



HAL
open science

Decision Under Uncertainty

Christophe Gonzales, Patrice Perny

► **To cite this version:**

Christophe Gonzales, Patrice Perny. Decision Under Uncertainty. Marquis, Pierre; Papini, Odile; Prade, Henri. A Guided Tour of Artificial Intelligence Research, I, Springer, pp.549-586, 2020, Knowledge Representation, Reasoning and Learning. hal-02860485

HAL Id: hal-02860485

<https://hal.science/hal-02860485>

Submitted on 12 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Decision under uncertainty

Christophe Gonzales¹ and Patrice Perny²

Abstract The goal of this chapter is to provide a general introduction to decision making under uncertainty. The mathematical foundations of the most popular models used in artificial intelligence are described, notably the *Expected Utility model* (EU), but also new decision making models, like *Rank Dependent Utility* (RDU), which significantly extend the descriptive power of EU. Decision making under uncertainty naturally involves risks when decisions are made. The notion of risk is formalized as well as the attitude of agents w.r.t. risk. For this purpose, probabilities are often exploited to model uncertainties. But there exist situations in which agents do not have sufficient knowledge or data available to determine these probability distributions. In this case, more general models of uncertainty are needed and this chapter describes some of them, notably *belief functions*. Finally, in most artificial intelligence problems, sequences of decisions need be made and, to get an optimal sequence, decisions must not be considered separately but as a whole. We thus study at the end of this chapter models of sequential decision making under uncertainty, notably the most widely used graphical models.

1 Introduction

Uncertainty and, more generally, decision making under uncertainty are central in artificial intelligence (AI). Indeed, even though AI addresses a wide range of problems, most of them involve to some extent uncertainties. This is the case, for instance, in diagnosis [Franklin et al., 1991; Jensen et al., 2001], prediction [Conati et al., 1997; Horvitz et al., 1998], robotics [Argall et al., 2009], planning [Puterman, 1994], machine learning and image processing [Doucet and Johansen, 2011]. *Deci-*

¹Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France,
e-mail: christophe.gonzales@lis-lab.fr

²Sorbonne Université, CNRS, LIP6, Paris, France,
e-mail: patrice.perny@lip6.fr

sion Theory proposes sophisticated mathematical tools to reason in such contexts. These tools can be roughly divided into two classes: *decision support systems* (or *decision aiding systems*) and *automatic decision systems*. The goal of the former is to help, to guide, human agents to make their decisions, especially when those are very complex (e.g., when many conflicting “decision criteria” need be taken into account). As such, they prove to be very useful in critical domains like, e.g., in medical decision making [Franklin et al., 1991; Bleichrodt, 1996], in space shuttle flying assistance [Horvitz and Barry, 1995] and in strategic applications like choosing the location of a future airport [Keeney and Raiffa, 1993]. As for *automatic decision systems*, they are designed to enable non-human agents (robots or software) to choose the best actions to reach their goals. They prove to be useful, for instance, in robotics [Argall et al., 2009], in missiles or drones control [Dasgupta, 2006], in serious games [Sordoni et al., 2010].

The decision theoretic tools developed in both decision aiding systems and automatic decision systems rely on mathematical models for representing the agents’ preferences. Those enable, for instance, to justify why the majority of agents who are asked which envelope they would prefer to get between two envelopes A and B containing 100 € and 200 € respectively answer they prefer B . Actually, Decision “choose envelope B ” yields the consequence of winning 200 €, which is often judged more preferable than the consequence of winning only 100 €. Unfortunately, in real-world applications, decisions are not so simple especially because, when decisions are made, their consequences are to some extent uncertain, i.e., they depend on the occurrences of some events that are still uncertain at the time the decision is made. For instance, when a physician must determine the best treatment to prescribe to his patient, his diagnosis does not allow him to know for sure the exact illness of the patient. Therefore, to take into account all the parameters involved in the decision making, decision models under uncertainty combine two components: a “preference model” and a “representation of uncertainties”.

Models and algorithms originating from the field of “decision theory under uncertainty” are widely used in artificial intelligence. This is essentially due to two reasons: i) these models have strong mathematical foundations; and ii) their axiomatic justifications rely on rationality arguments with which everybody seems to agree. As a consequence, the conclusions reached by these algorithms can be justified using essentially common sense arguments. The rationality justification incited the majority of the artificial intelligence community to adopt the “*expected utility*” (EU) criterion as *its* decision criterion. In Section 2, we describe this criterion and we focus on its axiomatic foundation. The notion of risk and how it translates in the EU model are also investigated. Finally, this section ends with a discussion on the descriptive limits of this model. This naturally calls for other models that can go beyond these limits: in Section 4, the emphasis is made on representing uncertainties outside the probabilistic framework while, in Section 3, the linear model of preferences itself is questioned. Finally, Section 5, addresses sequential decision problems, especially their representations and the issues raised by non-linear models.

2 The Expected Utility Criterion (EU)

Let \mathcal{D} denote the set of decisions that can be made by an agent. In the rest of this chapter, we assume that the agents have well defined preferences on \mathcal{D} and we denote by $\succsim_{\mathcal{D}}$ their preference relation. Thus, $d_1 \succsim_{\mathcal{D}} d_2$ means that the agent prefers Decision d_1 to Decision d_2 or she is indifferent between the two decisions. Strict preference is denoted as usual by $\succ_{\mathcal{D}}$. As we have seen in the introduction, when making decisions under *certainty*, preferring d_1 to d_2 amounts to prefer the consequence yielded by decision d_1 to that yielded by d_2 . Let \mathcal{X} denote the space of all the possible consequences. Preference relation $\succsim_{\mathcal{D}}$ over \mathcal{D} is thus induced by preference relation $\succsim_{\mathcal{X}}$ over \mathcal{X} as follows: $d_1 \succsim_{\mathcal{D}} d_2$ if and only if $x(d_1) \succsim_{\mathcal{X}} x(d_2)$, where $x(d)$ represents the consequence of decision d .

Under uncertainty, i.e., when the consequence of a decision is not fully known when the decision is made, the equivalence between $\succsim_{\mathcal{D}}$ and $\succsim_{\mathcal{X}}$ does not exist anymore. However, in this case, it is reasonable to assume that preference relation $\succsim_{\mathcal{D}}$ takes into account not only the agent's preferences over the consequences of the decisions but also her attitude w.r.t. the uncertainty over the fulfillment of these consequences. As an example, when asked to choose between an envelope A containing 100 € and an envelope B randomly chosen among a heap of 100 envelopes in which 97 contain 1 € and 3 contain 1000 €, most of people prefer envelope A because the probability of getting 1000 € with envelope B is too low. From this simple example, we can deduce that the agent translates the uncertainty over the 100 envelopes into an uncertainty over the amount of money contained in envelope B , i.e., on the consequence yielded by decision d_2 , and the decision is made taking into account the latter. Before investigating further how agents make their decisions, we need to define more precisely the notion of uncertainty from the agent's perspective. Similarly to probability theory, we need to define what are an event and an elementary event: an event is a set of possible results of a random experiment (above, the choice of an envelope) and an elementary event, which is called a *state of nature* in decision theory, corresponds to only one possible result. Thus, if the envelopes are numbered from 1 to 100, the fact that the envelope chosen is the 3rd one is a state of nature whereas the fact that it has an even number is a (non-elementary) event. Let \mathcal{S} and $\mathcal{A} = 2^{\mathcal{S}}$ denote the set of the states of natures and the set of events respectively. The above example of the envelopes suggests that probabilities are an attractive representation of the uncertainties with which the agent has to cope. This is the very representation exploited in the model presented in this section. Note however that this is not the only possible choice, as we will see later.

In the EU model, uncertainties are represented by probabilities and the decision criterion is simply the maximum of the expectation of the satisfaction provided by the decisions (a.k.a. a *utility*). Let $u : \mathcal{X} \mapsto \mathbb{R}$ be a function such that $x \succsim_{\mathcal{X}} y$ if and only if $u(x) \geq u(y)$. Such a function is called a *utility function* or, for short, a *utility*. A utility function therefore assigns to each consequence a real number such that the preferred the consequence, the higher the number. The utility expectation criterion has been popularized by Daniel Bernoulli in the 18th century [Bernoulli, 1738], although a letter by Gabriel Cramer to Nicolas Bernoulli seems to establish

that Cramer proposed this criterion earlier. But its modern axiomatic foundations are due, on one hand to von Neumann and Morgenstern [von Neumann and Morgenstern, 1944] and, on the other hand, to Savage [Savage, 1954]. These two axiomatics differ essentially by the fact that the former assume the existence of a probability distribution on $(\mathcal{S}, \mathcal{A})$ whereas the latter derives its existence from the rationality of the agent.

2.1 von Neumann-Morgenstern's Axiomatic Foundation

von Neumann and Morgenstern assume that $(\mathcal{S}, \mathcal{A}, P)$ is a probabilistic space. In other words, P is a probability distribution over $(\mathcal{S}, \mathcal{A})$. As we have seen before, for each decision d , this distribution induces another probability distribution P_d over the space of consequences $(\mathcal{X}, \mathcal{C})$, where $\mathcal{C} = 2^{\mathcal{X}}$. When the support of P_d is finite, i.e., when the number of possible consequences (those with nonzero probabilities) resulting from making decision d is finite, distribution P_d is called a *lottery*. A lottery can therefore be represented as a tuple $\langle x_1, p_1; \dots; x_n, p_n \rangle$, where the x_i 's are some consequences and the p_i 's correspond to their probability of occurrence. Note that a lottery is a representation, a summary, of what a decision really is: indeed it only represents synthetically what can result from making the decision. If this summary is faithful, then we can conclude that there exists an equivalence between the preferences of the agent over the decisions and those over their associated lotteries. Let \mathcal{L} be the set of all the possible lotteries and let \succsim be the preference relation of the agent over \mathcal{L} . Then we can conclude that $d_1 \succsim_{\mathcal{D}} d_2$ if and only if $P_{d_1} \succsim P_{d_2}$, where P_d represents the lottery associated with decision d . von Neumann and Morgenstern [von Neumann and Morgenstern, 1944] show that the preferences over \mathcal{L} (and thus over \mathcal{D}) of any *rational* agent necessarily follow the expected utility criterion:

$$\text{For all } P, Q \in \mathcal{L}, \quad P \succsim Q \iff \sum_{i=1}^n p_i u(x_i) \geq \sum_{j=1}^r q_j u(y_j), \quad (1)$$

where $P = \langle x_1, p_1; \dots; x_n, p_n \rangle$, $Q = \langle y_1, q_1; \dots; y_r, q_r \rangle$, and $u(x)$ is a utility function over the space of consequences \mathcal{X} (this function is called the *von Neumann-Morgenstern utility function*). The first axiom exploited by von Neumann and Morgenstern to prove this result is the following:

Axiom 1 (Complete weak order) \succsim is a complete weak order on \mathcal{L} . In other words, \succsim is reflexive (for all $P \in \mathcal{L}$, $P \succsim P$), transitive (for all $P, Q, R \in \mathcal{L}$, $(P \succsim Q) \wedge (Q \succsim R) \implies P \succsim R$) and complete (for all $P, Q \in \mathcal{L}$, $(P \succsim Q) \vee (Q \succsim P)$). In addition, \succsim is non-trivial, i.e., there exist $P, Q \in \mathcal{L}$ such that $P \succ Q$.

