ω. Opinions have been modeled by sets of abstract

arguments called forecast and mitigation arguments. Forecast arguments are justified claims in favour of the occurrence of events (thus supporting one or more scenarios) and mitigation arguments as justified claims undermining forecast or mitigation arguments. We have proved that dialectical probabilities are mathematically coherent and can be interpreted as coherent previsions. We have also established that they could be interpreted in the same way as in robust Bayes analysis. This has allowed us to study their ability to model probabilistic indeterminacy. Dialectical probabilities have been applied with success to trust computing.

Imprecise probabilities are not Bayesian and thus cannot be used directly to make decisions in decision theory. Therefore, we have dedicated Chapter 6 to explaining how dialectical probabilities could be used in practise to solve decision making problems, as well as problems in decision analysis, statistical inference and financial risk analysis. Then, we have studied in detail the main criterion for making decisions with imprecise probabilities (maximality criterion) and have provided a new algorithm based on semi-infinite programming for computing optimal decisions. We have also studied the exact relationships between the classical decision-theoretic criteria and the criteria used in imprecise probability theory. Argumentation, dialectical probabilities and the algorithm used to compute optimal decisions have been applied and validated within a financial setting, whereby we could demonstrate that expert abstract arguments drawn from technical analysis can significantly improve the performance of statistical stock trading strategies.

In summary, the main contributions of this thesis are

- a practical method for constructing decision models in decision theory,
- applications of that method to poker, trust computing and finance,
- a family of diagrams for the logical representation of situations of decision

making under strict uncertainty,

- a semantics for assumption-based argumentation and an algorithm based on hypergraphs to compute this semantics,
- a game-theoretic approach to the computation of the strength of arguments in abstract argumentation,
- formulae for building value functions from uncertain judgements of desirability,
- formulae for building belief and plausibility functions from statistical data and expert opinions,
- a general algorithm for computing maximal decisions under any imprecise probability,
- an analysis of the relationships between the decision criteria used in decision theory and those used in the theory of coherent previsions, and
- formulae to solve problems of statistical inference, risk and decision analysis when using statistical data and expert opinions.

The results in this thesis provide a foundation for the future design of expert systems. Two important challenges remain however to be addressed for the successful deployment of such systems. The first task would be to endow these systems with the capability to acquire, enhance and adapt their expertise with a high degree of autonomy. Achieving this would probably involve the use of advanced and maybe even novel techniques combining elements of machine learning and argumentation. The second issue would be to make these complex systems more understandable, influenceable and persuasive from the perspective of the people exploiting them. Argumentation is doubtlessly the key here.

Appendix

Theorem 5.3 *Let x be a parameter comprised between 0 and 1. Let $\underline{p}(w)$ and $\overline{p}(w)$ be the roots of the quadratic equation in the variable p*

$$\left(1+\frac{z^2}{N}\right)p^2-\left(2\hat{p}(w)+\frac{z^2}{N}\right)p+\hat{p}(w)^2=0$$

where we assume that $\underline{p}(w) \leq \overline{p}(w)$ and z corresponds to the $1-x/2$ percentile of the standard normal distribution $\mathcal{N}(0,1)$. Then, $\mathcal{I}(w) = \left[\underline{p}(w), \overline{p}(w)\right]$ is an approximate $(1-x)100\%$ confidence interval for the true probability $p(w)$.

Proof. Consider some fixed scenario $w \in \Omega$ and the random variable $X : \Omega \rightarrow \mathbb{R}$ defined $\forall w_i \in \Omega$ as $X(w_i) = 1$ if $w_i = w$ and 0 otherwise. Let $X_1, \ldots, X_N \in \{0,1\}$ be a random sample of N observations of X and $\overline{X} = \frac{1}{N}\sum_{i=1}^{N} X_i$ be the sample mean. If $p(w)$ is the true probability of w, then X has a unit binomial distribution with success $p(w)$. The mean and variance of this distribution exist and are respectively equal to $\mu = p(w)$ and $\sigma^2 = p(w)(1-p(w))$. If N is large, then by the central limit theorem, the random variable

$$Z = \sqrt{N}\left[\frac{\overline{X}-\mu}{\sigma}\right]$$

has an approximate standard normal distribution. Thus, $Pr(-z < Z < z) = 1-x$, or equivalently $Pr(Z^2 < z^2) = 1-x$, which yields by definition of the random variable Z

$$Pr\left(N\frac{(\overline{X}-p(w))^2}{p(w)(1-p(w))} < z^2\right) = 1-x$$

and after simplification

$$Pr\left(\left(1+\frac{z^2}{N}\right)p^2-\left(2\overline{X}+\frac{z^2}{N}\right)p+\overline{X}^2 < 0\right) = 1-x$$

hence the solutions of the above-mentioned quadratic equation. □

Property 5.1 *Every Choquet capacity of order $k > 2$ is also a Choquet capacity of order $k' = 2, 3, \ldots, k$.*