This axiom simply expresses the idea that, given any pair of lotteries, the agent is always capable of determining which one she prefers (completeness) and that if she prefers P to Q and Q to R , then, logically, she will also prefer P to R . This last property conveys some kind of rationality, although it is possible to find examples in

which rational decision makers have intransitive preferences [Anand, 1993]. Finally, non-triviality just guarantees that we study only situations in which all the decisions are not judged as equivalent by the agent, the decision maker (if this were not the case, a decision making model would be useless).

For the next two axioms, we need to define *mixtures of lotteries*: let P and Q be two lotteries and let $\lambda \in [0, 1]$ be a real number, then $R = \lambda P + (1 - \lambda)Q$, the mixture of P and Q w.r.t. λ , represents the lottery such that, for any consequence $x \in \mathcal{X}$, the probability of occurrence of x is $R(x) = \lambda P(x) + (1 - \lambda)Q(x)$. Intuitively, a mixture essentially amounts to create R in two steps: first, a coin with a probability λ to land on head (and therefore $1 - \lambda$ to land on tail) is flipped; second, if the coin landed on head, then we get lottery P , else we get lottery Q . The probability of occurrence of each consequence x is consequently $\lambda P(x) + (1 - \lambda)Q(x)$.

Axiom 2 (continuity) For all $P, Q, R \in \mathcal{L}$ such that $P \succ Q \succ R$, there exist $\alpha, \beta \in]0, 1[$ such that:

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

This axiom conveys the idea that if the agent strictly prefers P to Q , then a lottery resulting from a very small perturbation of P should still be preferred to Q . For instance, if $P = \langle 100, 1 \rangle$, i.e., P is the lottery which yields 100€ with certainty, and if $Q = \langle 10, 1 \rangle$ and $R = \langle 5, 1 \rangle$, then an agent who likes money should have the following preference relation: $P \succ Q \succ R$. If $\alpha = 1 - 10^{-20}$, lottery $\alpha P + (1 - \alpha)R$ is equal to $\langle 100, 1 - 10^{-20}; 5, 10^{-20} \rangle$. The chance of receiving 5€ is so small that the agent is almost assured to win 100€, which is preferable to Q . Therefore, it is very likely that the agent prefers $\alpha P + (1 - \alpha)R$ to Q . A similar argument can be used with β very close to 0. Here again, Axiom 2 seems quite reasonable. The last axiom used by von Neumann and Morgenstern is the following:

Axiom 3 (independence) For every $P, Q, R \in \mathcal{L}$ and every $\alpha \in]0, 1[$:

$$P \succsim Q \iff \alpha P + (1 - \alpha)R \succsim \alpha Q + (1 - \alpha)R.$$

The interpretation of this axiom follows that of mixtures. We have seen that $\alpha P + (1 - \alpha)R$ corresponds to a lottery created in two steps: first a coin is flipped, with probabilities α and $1 - \alpha$ to land on head and tail respectively and, then, depending on the side on which the coin landed, the agent receives lottery P or R . Following this principle, Axiom 3 can be interpreted as follows: if the coin lands on tail, from both lotteries $\alpha P + (1 - \alpha)R$ and $\alpha Q + (1 - \alpha)R$, the agent receives the same induced lottery R so, logically, in this case, she should be indifferent between $\alpha P + (1 - \alpha)R$ and $\alpha Q + (1 - \alpha)R$. If, on the other hand, the coin lands on head, then, from $\alpha P + (1 - \alpha)R$ and $\alpha Q + (1 - \alpha)R$, she receives lotteries P and Q respectively. As she (weakly) prefers P to Q or is indifferent between these two lotteries, she should also prefer $\alpha P + (1 - \alpha)R$ to $\alpha Q + (1 - \alpha)R$ or be indifferent between them, hence the axiom.

The three above axioms therefore express properties that can be expected from a rational agent. As the next theorem shows, they imply that there exists a unique

decision criterion representing the preferences of the agent, and this one is precisely the expected utility criterion:

Theorem 1 (von Neumann-Morgenstern). *The following two assertions are equivalent:*

1. *The preference relation \succsim over \mathcal{L} satisfies axioms 1, 2 and 3.*
2. *\succsim is representable by a utility function $U : \mathcal{L} \mapsto \mathbb{R}$ such that $U(P) = \sum_{i=1}^n p_i u(x_i)$, where $u(x_i) = U(\langle x_i, 1 \rangle)$.*

Function $u : \mathcal{X} \mapsto \mathbb{R}$ is called the von Neumann-Morgenstern utility function of the agent and is unique up to scale and location (i.e., up to strictly positive affine transforms).

This strong relationship between rationality and the EU criterion explains why EU is so popular in the artificial intelligence community but also among decision theorists and operations research scientists. Note also that the above theorem can be generalized, notably by using more general probability measures [Fishburn, 1970, 1982]. There also exist other axiomatics, like, e.g., the one provided in [Jensen, 1967], [Herstein and Milnor, 1953] or in [Fishburn and Roberts, 1978].

The von Neumann-Morgenstern axiomatics raises one issue: it assumes the existence of an “*objective*” probability distribution over the space of the states of nature—in decision theory, this situation is called “*decision under risk*” [Knight, 1921]—and one may wonder whether such a hypothesis is so reasonable in practical decision theory problems and, more generally, in artificial intelligence. As we will see, the answer to this question seems positive because the existence of a probability distribution over the states of nature necessarily follows from the rationality of the agent. This idea has been initially introduced in [Ramsey, 1931] but went largely unnoticed until the seminal book by Savage [Savage, 1954] got published. Sixty years later, the idea that probabilities are the only “rational” representation of uncertainties is so deeply anchored into people’s minds that, up to recently, it was very difficult in artificial intelligence to imagine a rational decision making process outside the EU framework¹. When the probability distribution over the states of nature results from the rationality of the agent, this distribution is said to be “*subjective*” and the decisional context is called “*decision under uncertainty*” instead of “*decision under risk*”, which is dedicated to the case of objective probabilities.

Let us now study Savage’s axiomatics [Savage, 1954], which has led to the decision model called “*Subjective Expected Utility*” (SEU). Of course, since the probability distribution over the space of the states of nature $(\mathcal{S}, \mathcal{A})$ results from the rationality of the decision maker, this is no more a primitive of the decisional language. The primitive, here, is called an “*act*”. Quite similarly to a lottery, this corresponds to the representation/summary of a decision but its description is more

¹ Outside the EU framework, the behavior of an agent cannot be rational (w.r.t. Savage’s meaning) and, therefore, it is thought in artificial intelligence that such a behavior must be proscribed. In the 70’s and 80’s, decision theorists, notably Kahneman, Tversky and Quiggin, suggested that Savage’s rationality was not the only possible form of rationality and they proposed to depart from the Savagian framework and developed their own kinds of “rationality”. This paved the way to new decision models like, e.g., RDU, that recently attracted the attention of AI researchers.

precise than a lottery. An *act* f is a function whose domain and codomain are \mathcal{S} and \mathcal{X} respectively. In other words, to each state of nature e , function f assigns the consequence that would result from the occurrence of e if the decision represented by f were made. For pedagogical purposes, let us first consider “*simple acts*”, i.e., finite-valued and \mathcal{A} -measurable functions. For any simple act f , there exists a finite partition $\{E_i, i \in I\}$ of \mathcal{S} such that, for all $i \in I$, $\{f(e) : e \in E_i\} = \{c_i\}$, where $c_i \in \mathcal{X}$ is a consequence. To put it differently, a simple act yields a finite set of possible consequences and those depend on the realization of some states of nature belonging to E_i . To draw a parallel with lotteries, a simple act corresponds to a description of a lottery with finite support, although this description is more precise than just the lottery. Table 1 highlights their differences: this table shows two acts, f_1 and f_2 . In the former, when state e_1 obtains, the resulting consequence is c_1 . For both acts, the probabilities of getting c_1 and c_2 are 0.3 and 0.7 respectively. Consequently, both acts correspond to the same lottery $\langle c_1, 0.3; c_2, 0.7 \rangle$. However, as can be seen in the table, act f_1 is different from f_2 . An act is therefore a description of a decision which is more precise than a lottery. In the sequel, we will explain the interpretations of new concepts and axioms using simple acts but those apply on general acts, not only on simple ones. Finally, let δ_c denote the “*constant*” act yielding consequence c , i.e., the act such that $\delta_c(\mathcal{S}) = \{c\}$. To illustrate visually what acts represent and how they will be combined, we will use figures in which the X and Y axes represent sets \mathcal{S} and \mathcal{X} respectively. In this setting, a simple act is just a stepwise function, as shown in Figure 1.

event	probability	act f_1	act f_2
e_1	0.3	cons. c_1	cons. c_2
e_2	0.3	cons. c_2	cons. c_1
e_3	0.4	cons. c_2	cons. c_2

Table 1 Comparison between acts and lotteries.

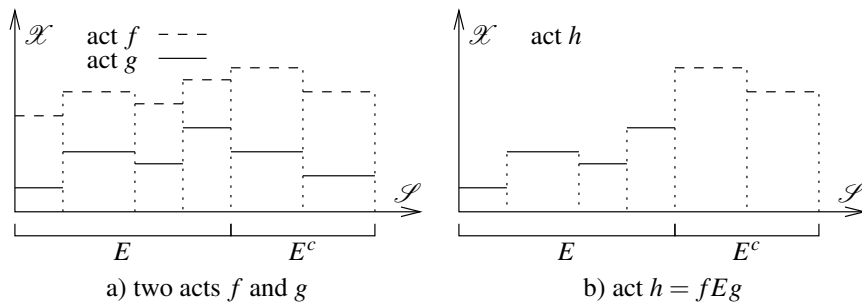


Fig. 1 The concept of splicing.

In the axiomatic theory of von Neumann-Morgenstern, one of the key ideas was the possibility to combine by “mixture” different lotteries to produce new ones. This was the core of their proof. Here, there exists an equivalent operation on acts, that we will call a “*splicing*” to distinguish it from mixtures². Let f and g be two acts and let $E \subseteq \mathcal{S}$ be an event. The splicing of f and g w.r.t. E , denoted by fEg is the act $h = fEg$ such that $h(s) = f(s)$ for all $s \in E$ and $h(s) = g(s)$ for all $s \in E^C = \mathcal{S} \setminus E$. Figure 1 illustrates this operation. All the primitives and operations necessary to describe the axiomatic theory of SEU are defined, we can now study the axioms provided by Savage. The first one corresponds in essence to Axiom 1 of von Neumann-Morgenstern.