Proof. It is sufficient to prove that for any integer $k > 2$, if μ is a Choquet capacity of order k then it is also a Choquet capacity of order $k-1$. Let then μ be a Choquet capacity of order k. Let also $A_1, ..., A_{k-1} \subseteq \Omega$ be a family of $k-1$ events. We may introduce $A_k = \emptyset$. We have by hypothesis on μ

$$\mu\left(\bigcup_{j=1}^{k-1} A_j\right) = \mu\left(\bigcup_{j=1}^{k} A_j\right) \geq \sum_{\substack{K \subseteq \{1,\ldots,k\} \\ K \neq \emptyset}} (-1)^{|K|+1} \mu\left(\bigcap_{j \in K} A_j\right)$$

If $k \in K$, then $\bigcap_{j \in K} A_j = \emptyset$, and since μ is a monotone measure, it satisfies in particular $\mu(\emptyset) = 0$. Therefore, we have in fact

$$\mu\left(\bigcup_{j=1}^{k-1} A_j\right) \geq \sum_{\substack{K \subseteq \{1,\ldots,k-1\} \\ K \neq \emptyset}} (-1)^{|K|+1} \mu\left(\bigcap_{j \in K} A_j\right)$$

and thus μ is also a Choquet capacity of order $k-1$. □

Property 5.2 *Every Choquet capacity of order 2 is a lower probability.*

Proof. Let μ be a Choquet capacity of order 2 and $\mathcal{D}$ the set of all probability distributions that dominate μ

$$\mathcal{D} = \{p \,|\, \forall E \subseteq \Omega, \mu(E) \leq p(E)\}$$

Let us introduce the function $q : 2^\Omega \to [0, 1]$ defined for every event E as

$$q(E) = \sum_{w \in E} \mu(\{w\})$$

Since μ is a Choquet capacity of order 2, we have for every event E

$$\mu(E) \leq \mu(\Omega) - \mu(\Omega - E) \leq 1 - \sum_{w \in E} \mu(\{w\}) = 1 - q(E)$$

Let us now define $p : \Omega \to \mathbb{R}$ for every scenario w as

$$p(w) = 1 - q(\{w\}) + \frac{|\Omega| - q(\Omega)}{|\Omega|}$$

We have $q(\{w\}) = \mu(\{w\}) \leq \mu(\Omega) = 1$ and $q(\Omega) \leq \mu(\Omega) = 1$ by monotonicity of μ, hence p is positive. Moreover, it is easy to check that p is normalised and thus defines a probability distribution. Finally, p dominates $1 - q$, which in turn dominates μ. Finally, $p \in \mathcal{D}$ and $\mathcal{D}$ is guaranteed to be a non-empty set. Let then $\underline{P}$ be the lower probability defined by the credal set $\mathcal{D}$. For every event E and every $p \in \mathcal{D}$, we have $\mu(E) \leq p(E)$ and consequently $\mu(E) \leq \underline{P}(E)$. For every event E, $\emptyset \subseteq E \subseteq \Omega$ and since μ is a monotone measure we have $0 \leq \mu(E) \leq 1$. There exists therefore a probability distribution p_E such that $p_E(E) = \mu(E)$. This implies notably that $p_E \in \mathcal{D}$. Therefore, $\underline{P}(E) \leq \mu(E)$. Finally, we have $\underline{P} = \mu$, so μ is indeed a lower probability. □

Property .1. *Given an imprecise probability with credal set $\mathcal{D}$, the lower and upper variances of a gamble X defined by [Walley, 1991] as*

$$\underline{V}(X) = \min_{P \in \mathcal{D}} V_P(X), \quad \overline{V}(X) = \max_{P \in \mathcal{D}} V_P(X)$$

where $V_P(X) = E_P((X - E_P(X))^2)$, are in the case of the dialectical probability equal to

$$\underline{V}(X) = (1 - \epsilon_A)\left(E_{\hat{p}_A}(X^2) - (1 - \epsilon_A)E^2_{\hat{p}_A}(X)\right) + \epsilon_A \min(\gamma(X))$$

$$\overline{V}(X) = -(1-\epsilon_A)\left(E_{\hat{p}_A}(X^2) - (1-\epsilon_A)E^2_{\hat{p}_A}(X)\right) - \epsilon_A \max(\gamma(X))$$

where $\gamma(X) = (1-\epsilon_A)X^2 - 2(1-\epsilon_A)E_{\hat{p}_A}(X)X$.

Proof. This result can be obtained by application of Lagrangian duality, see [Boyd and Vandenberghe, 2004]. The detailed calculus is omitted. □

Bibliography

[Ambler, 1996] S. Ambler. A categorical approach to the semantics of argumentation. *Mathematical Structures in Computer Science*, 6(2):167–188, 1996.

[Amgoud and Cayrol, 1998] L. Amgoud and C. Cayrol. On the Acceptability of Arguments in Preference-based Argumentation. In *14th Conference on Uncertainty in Artificial Intelligence*, 1998.