Axiom 4 (P1: Weak order on acts) *The set of all the acts is closed under splicing and there exists a complete weak order \succsim over the set of acts.*

But Savage’s key axiom is the “*Sure Thing Principle*”³.

Axiom 5 (P2: Sure Thing Principle) *For all acts $f, g, h, k \in \mathcal{X}^{\mathcal{S}}$ and all $E \subseteq \mathcal{S}$:*

$$fEh \succsim gEh \iff fEk \succsim gEk.$$

This axiom corresponds in spirit to the independence axiom of von Neumann-Morgenstern. Figure 2 provides an illustration: let fEh and gEh be two acts. They yield the same consequences over E^C . Consequently, if the state of nature that obtains belongs to E^C , the agent should be indifferent between both acts. So, if globally, she prefers fEh to gEh , this means that, over E , she prefers the consequences yielded by f to those by g . Now, substitute the common part of both acts h on E^C by another act k . Then, the resulting acts are fEk and gEk . These new acts yield precisely the same consequences over E^C , so the agent should still be indifferent between them if the state of nature that obtains belongs to E^C . And if the state that




Fig. 2 Illustration of the Sure Thing Principle.

² In his book, Savage did not name this operation. The term “splicing” was introduced in [Gilboa, 2009].

³ Most authors name P2 as the “sure thing principle” but it was pointed out by Peter Wakker that, in Savage’s book, the sure thing principle refers to axioms P2, P3 and P7.

obtains belongs to E , then both acts yield the same consequences as fEh and gEh , so, globally, if the agent preferred fEh to gEh , she should also prefer fEk to gEk . In other words, the sure thing principle states that, when comparing two acts, the agent only compares the acts on the events on which they differ. This axiom looks quite reasonable.

Note that Axiom P2 implies the existence of a weak order \succsim_E for every event E defined as $f \succsim_E g$ if and only if $fEh \succsim gEh$ for all h . The next axiom exploits this new preference relation to guarantee, using constant acts, that the agent has a well-defined preference relation $\succsim_{\mathcal{X}}$ over the space of consequences. This axiom relies on *non-null* events, i.e., on events E such that there exist at least two acts f and g such that $f \succ_E g$.

Axiom 6 (P3: Preferences among consequences) For all consequences $x, y \in \mathcal{X}$ and all non-null events $E \subseteq \mathcal{S}$, $\delta_x \succsim_E \delta_y$ if and only if $x \succsim_{\mathcal{X}} y$, where δ_x and δ_y are constant acts.

In the SEU framework, the existence of an “objective” probability distribution over the states of nature is never assumed. Rather, the existence of a “subjective” distribution results from the beliefs of the agent herself. The agent must therefore have beliefs that an event A is more or less likely to occur than another event B . This is exactly what the next axiom induces:

Axiom 7 (P4: Preferences over events) For all consequences $x, x', y, y' \in \mathcal{X}$ such that $x \succ_{\mathcal{X}} y$ and $x' \succ_{\mathcal{X}} y'$, and for all $A, B \subseteq \mathcal{S}$,

$$\delta_x A \delta_y \succ \delta_x B \delta_y \iff \delta_{x'} A \delta_{y'} \succ \delta_{x'} B \delta_{y'}.$$

Figure 3 illustrates this axiom: acts $\delta_x A \delta_y$ and $\delta_x B \delta_y$ differ only on the gray area. On this one, $\delta_x A \delta_y$ yields consequence x and $\delta_x B \delta_y$ yields y , which is not preferred to x . This explains why $\delta_x A \delta_y \succ \delta_x B \delta_y$. In this figure, the existence of the gray area results from the fact that A contains B and, consequently, it is more “probable” to happen than B . In general, it can be shown that, whenever the agent believes that A is more likely to happen than B , then the preferences of the agent satisfy Axiom P4. Axiom P5 below expresses the fact that all the consequences are not




Fig. 3 Interpretation of Axiom P4.

judged as equivalent by the agent (otherwise, it would be impossible to discriminate between acts and SEU would be useless to help the agent in her decision making process):

Axiom 8 (P5: Non-triviality of preferences over the consequences) *There exist two outcomes $x, y \in \mathcal{X}$ such that $\delta_x \succ \delta_y$.*

The five above axioms seem rather reasonable and do not seem too restrictive in the sense that they tend to win unanimous support from people. Yet, as Savage showed, their combination necessarily induces that the agent models uncertainties using a qualitative probability⁴. To establish the existence of a subjective probability, an additional axiom is needed, that is closely related to the continuity axiom of von Neumann-Morgenstern: assuming that E_i is a highly unlikely event, if f and g are two acts such that $f \succ g$ and if x is an arbitrary consequence, then $\delta_x E_i f$ should be very close to f and, therefore, as $f \succ g$, the agent should also prefer $\delta_x E_i f$ to g . For the same reason, she should also prefer f to $\delta_x E_i g$:

Axiom 9 (P6: Continuity) *For all acts $f, g \in \mathcal{X}^{\mathcal{S}}$ such that $f \succ g$ and for all $x \in \mathcal{X}$, there exists a finite partition $\{E_1, \dots, E_n\}$ of \mathcal{S} such that $\delta_x E_i f \succ g$ and $f \succ \delta_x E_i g$ for every $i \in \{1, \dots, n\}$.*

Adding Axiom P6 to the five other axioms necessarily induces the existence of a subjective probability distribution. In addition, all these axioms induce that the agent is an expected utility maximizer, as shown in the following theorem:

Theorem 2 (Savage, 1954⁵). *If the preferences of an agent satisfy axioms P1 to P6, then preference relation \succsim over the set of acts with finite support is representable by a utility function $U(f) = \sum_{s \in \mathcal{S}} p(s)u(f(s))$, where $p(s)$ is the subjective probability of the agent over the state of nature s . In addition, u , the utility function over the set of consequences, is unique up to scale and location.*

Savage has also extended this theorem, notably to the case in which acts are only constrained to be bounded [Savage, 1954]. Note that there also exist other axiomatics of the EU criterion under uncertainty, notably that of [Anscombe and Aumann, 1963]. All these axiomatics have however in common to rely on axioms that are easily justifiable and that, to some extent, reflect a logical reasoning. In this sense, they constitute the foundation of a rational behavior. From all these axiomatics, it could be easily inferred that only probabilities can “rationally” model uncertainties. This assertion has also been supported for a long time by what decision theorists call “*Dutch books*”, which are situations in which using a model of uncertainties different from probabilities inevitably leads the agent to loose some money. As an example, let us consider a bookmaker proposing bets on the three horses of a race.

⁴ Note that qualitative probabilities are slightly different from probabilities, see [Kraft et al., 1959] for a proof of this assertion.

⁵ Savage’s theorem is somewhat more general than the theorem mentioned here: acts need not have a finite support, it is sufficient that the set of consequences \mathcal{X} is finite. In this case, the summation needs to be substituted by an integral w.r.t. the subjective probability measure.

He offers the odds shown in Table 2. Note that the sum of the induced “probabilities” estimated by the bookmaker is equal to 0.95, not to 1. This deviation from a valid probability distribution implies that there is a possibility for gamblers to always win money from the bookmaker. Indeed, gamblers betting the amounts of money shown in the fourth column of the table are guaranteed to win 200 € even though they bet only 190 €. This type of money pump argument has also significantly contributed to establish probabilities as the only reasonable representation of uncertainties in a decision making context.

horse	odds	induced proba.	bet price	reimbursement
1	1 against 1	$\frac{1}{1+1} = 0.5$	100 €	100 € of bet + 100 € = 200 €
2	3 against 1	$\frac{1}{3+1} = 0.25$	50 €	50 € of bet + 150 € = 200 €
3	4 against 1	$\frac{1}{4+1} = 0.2$	40 €	40 € of bet + 160 € = 200 €

Table 2 Example of a Dutch book.

In the two axiomatics above, that of von Neumann-Morgenstern and that of Savage, the von Neumann-Morgenstern utility function, i.e., the utility representing the agent’s preferences over the consequences, is unique up to scale and location. But in decision under certainty, i.e., when the consequences of each action are known with certainty, utility functions (over outcomes) are unique only up to strictly positive increasing transforms. Consequently, we can deduce that von Neumann-Morgenstern utilities must implicitly include some factor related to uncertainties. We will see now that, in reality, this factor represents the attitude of the agent w.r.t. risk.

2.2 Risk Measures

Before defining formally the agent’s attitude w.r.t. risk, we need to define the concept of *risk*, and especially how the *quantity* of risk involved in a decision can be measured. A decision can be summarized by an act or a lottery $\langle x_1, p_1; \dots, x_n, p_n \rangle$. In a sense, the latter correspond to a random variable x whose domain is x_1, \dots, x_n and the usual risk measure of a real-valued random variable is its variance. So it is tempting to exploit variance as the measure of risk involved in a decision. This idea is supported by the celebrated Arrow-Pratt formula for approximating utility functions, which contains a component related to variance [Pratt, 1964; Arrow, 1965]. But as shows the following example in [Ingersoll, 1987], this measure is not very well suited: let $L_1 = \langle 0, 0.5; 4, 0.5 \rangle$ and $L_2 = \langle 1, 7/8; 9, 1/8 \rangle$ be two lotteries. Intuitively, observing L_1 and L_2 , lottery L_1 seems more risky than L_2 since its consequences are equiprobable whereas, in L_2 , it is very likely that the decision yields consequence 1. Unfortunately, the variances of both lotteries are equal.

In decision theory, the most commonly used risk measure is due to [Rotschild and Stiglitz, 1970, 1971]. It is much more robust than variance. It relies on the concept of “*mean-preserving risk increase*” or, as stated usually, “*Mean Preserving Spread*” (MPS). Let us consider the three lotteries P, Q, R of Table 3. Observe the only difference between P and Q : Lottery P yields consequence 4 with probability 0.3 whereas, Q yields consequences 3 and 5 with probability 0.15 (hence, globally, a probability of 0.3 to get consequence “3 or 5”). As a result, Q can be judged as more risky than P since, with a probability of 0.3, the consequence yielded by P is known (i.e., 4) whereas, in Q , with the same probability, we only know that 3 or 5 will be yielded, and there still exists a lottery $\langle 3, 0.5 ; 5, 0.5 \rangle$ to determine which consequence will be yielded. Remark that the expectations of random variables X and Y of Table 3 are equal. This explains why Y is said to be a mean-preserving (same expectation as X) risk increase (w.r.t. X) or, for short, a MPS of X . Similarly, Z is a MPS of Y because their expectations are equal and Y yields consequence 16 with probability 0.21 whereas Z induces lottery $\langle 12, 0.07 ; 18, 0.14 \rangle$ instead.

X	$P(X)$	Y	$Q(Y)$	Z	$R(Z)$
-2	0.09	-2	0.09	-2	0.09
4	0.30	3	0.15	3	0.15
		5	0.15	5	0.15
10	0.40	10	0.40	10	0.40
16	0.21	16	0.21	12	0.07
				18	0.14

Table 3 Mean preserving spread: $Y = \text{MPS}(X)$, $Z = \text{MPS}(Y)$ and $Z = \text{MPS}(X)$.

In the rest of this subsection, we will consider that \mathcal{X} is equal to \mathbb{R} and, more generally, that it is a monetary space (this will make the interpretations of the results easier to understand).

Definition 1 (Mean preserving spread). Let X and Y be two real-valued random variables. Y is said to be a *Mean Preserving Spread* of X if and only if there exists a white noise Θ , i.e., a random variable whose expectation is equal to 0, such that $Y = X + \Theta$.

Let us call F_X and F_Y the cumulative distribution functions (CDF) of random variables X and Y respectively. In other words, if P_X is the probability distribution of X , then $F_X(x) = P_X(z : z \leq x)$ for every $x \in X$. Figure 4 displays the CDFs of variables X and Z of Table 3. When $X, Z < 3$, the two CDFs are identical. Then, when $x \in [3, 4[$, we have that $F_Z(x) > F_X(x)$. Therefore, we also have that $\int_{x < 4} F_Z(x) dx > \int_{x < 4} F_X(x) dx$. When $x \in [4, 5[$, the difference $F_X(x) - F_Z(x)$ is positive, so the gap between the two integrals decreases but the two gray regions on the left of Figure 4 have the same area so, overall, the integral of F_Z is always greater than or equal to that of F_X . This property is general and provides an alternative characterization of MPS:

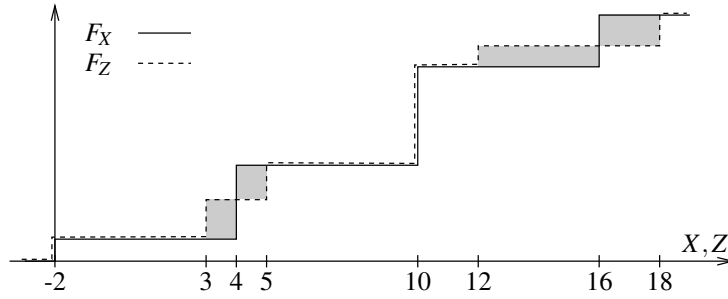


Fig. 4 Interpretation of MPS in terms of cumulative distributions.

Definition 2 (Mean preserving spread). Let X and Y be two real-valued random variables. Y is said to be a *Mean Preserving Spread* of X if and only if i) X and Y have the same expectations; and ii) X and Y satisfy the following equation:

$$\int_{-\infty}^T F_Y(x) dx \geq \int_{-\infty}^T F_X(x) dx \quad \text{for every } T \in \mathbb{R}. \quad (2)$$

Definition 3 (2nd order stochastic dominance). Let X and Y be two real-valued random variables. X dominates stochastically Y at the second order if and only if Equation (2) is satisfied.

Rothschild and Stiglitz proved that Definitions 1 and 2 are equivalent. They also provided a characterization in terms of risk aversion, as we will define it in the next subsection: Assertion 3 of the theorem below expresses the fact that Y is a MPS of X if and only if any weakly risk averse agent prefers X to Y .

Theorem 3 (Rothschild and Stiglitz, 1970). Let X and Y be two real-valued random variables with the same expectation. The following three assertions are equivalent:

1. $Y = \text{MPS}(X)$ (in the sense of Definition 2);
2. Y has the same distribution as $X + \Theta$, where Θ is a white noise;
3. for any increasing and concave function $u : \mathbb{R} \mapsto \mathbb{R}$, we have that $\int u(x) dF_X(x) \geq \int u(x) dF_Y(x)$.

We can now characterize the behavior of agents w.r.t. lotteries with different amounts of risk. Of special interest, we can now determine if the agent would prefer “taking risks” or not.

2.3 Attitude of Agents with respect to Risk

The simplest way to estimate whether an agent is risk seeking or risk averse consists of asking her which lottery she would prefer among one lottery X without any

risk (it can yield only one consequence, known for sure) and another lottery Y with the same expectation but containing some risk (the lottery can yield several consequences). Note that, as both lotteries have the same expectation, $Y = \text{MPS}(X)$. Assume now that the agent's von Neumann-Morgenstern utility is linear ($u(x) = x$ for simplicity). Then the expected utility of the lottery corresponding to Y is equal to the expectation of Y which, by definition, is equal to that of X and, also, to the expected utility of the lottery associated to X . An agent who is expected utility maximizer shall therefore be indifferent between X and Y . For instance, for the agent, $\langle \frac{x_1+x_2}{2}, 1 \rangle \sim \langle x_1, \frac{1}{2}; x_2, \frac{1}{2} \rangle$. These two lotteries have the same expectation (this is the reason why the agent is indifferent between them), but the first one is not risky while the second one is. So we can conclude that the preferences of the agent do not take into account the amount of risk involved in the lotteries. The agent is thus said to be “*risk neutral*”. Of course, if the agent had strictly preferred X to Y , we would say that she has some aversion w.r.t. risk and, therefore, he would be “*risk averse*”. Finally, if the agent had strictly preferred Y to X , she would be said to be “*risk seeking*”. Arrow and Pratt propose the following definition [Pratt, 1964; Arrow, 1965]:

Definition 4 (Weak risk attitudes). An agent is weakly risk averse if, for every real-valued random variable X , she prefers $E(X)$ to random variable X itself: $\langle E(X), 1 \rangle \succ X$. An agent is weakly risk neutral (resp. seeking) if $\langle E(X), 1 \rangle \sim X$ (resp. $X \succ \langle E(X), 1 \rangle$).

We have seen above that a linear von Neumann-Morgenstern utility implies that the agent is risk neutral. Arrow and Pratt have shown that, more generally, the agent's risk attitude is characterized by the concavity or convexity of the von Neumann-Morgenstern utility function:

Theorem 4. *An agent is (weakly) risk averse if and only if her von Neumann-Morgenstern utility function u is concave. She is (weakly) risk neutral if and only if u is linear. Finally, she is (weakly) risk seeking if and only if u is convex.*

Up to now, the risk attitude of the agent was characterized by comparing one risky lottery with a lottery involving no risk. It could be objected that such a comparison is extreme and could introduce some biases. So it might be more appropriate to compare only lotteries involving some risk, some being more risky than others. The concept of mean preserving spread allows to specify such lotteries: it is sufficient to compare lotteries X and Y such that one of them is an MPS of the other. In this case, an agent is risk averse if and only if she prefers lottery X to any $\text{MPS}(X)$:

Definition 5 (Strong risk attitudes). An agent is strongly risk averse if, for every real-valued random variable X , she prefers lottery X to any lottery Y such that $Y = \text{MPS}(X)$. An agent is strongly risk neutral (resp. seeking) if $X \sim Y$ (resp. $Y \succ X$).

Of course, by definition, strong risk aversion implies weak risk aversion. But in the EU model, the converse is also true:

Theorem 5 (Rotschild and Stiglitz, 1970). *In the EU model, the following three assertions are equivalent:*

1. the agent is weakly risk averse;
2. the agent is strongly risk averse;
3. the agent's von Neumann-Morgenstern utility is concave.

As the concavity of the von Neumann-Morgenstern utility function u characterizes the agent's aversion w.r.t. risk, it seems natural to define the intensity of this aversion in terms of properties of u . Arrow and Pratt have proposed to characterize it in terms of a *coefficient of absolute risk aversion*: assume that u is strictly increasing and twice continuously differentiable, with a strictly positive derivative. The coefficient of absolute risk aversion is defined as function $R_A : \mathbb{R} \mapsto \mathbb{R}$ such that $R_A(x) = -u''(x)/u'(x)$.

This definition can be easily interpreted by considering a risk averse agent. Assume that the set of consequences \mathcal{X} is a monetary space. A *common* agent prefers in general to win more money than less, so her utility $u(x)$ strictly increases with x and, consequently, $u'(x) > 0$. In addition, being risk averse, $u(x)$ is concave, hence $u''(x) < 0$. From these properties, it can be deduced that $R_A(x) > 0$. Consider now utility function $u_1(x) = \ln x$, which implies coefficient $R_A^1(x) = 1/x$. In Figure 5, it can be observed that the concavity rate of u_1 decreases with x . This translates in terms of coefficient of absolute risk aversion into a decreasing coefficient R_A^1 . The level of aversion w.r.t. risk therefore varies with x and, in practice, it is generally strictly decreasing. As a matter of fact, a poor agent is not often prone to take the risk of loosing some money in order to gain more money whereas a wealthy agent is inclined to take such a risk because the same loss of money seems to her relatively much less important than to the poor agent.

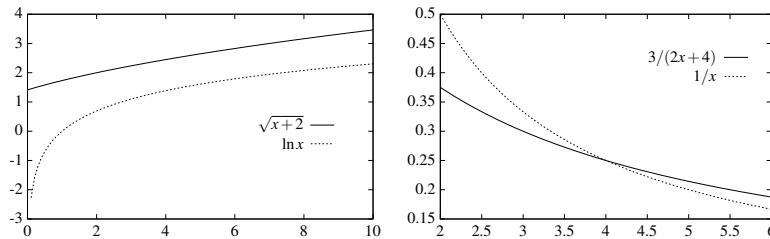


Fig. 5 Coefficients of absolute risk aversion.

Note that R_A can also be exploited to compare the aversions among several agents. Indeed, consider now two utility functions $u_1(x) = \ln x$ and $u_2(x) = \sqrt{x+2}$. These functions induce two coefficients $R_A^1(x) = 1/x$ and $R_A^2(x) = 3/(2x+4)$. Figure 5 displays functions u_1, u_2 as well as their respective coefficients of aversion. From this figure, it can be remarked that the second agent (u_2) is more risk averse for small amounts of money whereas this trend is inverted for larger amounts. Note that such a comparison is meaningful because von Neumann-Morgenstern utilities

being unique up to scale and location, R_A remains invariant w.r.t. affine transforms of u .

Clearly, the EU model presents very nice properties. As we have seen, it is justifiable from the viewpoint of the agent's rationality. In addition, its linearity allows for very efficient algorithms, notably in the context of sequential decision making and in that of preference elicitation [Keeney and Raiffa, 1993; Chajewska et al., 2000; Boutilier, 2002; Wang and Boutilier, 2003]. However, during the last decades, several criticisms were raised against this model, which led to alternative decision models. The next section shows some of the most important criticisms.