[Amgoud and Cayrol, 2002] L. Amgoud and C. Cayrol. A reasoning model based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence*, 34:197–215, 2002.

[Amgoud and de Saint-Cyr, 2008] L. Amgoud and F. Dupin de Saint-Cyr. On measuring persuasion dialogs quality. In *2nd International Conference on Computational Models of Argument*, 2008.

[Amgoud and Prade, 2004] L. Amgoud and H. Prade. Using arguments for making decisions: A possibilistic logic approach. In *20th Conference on Uncertainty in Artificial Intelligence*, 2004.

[Amgoud and Prade, 2006] L. Amgoud and H. Prade. Making decisions from weighted arguments. In *Decision theory and multi-agent planning*. 2006.

[Amgoud and Prade, 2008] L. Amgoud and H. Prade. Comparing decisions on the basis of a bipolar typology of arguments. In *Preferences and similarities.* 2008.

[Amgoud, 2005] L. Amgoud. A unified setting for inference and decision: An argumentation-based approach. In *21st Conference on Uncertainty in Artificial Intelligence*, 2005.

[Amgoud, 2009] L. Amgoud. Argumentation in decision making. *Argumentation in Artificial Intelligence*, pages 301–320, 2009.

[Artzner *et al.*, 1999] P. Artzner, F. Delbaen, J.M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.

[Arvey, 1994] R.D. Arvey. Mainstream Science on Intelligence. *Wall Street Journal*, 51, 1994.

[Bacchus and Grove, 1996] F. Bacchus and A.J. Grove. Utility Independence in a Qualitative Decision Theory. In *5th International Conference on Principles of Knowledge Representation and Reasoning*, 1996.

[Ben-Tal and Nemirovski, 2000] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming*, 88:411–424, 2000.

[Ben-Tal and Nemirovski, 2001] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, volume 2. SIAM, 2001.

[Bench-Capon and Dunne, 2002] T.J.M. Bench-Capon and P.E. Dunne. Value based argumentation frameworks. Technical Report ULCS-02-001, 2002.

[Bench-Capon and Dunne, 2007] T.J.M. Bench-Capon and P.E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

[Bench-Capon and Prakken, 2006] T.J.M. Bench-Capon and H. Prakken. Justifying Actions by Accruing Arguments. In *1st International Conference on Computational Models of Argument*, 2006.

[Bench-Capon, 2003] T.J.M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.

[Berger and Berliner, 1986] J.O. Berger and L.M. Berliner. Robust Bayes and empirical Bayes analysis with e-contaminated priors. *The Annals of Statistics*, 14:461–486, 1986.

[Berger, 1985] J.O. Berger. *Statistical Decision Theory and Bayesian Analysis.* Springer, 1985.

[Bertsimas and Brown, 2005] D. Bertsimas and D.B. Brown. Robust linear optimization and coherent risk measures. LIDS Publication #265, 2005.

[Bertsimas and Brown, 2008] D. Bertsimas and D.B. Brown. Constructing uncertainty sets for robust linear optimisation. *Operations Research*, 2008.

[Besnard and Hunter, 2000] P. Besnard and A. Hunter. A Logic-Based Theory of Deductive Arguments. *Artificial Intelligence*, 128(1–2):203–235, 2000.

[Besnard and Hunter, 2008] P. Besnard and A. Hunter. *Elements of Argumentation.* The MIT Press, 2008.

[Billings *et al.*, 1999] D. Billings, L. Peña Castillo, J. Schaeffer, and D. Szafron. Using Probabilistic Knowledge and Simulation to Play Poker. In *16th National Conference on Artificial Intelligence*, 1999.

[Biswas *et al.*, 1988] G. Biswas, M. Oliff, and R. Abramczyk. OASES (operations analysis expert system): an application in fiberglass manufacturing. *International Journal of Expert Systems: Research and Applications*, 1:193–216, 1988.

[Blankenship and Falk, 1976] J.W. Blankenship and J.E. Falk. Infinitely constrained optimization problems. *Journal of Optimization Theory and Applications*, 19:261–281, 1976.

[Blum and Rosenblatt, 1967] J.R. Blum and J. Rosenblatt. On partial a priori information in statistical inference. *Annals of Mathematical Statistics*, 38:1671–1678, 1967.

[Bondarenko *et al.*, 1997] A. Bondarenko, P.M. Dung, R.A. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1–2):63–101, 1997.

[Bonet and Geffner, 1996] B. Bonet and H. Geffner. Arguing for Decisions: A Qualitative Model of Decision Making. In *12th Conference on Uncertainty in Artificial Intelligence*, 1996.

[Borel, 1921] E. Borel. La théorie du jeu et les équations intégrales à noyau symétrique gauche. In *Comptes Rendus de l'Académie des Sciences*, 1921.