2.4 Some Descriptive Limits of the EU Model

Among the first detractors of the EU model, Allais proposed a celebrated example known as the “*Allais paradox*” [Allais, 1953], about which experimental studies have shown that the majority of the surveyed agents have preferences that violate the independence axiom (Axiom 3) and are, therefore, not representable in the EU model. Actually, consider the following two lotteries:

- $L_1 = \langle \text{win } 1\text{M } \text{€}, 1 \rangle$;
- $L_2 = \langle \text{win } 1\text{M } \text{€}, 0.89 ; 5\text{M } \text{€}, 0.1 ; 0 \text{€}, 0.01 \rangle$.

Most of the surveyed agents prefer L_1 to L_2 because the uncertainty contained in L_2 is not counterbalanced by the potential gain of $5\text{M } \text{€}$. When faced to the following alternatives:

- $L'_1 = \langle \text{win } 1\text{M } \text{€}, 0.11 ; 0 \text{€}, 0.89 \rangle$,
- $L'_2 = \langle \text{win } 5\text{M } \text{€}, 0.10 ; 0 \text{€}, 0.90 \rangle$,

the same agents usually prefer L'_2 to L'_1 because the difference in probability between 0.11 and 0.10 is judged as relatively low and the agents therefore base essentially their preferences on the potential gains of the lotteries. But, if we set: $P = \langle 1\text{M } \text{€}, 1 \rangle$, $Q = \langle 5\text{M } \text{€}, 10/11 ; 0 \text{€}, 1/11 \rangle$, $R = \langle 1\text{M } \text{€}, 1 \rangle$ and $S = \langle 0 \text{€}, 1 \rangle$, then:

$$\begin{array}{ll} L_1 = 0.11P + 0.89R & L_2 = 0.11Q + 0.89R \\ L'_1 = 0.11P + 0.89S & L'_2 = 0.11Q + 0.89S. \end{array}$$

Therefore, according to the independence axiom, if $L_1 \succ L_2$, the agent shall also have the following preference: $L'_1 \succ L'_2$. Obviously, this is not observed experimentally. This example is quite unsettling because this preference reversal can be explained easily and does not seem to result from some irrational behavior. As we will see in the next section, this example has led researchers to develop new decision models based on different rationality criteria. These models have a higher descriptive power than the EU model and are notably capable of explaining why people tend to prefer L_1 to L_2 and L'_2 to L'_1 . Other experimental studies, in particular [Kahneman and Tversky, 1972, 1979], highlight other biases w.r.t. the predictions made by the EU model. This is the case, for instance, of the certainty effects.

The second criticism addressed against the EU model concerns the interpretation of the concavity of the von Neumann-Morgenstern utility function u . Indeed, we have seen that in this model a concave utility represents an aversion w.r.t. risk. But u represents the agent's preferences over the space of the consequences and, in general, agents have decreasing marginal preferences over money, i.e., the amount of increase of the agent's satisfaction (as measured by the utility function) tends to decrease when the amounts of money tend to rise. Thus, the satisfaction to increase the agent's wealth from 10 € to 20 € is higher than that to increase it from 10010 € to 10020 €. In terms of preferences, this decrease necessarily induces the concavity of u . This double interpretation of u 's concavity implies that the EU model is unable to describe the behavior of agents that are at the same time risk averse and that have decreasing marginal preferences.

The third main criticism against the EU model lies in its lack of flexibility to model different types of risk aversions. Indeed, in EU, it is impossible to model an agent who is weakly but not strongly risk averse. But this kind of agent can exist and, more generally, there exist several notions of risk aversion that are not necessarily all equivalent [Chateauneuf et al., 2004]. We will see in the next section some "new" decision models that can cope with this lack of flexibility.

The set of criticisms presented here cannot be exhaustive due to lack of space. However, we shall mention two important additional criticisms. First, the formula of the expected utility model combines through multiplications the probabilities of occurrence of the consequences with the utilities. As a consequence, EU necessarily requires the commensurability of preferences and uncertainties: one can "trade" uncertainty for preference satisfaction. For instance, if $\langle x_1, 0.5 ; x_2, 0.5 \rangle \sim \langle x_3, 1 \rangle$, the agent is willing to trade/discard some uncertainty (0.5) for a change in consequences (winning x_3 instead of x_1 or x_2 , hence a modification in her satisfaction). In addition, even though commensurability may be a reasonable assumption in some practical applications, is it always sensible to model uncertainties by probabilities? According to Savage, this is the only rational representation. However, when considering the example of the Ellsberg's urn [Ellsberg, 1961], this justification seems far from being convincing: consider an urn containing red, yellow and black balls. The only information available about these balls is that one third are red and the two remaining third are either yellow or black (but we do not know their respective proportions). With so few information available, it seems difficult for a "rational" agent to estimate the underlying probability distribution over the colors of the balls, and experimental studies highlight this fact. When agents are invited to determine the alternative they prefer among the following ones, whose outcome depends on the color of a ball drawn randomly from the urn:

- Alternative A : win 1M € if the ball is red, else 0 €,
- Alternative B : win 1M € if the ball is black, else 0 €,

most of the agents prefer A to B because, potentially, the urn contains no black ball whereas the urn is guaranteed to contains 1/3 of red balls. On the other hand, when facing the following alternatives:

- Alternative C : win 1M € if the ball is red or yellow, else 0 €,

- Alternative D : win $1M \text{€}$ if the ball is black or yellow, else 0€ ,

the agents prefer in general alternative D to C . But this kind of behavior is incompatible with the EU model because it violates the Sure Thing Principle. Indeed, if E represents the event “the drawn ball is red or black”, if a_1 and a_2 represent the acts yielding “ $1M \text{€}$ if red ball, else 0€ ” and “ $1M \text{€}$ if black ball, else 0€ ”, and if δ_h and δ_k represent the “constant” acts yielding with certainty 0€ and $1M \text{€}$ respectively, then alternatives A and B can be represented by acts $a_1 E \delta_h$ and $a_2 E \delta_h$ respectively, whereas alternatives C and D correspond to acts $a_1 E \delta_k$ and $a_2 E \delta_k$ respectively. According to the Sure Thing Principle, one of the fundamental principles underlying EU, $A \succ B$ should imply $C \succ D$, which is not the case observed experimentally.

All these descriptive limits have led researchers to propose new models, also relying on rationality criteria, but with a higher expressive power. We will now briefly describe some of them.

3 Non-linear Models for Decision under Risk

The descriptive limits mentioned above first led decision making researchers to propose models quite similar to EU but, still, weakening one or several axioms of von Neumann-Morgenstern (or of Savage). Let us cite for instance the model proposed in [Machina, 1982] which discards the independence axiom but is still locally coherent with EU. There also exist models based on security levels like, e.g., that of [Jaffray, 1988] in which the independence axiom is defined only on pairs of probability distributions that share the same worst consequence.

However, these models have been replaced by what decision theorists call “new” models, which are generalizations of EU. Among the first new models proposed, “*Prospect Theory*” consists of deforming probabilities using an increasing transform [Kahneman and Tversky, 1979] in order not to take into account the true probabilities but rather the way agents perceive these probabilities. Although seminal, this model is not used anymore, essentially because it could sometimes propose to the agent to make dominated decisions, i.e., to choose an alternative D_1 such that there existed another alternative D_2 such that, whatever the state of nature that could occur, the consequence yielded by D_2 was judged at least as good as that yielded by D_1 (and it was judged strictly better for at least one state of nature). This feature being very difficult to justify from a rationality point of view, the model is not used anymore. However, it paved the way for the new models, notably for “*Rank Dependent Utility*” (RDU), that we will now describe [Quiggin, 1982, 1993].

Let x_1, x_2, x_3 be three consequences. Without loss of generality, let us assume that $u(x_2) < u(x_1) < u(x_3)$. According to the EU model, lottery $L = \langle x_1, p_1; x_2, p_2; x_3, p_3 \rangle$ is evaluated as $EU(L) = p_1 u(x_1) + p_2 u(x_2) + p_3 u(x_3)$. It is easy to show that this expression is equivalent to:

$$EU(L) = (p_1 + p_2 + p_3)u(x_2) + (p_1 + p_3)[u(x_1) - u(x_2)] + p_3[u(x_3) - u(x_1)]. \quad (3)$$

This new expression can be interpreted as follows: at worst, the agent is guaranteed with probability $p_1 + p_2 + p_3 = 1$ to win consequence x_2 . Then, the probability that she gets a consequence strictly better than x_2 , i.e., at least as good as consequence x_1 is $p_1 + p_3$. Finally, the probability to win something better than x_1 , i.e., x_3 , is p_3 . The key idea of RDU is to combine this expression with the probability transformation principle of the Prospect Theory. Thus, in its decision making process, RDU does not take into account the true probabilities but only their perceptions by the agent. The score assigned to L by RDU is therefore:

$$RDU(L) = \varphi(p_1 + p_2 + p_3)u(x_2) + \varphi(p_1 + p_3)[u(x_1) - u(x_2)] + \varphi(p_3)[u(x_3) - u(x_1)], \quad (4)$$

where φ is an increasing function from $[0, 1]$ to $[0, 1]$. Experimental studies by Kahneman and Tversky have shown that this function is, in general, similar to that of Figure 6, whose equation is $\varphi(x) = e^{-\sqrt{-\ln(x)}}$.

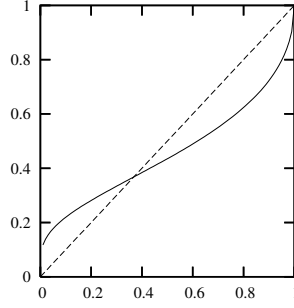


Fig. 6 Probability transformation function.

Definition 6 (Rank Dependent Utility (RDU)). An agent behaves according to the RDU model if her preference relation over the set of lotteries \mathcal{L} is representable by two functions u and φ , where u is the von Neumann-Morgenstern utility over the set of consequences and $\varphi : [0, 1] \mapsto [0, 1]$ is an increasing function such that $\varphi(0) = 0$ and $\varphi(1) = 1$. The agent assigns to every lottery $L = \langle x_1, p_1 ; \dots, x_n, p_n \rangle$ such that $u(x_1) \leq u(x_2) \leq \dots \leq u(x_n)$ utility:

$$RDU(L) = u(x_1) + \sum_{i=2}^n \left[\varphi \left(\sum_{k=i}^n p(x_k) \right) [u(x_i) - u(x_{i-1})] \right]. \quad (5)$$

As an example, if $u(x) = x/2$ and $\varphi(x) = x^2$, then, to compute the RDU value of lottery $L = \langle 3, 0.2 ; 10, 0.4 ; 5, 0.1 ; 9, 0.3 \rangle$, consequences must first be sorted in increasing utility order: $L = \langle 3, 0.2 ; 5, 0.1 ; 9, 0.3 ; 10, 0.4 \rangle$. Then, the application of Equation (5) yields:

$$RDU(L) = \varphi(1) \times \frac{3}{2} + \varphi(0.8) \times \left[\frac{5}{2} - \frac{3}{2} \right] + \varphi(0.7) \times \left[\frac{9}{2} - \frac{5}{2} \right] + \varphi(0.4) \times \left[\frac{10}{2} - \frac{9}{2} \right].$$

There exist alternative definitions of RDU. Let us show one of them that will prove useful for highlighting the connection between RDU and another more general model: Choquet expected utility.