[Boros *et al.*, 2000] E. Boros, K.M. Elbassioni, V. Gurvich, and L. Khachiyan. An efficient incremental algorithm for generating all maximal independent sets in hypergraphs of bounded dimension. *Parallel Processing Letters*, 10(4):253–266, 2000.

[Boutilier *et al.*, 1999] C. Boutilier, R.I. Brafman, H.H. Hoos, and D. Poole. Reasoning with conditional ceteris paribus preference statements. In *15th Conference on Uncertainty in Artificial Intelligence*, 1999.

[Boutilier, 1994] C. Boutilier. Toward a logic for qualitative decision theory. In *4th International Conference on Principles of Knowledge Representation and Reasoning*, 1994.

[Boyd and Vandenberghe, 2004] S. Boyd and L. Vandenberghe. *Convex Optimization.* Cambridge University Press, 2004.

[Brachman and Levesque, 2004] R.J. Brachman and H.J. Levesque. *Knowledge Representation and Reasoning.* Morgan Kaufmann Publishers, 2004.

[Caminada and Amgoud, 2005] M. Caminada and L. Amgoud. An axiomatic account of formal argumentation. In *20th National Conference on Artificial Intelligence*, 2005.

[Caminada, 2006] M. Caminada. Semi-stable semantics. In *1st International Conference on Computational Models of Argument*, 2006.

[Caro, 2003] M. Caro. *Caro's Book of Poker Tells.* Cardoza, 2003.

[Cayrol and Lagasquie-Schiex, 2005] C. Cayrol and M.-C. Lagasquie-Schiex. Graduality in argumentation. *Journal of Artificial Intelligence Research*, 23:245–297, 2005.

[Chesley, 1975] G.R. Chesley. Elicitation of subjective probabilities: a review. *The Accounting Review*, 1975.

[Chesñevar *et al.*, 2000] C.I. Chesñevar, A.G. Maguitman, and R.P. Loui. Logical models of argument. *ACM Computing Surveys*, 32(4):337–383, 2000.

[Choquet, 1953] G. Choquet. Theory of capacities. *Annales de l'Institut Fourier*, 5:131–295, 1953.

[Cobb and Shenoy, 2006] B.R. Cobb and P.P. Shenoy. On the plausibility transformation method for translating belief function models to probability models. *International Journal of Approximate Reasoning*, 41:314–330, 2006.

[Cragg and Malkiel, 1968] J.G. Cragg and B.G. Malkiel. The consensus and accuracy of some predictions of the growth of corporate earnings. *Journal of Finance*, 23:67–84, 1968.

[Daniel, 2006] M. Daniel. On transformations of belief functions to probabilities. *International Journal of Intelligent Systems*, 21:261–282, 2006.

[Dantzig *et al.*, 1955] G. B. Dantzig, A. Orden, and P. Wolfe. The generalized simplex method for minimizing a linear form under linear inequality constraints. *Pacific Journal of Mathematics*, 5(2):183–195, 1955.

[de Cooman, 1997] G. de Cooman. Possibility theory – I, II, III. *International Journal of General Systems*, 25:291–371, 1997.

[de Moivre, 1718] A. de Moivre. *The Doctrine of Chances*. H. Woodfall, 1718.

[Dimopoulos *et al.*, 2002] Y. Dimopoulos, B. Nebel, and F. Toni. On the computational complexity of assumption-based argumentation for default reasoning. *Artificial Intelligence*, 141(1–2):57–78, 2002.

[Dourra and Siy, 2002] H. Dourra and P. Siy. Investment using technical analysis and fuzzy logic. *Fuzzy Sets and Systems*, 127(2):221–240, 2002.

[Dubois and Prade, 1988] D. Dubois and H. Prade. An introduction to possibilistic and fuzzy logics. In *Non-Standard Logics for Automated Reasoning*. 1988.

[Dubois and Prade, 1995] D. Dubois and H. Prade. Possibility theory as a basis for qualitative decision theory. In *14th International Joint Conference on Artificial Intelligence*, 1995.

[Dung *et al.*, 2002] P.M. Dung, P. Mancarella, and F. Toni. Argumentation-Based Proof Procedures for Credulous and Sceptical Non-monotonic Reasoning. In *Computational Logic: Logic Programming and Beyond, Essays in Honour of Robert A. Kowalski, Part II*, 2002.

[Dung *et al.*, 2006] P.M. Dung, R.A. Kowalski, and F. Toni. Dialectic proof procedures for assumption-based, admissible argumentation. *Artificial Intelligence*, 170(2):114–159, 2006.

[Dung *et al.*, 2007] P.M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 171(10–15):642–674, 2007.

[Dung, 1995] P.M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. *Artificial Intelligence*, 77(2):321–257, 1995.

[Dunne and Bench-Capon, 2003] P.E. Dunne and T. Bench-Capon. Two Party Immediate Response Disputes: Properties and Efficiency. *Artificial Intelligence*, 149(2):221–230, 2003.

[Dunne, 2007] P.E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171(10–15):701–729, 2007.