Definition 7 (Rank dependent utility (RDU)). Let u and φ be the functions defined in Definition 6. Let X be a random variable whose probability distribution is P . Then:

$$RDU(X) = \int_{-\infty}^0 [\varphi(P(u(X) > t)) - 1] dt + \int_0^{\infty} \varphi(P(u(X) > t)) dt.$$

Note that the Allais paradox can be explained by RDU. This is notably the case when utility u is linear and the probability transform φ is like the one suggested by Kahneman and Tversky: $\varphi(x) = e^{-\sqrt{-\ln(x)}}$. The expressive power of RDU is therefore higher than that of EU. It generalizes the latter since, when $\varphi(x) = x$, RDU boils down to EU. Note also that, when $\varphi(p) \leq p$ for every p , the agent always underestimate the probabilities of the utility increases $u(x_i) - u(x_{i-1})$ (see Equations (3) and (4)). This can be interpreted as a kind of pessimism under risk (since the agent takes more into account the worst consequences than the best ones).

The axiomatic foundations of RDU are quite complicated [Quiggin, 1982; Wakker, 1994; Chateauneuf, 1999], so in this chapter, we will not detail them. However, to let the reader understand the key feature of RDU, we will now focus on RDU's main properties: the comonotonic independence axiom in von Neumann-Morgenstern's framework and the comonotonic sure thing principle in Savage's framework [Chew and Wakker, 1996]. Here, we chose to present only the latter because it is somewhat simpler to understand than the former. For this purpose, we need to define "comonotonic acts": two acts f and g are said to be comonotonic if there exists no pair of states of nature $s, s' \in \mathcal{S}$ such that $f(s) \succ_{\mathcal{X}} f(s')$ and $g(s) \prec_{\mathcal{X}} g(s')$. Intuitively, two acts are comonotonic if their variations (in terms of preferences over the consequences) do not evolve in the opposite directions when moving from one state of nature to another. For instance, in Figure 7, in which preferences over the consequences increase along the vertical axis, f and g are comonotonic, as well as g and k , and h and k . But g and h are not comonotonic because $g(s_3) \succ_{\mathcal{X}} g(s_2)$ and $h(s_2) \succ_{\mathcal{X}} h(s_3)$. Note that comonotonicity is not a transitive property since g and k are comonotonic, as well as k and h , but g and h are not comonotonic. The key idea of RDU consists of imposing the "Sure Thing Principle" only over comonotonic acts:

Axiom 10 (comonotonic sure thing principle) Let $\{A_1, \dots, A_n\}$ be a partition of \mathcal{S} and let $f : A_i \mapsto x_i$ and $g : A_i \mapsto y_i$ be two acts such that $x_1 \leq x_2 \leq \dots \leq x_n$ and $y_1 \leq y_2 \leq \dots \leq y_n$. Assume that there exists $i_0 \in \{1, \dots, n\}$ such that $x_{i_0} = y_{i_0}$. Let $f' : A_i \mapsto x'_i$ and $g' : A_i \mapsto y'_i$ be two other acts such that:

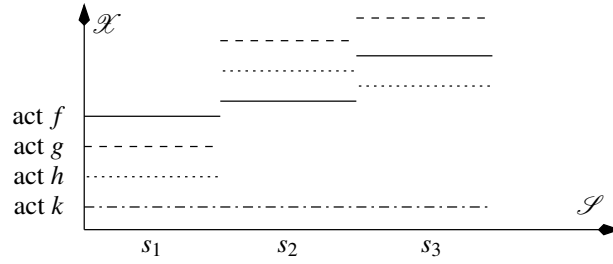


Fig. 7 Illustration of comonotonicity.

$$\begin{cases} x'_{i_0} = y'_{i_0}; & \text{and } x'_i = x_i \text{ and } y'_i = y_i \text{ for every } i \neq i_0, \\ x'_1 \leq \dots \leq x'_n \text{ and } y'_1 \leq \dots \leq y'_n. \end{cases}$$

Then $f \succsim g \implies f' \succsim g'$.

This principle is illustrated in Figure 8: The common part of acts f and g can vary only between points A and B . Thus, acts f' and g' satisfy the constraints of the above definition, which is not the case for acts f'' and g'' . Table 4 shows the acts corresponding to the Allais paradox mentioned in the preceding section. In this table, the A_i 's are sorted in such a way that acts L_1 and L_2 correspond to f and g of Axiom 10. It can be seen that quadruple (L_1, L_2, L'_1, L'_2) does not satisfy the premises of Axiom 10 (see the difference between L_1 and L'_1). As a consequence, the Allais paradox does not violate the comonotonic sure thing principle. This is the reason why RDU can explain why agents prefer L_1 to L_2 and L'_2 to L'_1 .

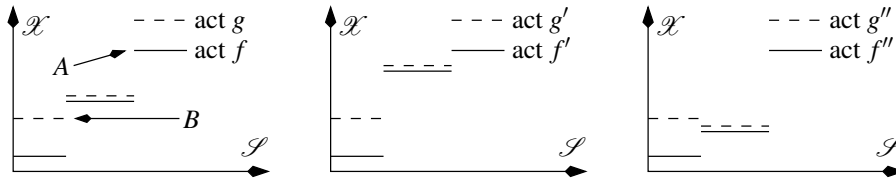


Fig. 8 The comonotonic sure thing principle.

The RDU model is in fact a particular case of a more general model: Choquet expected utility (CEU), that we will briefly describe after introducing the concept of capacity:

Definition 8 (Capacity). A capacity $\mu : 2^{\mathcal{S}} \mapsto [0, 1]$, where \mathcal{S} is the set of states of nature, is a function satisfying the following two properties:

act	$A_1 (P(A_1) = 0.01)$	$A_2 (P(A_2) = 0.89)$	$A_3 (P(A_3) = 0.10)$
L_1	1M€	1M€	1M€
L_2	0€	1M€	5M€
L'_1	1M€	0€	1M€
L'_2	0€	0€	5M€

Table 4 The Allais paradox and the comonotonic acts.

1. $\mu(\emptyset) = 0$ and $\mu(\mathcal{S}) = 1$;
2. For every pair $A, B \subseteq \mathcal{S}$, we have that $A \subseteq B \implies \mu(A) \leq \mu(B)$.

Here, a capacity must be understood as a generalization of the concept of probability distribution⁶. Indeed, any probability distribution satisfies properties 1) and 2) above. This is also the case for all the probability transforms of the RDU model. Therefore, capacities allow to define a more general decision model:

Definition 9 (Choquet expected utility (CEU)). An agent behaves according to the CEU model if her preference relation over the set of acts $\mathcal{X}^{\mathcal{S}}$ is representable using two functions u and μ , where u is the utility function over the consequences and $\mu : 2^{\mathcal{S}} \mapsto [0, 1]$ is a capacity. The agent assigns to each act f utility:

$$CEU(f) = \int_{C_h} u(f) d\mu = \int_{-\infty}^0 [\mu(u(f) > t) - 1] dt + \int_0^{\infty} \mu(u(f) > t) dt. \quad (6)$$

It has been proved in [Wakker, 1990] that CEU reduces to RDU when Axiom 11 below is added to the axiomatics of CEU [Schmeidler, 1986; Gilboa, 1987; Wakker, 1990]. It is generally believed that this axiom is attractive for a “rational” decision model since it expresses the fact that if, for every consequence x , the probability of winning at least x is higher with act f than with act g , the agent should prefer f to g .

Definition 10 (First order stochastic dominance). For every act h , let $F_h(x) = P(\{s \in \mathcal{S} : h(s) \leq x\})$ denote the cumulative distribution of h . Let f and g be two acts and let F_f and F_g be their respective cumulative distributions. Then f stochastically dominates g at the first order if, for every $x \in \mathbb{R}$, we have that $F_f(x) \leq F_g(x)$.

Axiom 11 (First order stochastic dominance) *Let f and g be two acts. If f stochastically dominates g at the first order, then $f \succsim g$.*

We will see again the CEU model and its usefulness for decision making under uncertainty in the next section. To complete our overview of RDU, we must mention some results about risk aversion. We have seen earlier that, in the EU model, strong risk aversion is equivalent to weak risk aversion, which also corresponds to the concavity of the von Neumann-Morgenstern utility u . Is this also the case in RDU? A

⁶ For an interpretation in terms of weights of agents’ coalitions or of criteria, see Chapter 16 of this volume.

first answer to this question can be found in [Chew et al., 1987], where it is proved that a RDU agent is strongly risk averse if and only if her utility u is concave and her probability transform φ is convex. Similarly, the agent is strongly risk seeking if and only if u is convex and φ is concave. To our knowledge, there does not exist yet any complete characterization of weak risk aversion in the RDU model. Only sufficient conditions have been proposed and those do not require the concavity of u [Chateauneuf and Cohen, 1994]. In terms of risk aversion, the expressive power of RDU is therefore higher than that of EU. Finally, note that other concepts of risk aversion designed specifically for RDU have been proposed. Those are different from both strong and weak risk aversions. For instance, Quiggin suggested to substitute strong risk aversion by monotonic risk aversion [Quiggin, 1992]: let X and Y be two random variables. Y is said to be a monotonic mean preserving spread (MMPS) of X if $Y = X + Z$, where Z is a white noise, and X and Z are comonotonic. An agent is monotonic risk averse if she does not like monotonic risk increase, i.e., if $Y = \text{MMPS}(X)$, then $X \succsim Y$.

Up to now, we have only studied decision models relying on the existence of probability distributions to model uncertainties. But, what can we do if there does not exist sufficient information to construct one, like in the Ellsberg's urn example? The goal of the next section is to provide some keys to answer to this question.

4 Decision Models outside the Probabilistic Framework

Let us recall the Ellsberg's urn problem: this is an urn containing 99 balls, which can be either red, yellow or black. The only information available to the agent is that one third of the balls is red and the remaining two third are either yellow or black (but their respective proportions are unknown). Agents bet on the color of a ball to be drawn from the urn. Thus, an agent is asked which alternative she prefers between alternatives A and B below, and which one she prefers between C and D :

- Alternative A : win $1M\text{€}$ if the drawn ball is red, else win 0€ ,
- Alternative B : win $1M\text{€}$ if the drawn ball is black, else win 0€ ,
- Alternative C : win $1M\text{€}$ if the drawn ball is red or yellow, else win 0€ ,
- Alternative D : win $1M\text{€}$ if the drawn ball is black or yellow, else win 0€ .