[Dunne, 2008] P.E. Dunne. The Computational Complexity of Ideal Semantics I: Abstract Argumentation Frameworks. In *2nd International Conference on Computational Models of Argument*, 2008.

[Dunne, 2009] P.E. Dunne. The computational complexity of ideal semantics. *Artificial Intelligence*, (in press), 2009.

[Dyer and Sarin, 1978] J.S. Dyer and R.K. Sarin. Measurable multiattribute value functions. *Operations Research*, 27:810–822, 1978.

[Eiter and Gottlob, 1995] T. Eiter and G. Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing*, 24:1278–1304, 1995.

[Eiter and Gottlob, 2002] T. Eiter and G. Gottlob. Hypergraph transversal computation and related problems in logic and AI. In *European Conference on Logics in Artificial Intelligence*, 2002.

[Eliashberg, 1980] J. Eliashberg. Consumer Preference Judgements: An Exposition with Empirical Applications. *Management Science*, 26(1):60–77, 1980.

[Farquhar, 1977] P.H. Farquhar. A survey of multiattribute utility theory and applications. *Studies in Management Science*, 6:59–80, 1977.

[Fishburn, 1964] P.C. Fishburn. *Decision and Value Theory.* John Wiley and Sons, 1964.

[Fishburn, 1979] P.C. Fishburn. Two-Piece von Neumann-Morgenstern Utility Functions. *Decision Sciences*, 10:503–518, 1979.

[Fleming, 1952] M. Fleming. A cardinal concept of welfare. *Quarterly Journal of Economics*, 66:366–384, 1952.

[Fox *et al.*, 1993] J. Fox, P. Krause, and M. Elvang-Gøransson. Argumentation as a General Framework for Uncertain Reasoning. In *9th Conference on Uncertainty in Artificial Intelligence*, 1993.

[Fox *et al.*, 1998] J. Fox, N. Johns, and A. Rahmanzadeh. Disseminating medical knowledge: the PROforma approach. *Artificial Intelligence in Medicine*, 14:157–182, 1998.

[Fox, 1993] J. Fox. On the soundness and safety of expert systems. *Artificial Intelligence in Medicine*, 5(2):159–179, 1993.

[Fox, 2003] J. Fox. Probability, logic and the cognitive foundations of rational belief. *Journal of Applied Logic*, 1:197–224, 2003.

[French, 1987] S. French. *Decision theory: an introduction to the mathematics of rationality.* Ellis Horwood, 1987.

[Gately, 1996] E. Gately. *Neural Networks for Financial Forecasting.* Wiley, 1996.

[Goodman *et al.*, 1979] B. Goodman, M. Saltzman, W. Edwards, and D.H. Krantz. Prediction of bids for two-outcome gambles in a casino setting. *Organizational, Behavioural and Human Performance*, 24:382–399, 1979.

[Harsanyi, 1955] J.C. Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63:309–321, 1955.

[Hartley, 1928] R.V.L. Hartley. Transmission of information. *The Bell System Technical Journal*, 7(3):535–563, 1928.

[Hauser and Urban, 1975] J.R. Hauser and G.L. Urban. A normative methodology for modelling consumer response to innovation. *Operations Research*, 25:579–619, 1975.

[Hillier and Lieberman, 1995] F.S. Hillier and G.J. Lieberman. *Introduction to Operations Research - 6th Edition.* McGraw-Hill, 1995.

[Howard and Matheson, 1981] R.A. Howard and J.E. Matheson. Influence diagrams. In *Readings on the Principles and Applications of Decision Analysis.* 1981.

[Hsia and Shenoy, 1989] Y.-T. Hsia and P.P. Shenoy. An evidential language for expert systems. *Methodologies for Intelligent Systems*, 4:9–16, 1989.

[Huber, 1981] P.J. Huber, editor. *Robust Statistics.* Wiley series in probability and mathematical statistics. John Wiley & Sons, 1981.

[Hunter, 2004] A. Hunter. Towards higher impact argumentation. In *19th American National Conference on Artificial Intelligence*, 2004.

[Hutter, 2003] M. Hutter. Robust Estimators under the Imprecise Dirichlet Model. Technical Report IDSIA-03-03, 2003.

[Ibramsha and Rajaraman, 1978] M. Ibramsha and V. Rajaraman. Detection of Logical Errors in Decision Table Programs. *Communications of the ACM*, 21(12):1016–1025, 1978.

[Jaffray, 1999] J.-Y. Jaffray. Rational decision making with imprecise probabilities. In *1st International Symposium on Imprecise Probabilities and Their Applications*, 1999.

[Jakobovits and Vermeir, 1999] H. Jakobovits and D. Vermeir. Robust Semantics for Argumentation Frameworks. *Journal of Logic and Computation*, 9:215–261, 1999.

[Jefferys *et al.*, 2006] B.R. Jefferys, L.A. Kelley, M.J. Sergot, J. Fox, and M.J. E. Sternberg. Capturing expert knowledge with argumentation: a case study in bioinformatics. *Bioinformatics*, 22(8):924–933, 2006.

[Kak *et al.*, 1990] A.C. Kak, K.M. Andress, C. Lopez-Abadia, M.S. Carroll, and J.R. Lewis. Hierarchical Evidence Accumulation in the Pseiki System and Experiments in Model-Driven Mobile Robot Navigation. In *5th International Conference on Uncertainty in Artificial Intelligence.* 1990.

[Kane, 1999] S.A. Kane. What happened to the utility functions ? Imprecise expectations of security prices. *International Review of Financial Analysis*, 8:2:165–175, 1999.

[Karmarkar, 1984] N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4:373–395, 1984.

[Keeney and Raiffa, 1976] R.L. Keeney and H. Raiffa. *Decisions With Multiple Objectives.* John Wiley and Sons, 1976.

[Khachiyan *et al.*, 2005] L. Khachiyan, E. Boros, K. Elbassioni, and V. Gurvich. A new algorithm for the hypergraph transversal problem. In *Annual International Conference on Computing and Combinatorics*, 2005.

[Khachiyan, 1979] L. Khachiyan. A polynomial time algorithm for linear programming. *Dokl. Akad. Nauk. SSSR*, 244:1093–1096, 1979.

[King, 1968] P.J.H. King. Ambiguity in limited entry decision tables. *Communications of the ACM*, 11(10):680–684, 1968.

[Klir, 2006] G.J. Klir. *Uncertainty and Information.* John Wiley and Sons, 2006.

[Knight, 1921] F.H. Knight. *Risk, Uncertainty and Profit.* Houghton Mufflin, 1921.

[Kohlas *et al.*, 2002] J. Kohlas, D. Berzati, and R. Haenni. Probabilistic Argumentation Systems and Abduction. *Annals of Mathematics and Artificial Intelligence*, 34(1-3):177–195, 2002.

[Kowalski, 1979a] R.A. Kowalski. Algorithm= logic + control. *Communications of the ACM*, 22:424–431, 1979.

[Kowalski, 1979b] R.A. Kowalski. *Logic for Problem Solving.* North Holland, 1979.

[Kozine and Utkin, 2005] I.O. Kozine and L.V. Utkin. Constructing imprecise probability distributions. *International Journal of General Systems*, 34(4):401–408, 2005.

[Krause and Clark, 1993] P. Krause and D. Clark. *Representing uncertain knowledge.* Kluwer Academic Press, 1993.

[Krause *et al.*, 1995] P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11:113–131, 1995.

[Laplace, 1812] P.S. Laplace. *Théorie Analytique des Probabilités.* Courcier, 1812.

[Lee, 1971] W. Lee. *Decision Theory and Human Behavior.* John Wiley and Sons, 1971.

[Levi, 1980] I. Levi. *The Entreprise of Knowledge.* MIT Press, 1980.

[Lew and Tamanaha, 1976] A. Lew and D. Tamanaha. Decision Table Programming and Reliability. In *2nd international Conference on Software Engineering*, 1976.

[Lew, 1984] A. Lew. Proof of Correctness of Decision Table Programs. *Computer Journal*, 27(3):230–237, 1984.

[Lhabitant, 2001] F.-S. Lhabitant. *Hedge Funds: Quantitative Insights.* Wiley, 2001.

[Ljapunov, 1901] A.M. Ljapunov. *Nouvelle forme du théoreme de la limite centrale.* Academie Imperiale de Sciences de St.-Petersbourg, 1901.

[Lo *et al.*, 2000] A.W. Lo, H. Mamaysky, and J. Wang. Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation. *Journal of Finance*, LV(4), 2000.

[Markowitz, 1952] H. Markowitz. Portfolio selection. *Journal of Finance*, 7:77–91, 1952.

[Matt and Toni, 2008a] P.-A. Matt and F. Toni. Basic influence diagrams and the liberal stable semantics. In *2nd International Conference on Computational Models of Argument*, 2008.

[Matt and Toni, 2008b] P.-A. Matt and F. Toni. A game-theoretic measure of argument strength for abstract argumentation. In *11th European Conference on Logics in Artificial Intelligence*, 2008.

[McCarthy, 1990] J. McCarthy. Some Expert Systems Need Common Sense. In *Formalizing Common Sense: Papers by John McCarthy.* 1990.

[McGeachie and Doyle, 2002] M. McGeachie and J. Doyle. Efficient Utility Functions for Ceteris Paribus Preferences. In *18th National Conference on Artificial Intelligence*, 2002.

[Mendelson, 2005] P. Mendelson. *Texas Hold'em Poker: Begin and Win.* Elliot Right Way Books, 2005.

[Milner, 1977] D.A. Milner. The use of decision table logic in a manufacturing system. *International Journal of Production Research*, 15(1), 1977.

[Moore and Cullinane, 1969] J.M. Moore and T. Cullinane. Networks put your problems in computer language. *Industrial Engineering*, 28(1), 1969.