Most of the human agents prefer A to B and D to C . As we have seen before, EU cannot account for such preferences (violation of the sure thing principle). RDU can neither model these preferences. Indeed, if it could then, assuming that the agent prefers winning more money than less, and denoting by P_r, P_y, P_b the probabilities that the drawn ball is red, yellow and black respectively, we have that $A \succ B \iff \text{RDU}(A) > \text{RDU}(B) \iff \varphi(P_r) > \varphi(P_b)$ and $D \succ C \iff \varphi(P_b + P_y) > \varphi(P_r + P_y)$. But this is impossible to have both inequalities satisfied because φ is an increasing function. Here, the problem is that there does not exist a unique probability distribution compatible with the information available to the agent. Therefore, in this case, we should not try to use a decision model that relies on a unique probabil-

ity distribution but rather on a model that relies on the set of all the distributions compatible with the available information. Here, it is easy to see that this set is convex: if P and Q are two compatible probability distributions, for every $\alpha \in [0, 1]$, we have that $\alpha P + (1 - \alpha)Q$ is also compatible with the available information. As a consequence, to represent the uncertainties in the Ellsberg's urn, it is sufficient to know the boundary of the convex hull of all the compatible distributions. But since the probability of any event and that of its complementary event sum always to 1, the lower bounds on the probabilities are sufficient to characterize all the convex hull. Those correspond to a function $\mu : 2^{\mathcal{S}} \mapsto [0, 1]$ such that, for every $A \subseteq \mathcal{S}$, $\mu(A) = \min_{\{P \text{ compatibles}\}} P(A)$. For the Ellsberg's urn, this function μ , also called a "belief function", is described in Table 5. Indeed, the min of $P(Y)$ is equal to 0 because it is possible that the urn contains no yellow ball. On the other hand, $\min P(Y, B) = 2/3$ because, for all the probability distributions P compatible with the Ellsberg's urn, we have $P(Y, B) = 2/3$. More formally, belief functions [Dempster, 1967; Shafer, 1976] are defined as follows (see Chapter 4 of this volume):

Definition 11 (Belief function). A belief function μ is a capacity (in the sense of Choquet) which is ∞ -monotone, i.e., it is such that for all $n \geq 2$, and for all $A_1, \dots, A_n \in 2^{\mathcal{S}}$:

$$\mu\left(\bigcup_{i=1}^n A_i\right) \geq \sum_{\emptyset \subset I \subseteq \{1, \dots, n\}} (-1)^{|I|+1} \mu\left(\bigcap_{i \in I} A_i\right).$$

To any capacity (and *a fortiori* to any belief function) is associated its *Möbius inverse* ϕ defined by: $\phi(A) = \sum_{B \subseteq A} (-1)^{|A \setminus B|} \mu(B)$ for every $A \subseteq \mathcal{S}$. Intuitively, ϕ represents the information/the belief about the realization of event A that is not captured in its subevents. For instance, in Table 5, $\phi(\{R, Y\}) = 0$ because the agent has no more information about the chances of realization of R or Y than she has of R alone because there is no information available on the proportion of yellow balls in the urn. Above, we have characterized ϕ in terms of μ but it is also possible to characterize μ in terms of ϕ . Indeed, it is not difficult to show that $\mu(A) = \sum_{B \subseteq A} \phi(B)$ for all events A . This formula simply states that the agent's belief about event A corresponds to the sum of all her "elementary" beliefs on the realizations of A 's subevents. Thus, Belief $\mu(\{R, Y\})$ about the realization of event $\{R, Y\}$ corresponds to the belief generated by the sum of the information available about R alone, Y alone and the compound (R or Y) that could not be captured in singletons $\{R\}$ and $\{Y\}$. From a mathematical point of view, this translates as $\mu(\{R, Y\}) = \phi(\{R\}) + \phi(\{Y\}) + \phi(\{R, Y\})$.

Evt	\emptyset	$\{R\}$	$\{Y\}$	$\{B\}$	$\{R, Y\}$	$\{R, B\}$	$\{Y, B\}$	\mathcal{S}
f	0	1/3	0	0	1/3	1/3	2/3	1
ϕ	0	1/3	0	0	0	0	2/3	0

Table 5 The belief function of the Ellsberg's urn and its Möbius inverse.

In [Jaffray, 1989], Jaffray observed that the set of all the belief functions is a *mixture set*, i.e., it is closed under mixture operations or, in other words, any convex combination of belief functions is another belief function. In addition, he remarked that this is the key property used by von Neumann and Morgenstern to develop their axiomatic foundation of EU. As a consequence, it is possible to substitute in each axiom probabilities by belief functions. Expected utility thus boils down to a Choquet integral with respect to capacity μ . More precisely, in von Neumann-Morgenstern's framework, the probability distribution over the space of the states of nature generates, for each decision, a probability distribution over the outcomes of the decision, which is translated as a lottery. Here, Jaffray showed that the belief function over the space of the states of nature generates, for each decision, a belief function over the space of consequences. Let us call \mathcal{G} the space of these functions.

Theorem 6 (Jaffray, 1989). *The two assertions below are equivalent:*

1. Preference relation \succsim over \mathcal{G} satisfies axioms 1,2,3, where lotteries over \mathcal{L} are substituted by \mathcal{G} , the set of belief functions over the space of consequences.
2. \succsim is representable by a utility function

$$U : \mathcal{G} \mapsto \mathbb{R}$$

such that $U(\mu) = \int u d\mu$.

Function $u : \mathcal{X} \mapsto \mathbb{R}$ is called the von Neumann-Morgenstern utility function of the agent and is unique up to scale and location.

Therefore, the Choquet integral provides an attractive decision framework for situations in which probabilities are inadequate to model uncertainties. Thanks to the following definitions, it can be appropriately redefined in terms of Möbius inverses rather than belief functions: a belief function e_B is said to be elementary and concentrated on B if $e_B(A) = 1$ when $A \supseteq B$ and $e_B(A) = 0$ otherwise. In other words, its Möbius inverse ϕ_B is such that $\phi_B(B) = 1$ and $\phi_B(A) = 0$ for every $A \neq B$. Let μ be a belief function whose Möbius inverse is ϕ . The focal set \mathcal{C}_μ of μ is defined as $\mathcal{C}_\mu = \{B : \phi(B) > 0\}$. From these two definitions, it can be inferred that, for every belief function μ , and for every consequence set A , $\mu(A) = \sum_{B \subseteq A} \phi(B) = \sum_{B \in \mathcal{C}_\mu} \phi(B) e_B(A)$. But Theorem 6 trivially implies that, for every convex combination $\{\lambda_i, i = 1, \dots, n : \lambda_i \geq 0 \text{ and } \sum_{i=1}^n \lambda_i = 1\}$, $U(\sum_{i=1}^n \lambda_i \mu_i) = \sum_{i=1}^n \lambda_i U(\mu_i)$. As a consequence, if $\mu = \sum_{B \in \mathcal{C}_\mu} \phi(B) e_B$, $U(\mu) = \sum_{B \in \mathcal{C}_\mu} \phi(B) U(e_B)$. Let us denote by $u(B) = U(e_B)$ the utility of set of consequences B . Then, we get a linear utility model called *Belief expected utility (BEU)*:

Theorem 7 (Belief expected utility (BEU) – Jaffray, 1989). *The following two assertions are equivalent:*

1. Preference relation \succsim over \mathcal{G} satisfies axioms 1,2,3, where lotteries over \mathcal{L} are substituted by belief functions over \mathcal{G} .
2. \succsim is representable by a utility function $U : \mathcal{G} \mapsto \mathbb{R}$ such that $U(\mu) = \sum_{B \in \mathcal{C}_\mu} \phi(B) u(B)$, where $u(B)$ is the utility of set of consequences B and ϕ is the Möbius inverse of μ .

Table 6 illustrates the computation of U on the four alternatives A, B, C, D of the Ellsberg's urn. Assume that $u(\{0\}) = 0$, $u(\{1M\}) = 1$ and $u(\{0,1M\}) = \alpha$. Then $A \succ B$ and $D \succ C$ is equivalent to $\alpha < 1/2$. Therefore BEU is capable of representing "common" agent's preferences on the Ellsberg's urn.

Evs	0	1M	$\{0,1M\}$
balls	$\{B, Y\}$	$\{R\}$	\mathcal{S}
μ	2/3	1/3	1
ϕ	2/3	1/3	0

$$BEU(A) = 2/3u(\{0\}) + 1/3u(\{1M\}) = 1/3$$

Evs	0	1M	$\{0,1M\}$
balls	$\{R, Y\}$	$\{B\}$	\mathcal{S}
μ	1/3	0	1
ϕ	1/3	0	2/3

$$BEU(B) = 1/3u(\{0\}) + 2/3u(\{0,1M\}) = 2/3\alpha$$

Evs	0	1M	$\{0,1M\}$
balls	$\{B\}$	$\{R, Y\}$	\mathcal{S}
μ	0	1/3	1
ϕ	0	1/3	2/3

$$BEU(C) = 1/3u(\{1M\}) + 2/3u(\{0,1M\})$$

Evs	0	1M	$\{0,1M\}$
balls	$\{R\}$	$\{B, Y\}$	\mathcal{S}
μ	1/3	2/3	1
ϕ	1/3	2/3	0

$$BEU(D) = 1/3u(\{0\}) + 2/3u(\{1M\}) = 2/3$$

Table 6 BEU utilities for the Ellsberg's urn problem.

However, the BEU formula clearly highlights its limits w.r.t. EU: in the EU model, the agent's utility function u needs to be elicited only over the space of consequences \mathcal{X} whereas with BEU or CEU, it must be elicited on $2^{\mathcal{X}}$. Unfortunately, elicitation, i.e., the learning of the agent's preferences, is a complex and time consuming process. Therefore, to fix this problem, Jaffray proposed to add a new axiom called a "dominance" axiom to BEU. This axiom expresses the fact that, without any knowledge, within a set of consequences $\{x_1, \dots, x_k\}$, the agent has no reason to believe that a consequence is more likely to be yielded than any other. So the agent can summarize the information about the set of consequences by defining her preferences taking into account only the worst and the best consequences of the set. Consequently, utility $u(B)$ of a set of consequences B boils down to utility $u(m_B, M_B)$, where m_B and M_B denote the worst and the best consequences of B respectively.

Axiom 12 (Dominance) For every set of consequences $B \subseteq \mathcal{X}$, let m_B and M_B denote the worst and the best consequences of B respectively. Let e_B be the elementary belief function concentrated on B . Then, for every $B, B' \subseteq \mathcal{X}$, if $m_B \succsim_{\mathcal{X}} m_{B'}$ and $M_B \succsim_{\mathcal{X}} M_{B'}$ then $e_B \succsim e_{B'}$.