[Morge and Mancarella, 2007] M. Morge and P. Mancarella. The hedgehog and the fox. An argumentation-based decision support system. In *4th International Workshop on Argumentation in Multi-Agent Systems*, 2007.

[Mozina *et al.*, 2007] M. Mozina, J. Zabkar, and I. Bratko. Argument based machine learning. *Artificial Intelligence*, 171(10-15):922–937, 2007.

[Natarajan *et al.*, 2005] K. Natarajan, D. Pachamanova, and M. Sim. Constructing risk measures from uncertainty sets. Working Paper, 2005.

[Neyman, 1950] J. Neyman. *First Course in Probability and Statistics.* Henry Holt, 1950.

[Nielsen and Parsons, 2006] S.H. Nielsen and S. Parsons. Computing preferred extensions for argumentation systems with sets of attacking arguments. In *1st International Conference on Computational Models of Argument*, 2006.

[Nielsen and Parsons, 2007] S.H. Nielsen and S. Parsons. An application of formal argumentation: Fusing bayesian networks in multi-agent systems. *Artificial Intelligence*, 171:754–775, 2007.

[Osler, 2003] C.L. Osler. Currency Orders and Exchange Rate Dynamics: An Explanation for the Predictive Success of Technical Analysis. *The Journal of Finance*, 58(5):1791–1819, 2003.

[Parsons, 1997] S. Parsons. Normative Argumentation and Qualitative Probability. In *1st International Joint Conference on Qualitative and Quantitative Practical Reasoning*, 1997.

[Pearl, 1986] J. Pearl. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29(3):241–288, 1986.

[Pelessoni and Vicig, 2001] R. Pelessoni and P. Vicig. Coherent risk measures and upper previsions. In *2nd International Symposium on Imprecise Probabilities and Their Applications*, 2001.

[Perelman and Olbrechts-Tyteca, 1969] C. Perelman and L. Olbrechts-Tyteca. *The new rhetoric: A treatise on argumentation.* University of Notre Dame Press, 1969.

[Pollock, 1992] J.L. Pollock. How to reason defeasibly. *Artificial Intelligence*, 57:1–42, 1992.

[Pollock, 2001] John L. Pollock. Defeasible reasoning with variable degrees of justification. *Artificial Intelligence*, 133(1–2):233–282, 2001.

[Poole, 1993] D. Poole. Probabilistic Horn Abduction and Bayesian Networks. *Artificial Intelligence*, 64(1):81–129, 1993.

[Prakken and Sartor, 1997] H. Prakken and G. Sartor. Argument-based extended logic programming with defeasible priorities. *Journal of Applied Non-classical Logics*, 7:25–75, 1997.

[Prakken, 2006] H. Prakken. Combining sceptical epistemic reasoning with credulous practical reasoning. In *1st International Conference on Computational Models of Argument*, 2006.

[Rahwan and Larson, 2008] I. Rahwan and K. Larson. Mechanism design for abstract argumentation. In *7th International Conference on Autonomous Agents and Multiagent Systems*, 2008.

[Raiffa, 1968] H. Raiffa. *Decision Analysis: Introductory Lectures on Choice Under Uncertainty.* Addision-Wesley, 1968.

[Refenes *et al.*, 1997] A.N. Refenes, N. Burgess, and Y. Bentz. Neural networks in financial engineering: a study in methodology. *IEEE Trans. Neural Networks*, 6(8):1222–1267, 1997.

[Reilly *et al.*, 1987] K.D. Reilly, A. Salah, and C.-C. Yang. A Logic Programming Perspective on Decision Table Theory and Practice. *Data Knowledge Engineering*, 2:191–212, 1987.

[Rose, 1998] K. Rose. Deterministic annealing for clustering, compression, classification, regression, and related optimization problems. In *Proceedings of the IEEE*, 1998.

[Russel and Norvig, 2002] S. Russel and P. Norvig. *Artificial Intelligence: A Modern Approach.* Second Edition, Prentice-Hall, 2002.

[Saffiotti and Umkehrer, 1991] A. Saffiotti and E. Umkehrer. Pulcinella: A General Tool for Propagating Uncertainty in Valuation Networks. In *7th International Conference on Uncertainty in Artificial Intelligence*, 1991.

[Savage, 1972] L.J. Savage. *The foundations of statistics.* Dover, 1972.

[Schmidt and Kavanagh, 1964] D.T. Schmidt and T.F. Kavanagh. Using decision structure table. *Datamation*, 42(10), 1964.

[Sentz and Scott, 2002] K. Sentz and F. Scott. *Combination of Evidence in Dempster-Shafer Theory.* Sandia National Laboratories, 2002.

[Shafer, 1976] G. Shafer. *A Mathematical Theory of Evidence.* Princeton University Press, 1976.

[Shannon, 1948] C.E. Shannon. Mathematical Theory of Communication. *The Bell System Technical Journal*, 27:379–656, 1948.

[Shapiro and Gross, 1981] S.S. Shapiro and A.J. Gross. *Statistical Modeling Techniques.* M. Dekker, 1981.

[Sharpe, 1966] W.F. Sharpe. Mutual Fund Performance. *The Journal of Business*, 39:119–138, 1966.

[Shwayder, 1975] K. Shwayder. Combining Decision Rules in a Decision Table. *Communications of the ACM*, 18(8):476–480, 1975.

[Simutis, 2000] R. Simutis. Fuzzy logic based stock trading system. In *IEEE/IAFE Conference on Computational Intelligence of Financial Engineering*, 2000.

[Smets, 1990] P. Smets. Constructing the pignistic probability function in a context of uncertainty. In *5th International Conference on Uncertainty in Artificial Intelligence*, 1990.

[Takahashi and Sawamura, 2004] T. Takahashi and H. Sawamura. A logic of multiple-valued argumentation. In *3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 2004.

[Tan and Pearl, 1994] S. Tan and J. Pearl. Qualitative Decision Theory. In *12th National Conference on Artificial Intelligence*, pages 928–932, 1994.

[Toni, 2007] F. Toni. Assumption-based argumentation for closed and consistent defeasible reasoning. In *21st Annual Conference of The Japanese Society for Artificial Intelligence*, 2007.

[Toni, 2008] F. Toni. Assumption-based argumentation for epistemic and practical reasoning. In *Computable Models of the Law*, Lecture Notes in Computer Science, 2008.

[Toulmin, 1958] S.E. Toulmin. *The Uses of Argument.* Cambridge University Press, 1958.

[Treynor and Ferguson, 1985] J.L. Treynor and R. Ferguson. In defense of technical analysis. *Journal of Finance*, 40, 1985.

[Tversky, 1967] A. Tversky. Additivity, utility, and subjective probability. *Journal of Mathematical Psychology*, 4:175–201, 1967.

[Utkin, 2006] L.V. Utkin. A method for processing the unreliable expert judgments about parameters of probability distributions. *European Journal of Operational Research*, 175(1):385–398, 2006.

[Vanthienen and Dries, 1993] J. Vanthienen and E. Dries. Illustration of a Decision Table Tool for Specifying and Implementing Knowledge Based Systems. In *5th International Conference on Tools with Artificial Intelligence*, 1993.

[Vanthienen and Robben, 1993] J. Vanthienen and F. Robben. Developing Legal Knowledge Based Systems Using Decision Tables. In *4th International Conference on Artificial Intelligence and Law*, 1993.

[Verhelst, 1980] M. Verhelst. *De Pratik van Enstcheidungstabelle.* Kluwer, 1980.

[Vicig, 2008] P. Vicig. Financial risk measurement with imprecise probabilities. *International Journal of Approximate Reasoning*, 49(1):159–174, 2008.

[von Neumann and Morgenstern, 1944] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, 1944.

[von Neumann, 1928] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.

[Vreeswijk and Prakken, 2000] G.A.W. Vreeswijk and H. Prakken. Credulous and Sceptical Argument Games for Preferred Semantics. In *7th European Conference on Logics in Artificial Intelligence*, 2000.

[Walley and de Cooman, 2001] P. Walley and G. de Cooman. A behavioral model for linguistic uncertainty. *Information Sciences*, 134:1–37, 2001.

[Walley and Fine, 1979] P. Walley and T. L. Fine. Varieties of modal (classificatory) and comparative probability. *Synthese*, 41:321–374, 1979.

[Walley, 1991] P. Walley. *Statistical reasoning with Imprecise Probabilities.* Chapman and Hall, 1991.

[Walley, 1996] P. Walley. Measures of Uncertainty in Expert Systems. *Artificial Intelligence*, 83(1):1–58, 1996.

[Walley, 2000] P. Walley. Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning*, 24:125–148, 2000.

[Wigmore, 1940] J.H. Wigmore. *Evidence in trials at common law.* Little, Brown, 1940.

[Wilson, 1927] E. B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22:209–212, 1927.

[Yu and Singh, 2002] B. Yu and M.P. Singh. Distributed Reputation Management for Electronic Commerce. *Computational Intelligence*, 18(4):535–549, 2002.

[Zadeh, 1978] L.A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.

[Zaffalon, 2002a] M. Zaffalon. Exact treatment of missing data. *Journal of Statistical Planning and Inference*, 105:105–122, 2002.

[Zaffalon, 2002b] M. Zaffalon. The naive credal classifier. *Journal of Statistical Planning and Inference*, 105:5–21, 2002.

[Zaffalon, 2005] M. Zaffalon. Credibility via imprecise probability. *International Journal of Approximate Reasoning*, 39(2-3):115–121, 2005.

[Zuckerman, 2009] I. Zuckerman. Towards probabilistic argumentation. *Argumentation in Artificial Intelligence*, pages 443–462, 2009.