Theorem 8 (Jaffray's model, 1989). The following two assertions are equivalent:

1. Preference relation \succsim over \mathcal{G} satisfies axioms 1,2,3 and 12, where lotteries over \mathcal{L} are substituted by belief functions over \mathcal{G} .
2. \succsim is representable by a utility function $U : \mathcal{G} \mapsto \mathbb{R}$ such that

$$U(\mu) = \sum_{B \in \mathcal{E}_\mu} \phi(B)u(m_B, M_B).$$

Functions U and u are unique up to scale and location. In addition, u is a non-decreasing function of m and M and the von Neumann-Morgenstern utility $u(x)$ of consequence x is equal to $u(x, x)$.

As a consequence of this theorem, utility $u(m, M)$ takes into account two factors: i) the attitude of the agent w.r.t. risk (concavity of $u(x, x)$), but also ii) the attitude w.r.t. ambiguity when $M \neq m$. The model can be further refined using the Hurwicz criterion [Hurwicz, 1951]:

Definition 12 (Hurwicz criterion). for every (m, M) , let us call the “local optimism/pessimism criterion” the value $\alpha(m, M)$ for which the agent is indifferent between the following two alternatives:

1. winning m with probability $\alpha(m, M)$ and M with probability $1 - \alpha(m, M)$,
2. winning at least m and at most M , without any further information.

Thanks to this criterion, utility $u(m, M)$ can be redefined as $\alpha(m, M)u(m) + [1 - \alpha(m, M)]u(M)$, with $u(x)$ the von Neumann-Morgenstern utility function of the agent. In this context, coefficient α expresses the attitude of the agent w.r.t. ambiguity and the concavity of u expresses the agent’s attitude w.r.t. risk. Now, the task of eliciting the agent’s preferences (the learning of function u) has a complexity similar to that in the EU model.

4.1 Qualitative Decision Models under Uncertainty

In parallel to the research works made in the field of mathematical economics, decision under uncertainty has received attention in artificial intelligence. In particular, researchers investigated qualitative models, which describe preferences only through ordinal information [Tan and Pearl, 1994; Boutilier, 1994; Dubois and Prade, 1995; Brafman and Tennenholtz, 1996; Lehmann, 1996; Dubois et al., 1997]. Thus, within the framework of possibilistic lotteries [Dubois and Prade, 1995], Dubois and Prade proposed a counterpart to the von Neumann-Morgenstern axiomatic foundation: they axiomatized “qualitative utilities”, which generalize Wald criterion [Wald, 1950] for comparing possibility distributions. A possibility distribution is characterized by a function π which assigns to each consequence x its possibility $\pi(x) \in L$, L being an ordered set. The pessimistic qualitative utility model is based on an L -valued utility function u defined over the set of consequences \mathcal{X} , with L an ordered set. This function assigns to every possibilistic lottery π the following value:

$$U^-(\pi) = \min_{x \in \mathcal{X}} \max\{n(\pi(x)), u(x)\}$$

where n is a decreasing function which inverses the scale of L . Typically, when $L = [0, 1]$, n is chosen as $n(x) = 1 - x$. Value U^- indicates to which extent, by choosing π , the agent is sure to get a consequence having a “good” utility value. In the same possibilistic framework, there exists a more optimistic version which

evaluates to which extent it is possible that the agent gets a consequence with a “good” utility value. This version consists of assigning to every possibilistic lottery the following quantity:

$$U^+(\pi) = \max_{x \in \mathcal{X}} \min\{\pi(x), u(x)\}$$

The axiomatic foundation of Savage has also been revisited in order to propose qualitative counterparts to the EU model. Thus, Dubois, Prade and Sabbadin [Dubois et al., 1998] proposed axiomatic justifications for the optimistic and pessimistic qualitative utility criteria when comparing acts in the sense of Savage. This led to the following models:

$$U^-(f) = \min_{s \in \mathcal{S}} \max\{n(\pi(s)), u(f(s))\}$$

$$U^+(f) = \max_{s \in \mathcal{S}} \min\{\pi(s), u(f(s))\}$$

For every act f in $\mathcal{X}^{\mathcal{S}}$. $U^+(f)$ evaluates to which extent there exists a consequence of f which is at the same time very good and very plausible. On the other hand, $U^-(f)$ evaluates to which extent all the consequences in act f are plausible and good. These formula are therefore the numerical translations of logic principles. For more details, see [Dubois et al., 1999]. Dubois, Prade and Sabbadin have also proposed an axiomatic foundation of the Sugeno integral for comparing acts [Dubois et al., 1998], which led to the following model:

$$S_v(f) = \max_{x \in \mathcal{X}} \min\{v(F_x), u(x)\}$$

where $F_x = \{s \in \mathcal{S} : f(s) \geq x\}$ and v is a capacity defined on $2^{\mathcal{S}}$.

These models depart from EU notably by their weakening of the “sure thing principle”, which becomes the “weak sure thing principle”:

Axiom 13 (Weak Sure Thing Principle) *For every $f, g, h, h' \in \mathcal{X}^{\mathcal{S}}$ and for every $A \in 2^{\mathcal{S}}$, we have that $fAh \succ gAh \Rightarrow fAh' \succsim gAh'$.*

This axiom is important because, although it is weaker than the sure thing principle, it is sufficient to enable the computation of optimal policies in dynamic decision problems by backward induction. For more details on this point, see [Sabbadin, 1998].

Finally, pure ordinal aggregation rules (derived from majority rules used in voting) have been proposed under the name of “*lifting rules*” [Dubois et al., 2002, 2003]. To compare acts, they only rely on relative events likelihoods and on a preference relation over the consequences. They are defined as:

$$f \succsim g \text{ if and only if } \{s \in \mathcal{S} : f(s) \succsim_{\mathcal{X}} g(s)\} \supseteq \{s \in \mathcal{S} : g(s) \succsim_{\mathcal{X}} f(s)\}$$

where $\succsim_{\mathcal{X}}$ is the projection on the consequence scale of preference relation \succsim restricted to the constant acts, and \triangleright is a relative likelihood relation over the events. Their axiomatic justification is based on the introduction, in Savage's framework, of an axiom compelling the purely ordinal nature of the rule [Dubois et al., 2002, 2003]:

Axiom 14 (Ordinal invariance) [*for every $s \in \mathcal{S}$, $(f(s) \succsim_{\mathcal{X}} g(s)$ if and only if $f'(s) \succsim_{\mathcal{X}} g'(s)$) and $(g(s) \succsim_{\mathcal{X}} f(s)$ if and only if $g'(s) \succsim_{\mathcal{X}} f'(s)$)*]
 $\implies (f \succsim g \text{ if and only if } f' \succsim g')$.

This axiom states that preference $f \succsim g$ among two acts f and g , characterized by consequence vectors $(f(s_1), \dots, f(s_n))$ and $(g(s_1), \dots, g(s_n))$ respectively, does not depend on the relative positions of these consequences in the agent's preference scale, i.e., it only depends on preferences $f(s) \succsim_{\mathcal{X}} g(s)$ and $g(s) \succsim_{\mathcal{X}} f(s)$ for all the states of nature $s \in \mathcal{S}$. This model reminds of the relative concordance rules introduced in Chapter 16 of this volume in multicriteria decision making. Such rules do not necessarily induce transitive preferences, except when the beliefs over the events are highly hierarchical systems (see [Dubois et al., 2002, 2003] for more details). Here again, in order to obtain transitive preferences without constraining arbitrarily the beliefs over the events, it can be advantageous to introduce reference points in the model and to propose rules like:

$$f \succsim_r g \text{ if and only if } \{s \in \mathcal{S} : f(s) \succsim_{\mathcal{X}} r\} \triangleright \{s \in \mathcal{S} : g(s) \succsim_{\mathcal{X}} r\}$$

in which r represents a reference consequence on scale \mathcal{X} . For more details on this type of models, see [Perny and Rolland, 2006].

5 Sequential Decision Models

In practical situations, a decision is seldom made independently of the other decisions. Therefore, agents often have to choose among sets of decisions that must be made consecutively, each one having some impact on the next ones. In this section, we will study such problems and some decision models that were designed for that purpose.

Graphical models are well-suited for this task. "*Decision trees*" are certainly one of the most popular models. Their graphs contain two types of nodes: "*decision nodes*", drawn as rectangles, which represent the alternatives among which the agent has to choose; and "*chance nodes*", drawn as circles, which represent the uncertainties about the events. All these nodes are put into the graph in such a way that time always increases from the left to the right of the graph. Finally, to the leaves of the tree are assigned the utilities of the consequences resulting from the sequence of decisions and the set of events made from the root of the tree up to the leaves. Figure 9 represents a simple decision tree corresponding to the following problem [Raiffa, 1968]: An oil wildcatter must decide either to drill or not. He is uncertain whether

the hole is dry, wet or soaking. If he decides to drill, then, his gain will depend on the quantity of oil in the hole: if the hole is dry (no oil), he will lose 1M€; if the hole is wet, he will win 2M€; finally, if the hole is soaking, he will win 10M€. At a cost of 10K€, the wildcatter can make seismic soundings which help determine the geological structure of the site. The soundings will disclose whether the terrain below has no structure (*NoS*), in which case there is not much chance that the hole contains some oil, or open structure (*OpS*), in which case the presence of oil is somewhat more probable, or closed structure (*CIS*), in which case there are high chances that the hole contains a lot of oil. This problem can be modeled by a decision tree in the following way: The first decision to be made consists of making or not seismic soundings. This decision is represented by node T in Figure 9. If the oil wildcatter decides to make the soundings, we pass through the upper branch, else in the lower branch. Once the test is made, the wildcatter gets back the result R of the test. Of course, this result is only known after making the seismic soundings and, therefore, after making the decision to make the seismic soundings. This is the reason why node R must be located on the right of node T (time increases from left to right). Whatever the result of the test, upon knowing this result, the oil wildcatter must decide whether he will drill or not (nodes F_1). If he decides not to drill, then he will have lost the price of the seismic soundings, i.e., 10K€. This information can be found on the leaves of the tree. If the oil wildcatter decides to drill, then he will win the amount of money depending on the quantity of oil in the hole minus the price of the seismic soundings. This quantity (the E_i 's) is unknown when the agent makes the decision to drill or not, hence the E_i 's must be located on the right of F_1 in the decision tree. In the end, we get Figure 9. In general, on the branches outgoing chance nodes, are indicated the beliefs the agent has that the events will occur. Those are often the conditional probability that the event will occur given the values taken by all the preceding nodes, i.e., all the nodes to the left, up to the root. For instance, on the upper branch on the right of E_3 should be stored $P(E_3 = \text{dry} | F_3 = \text{yes}, R = \text{CIS}, T = \text{yes})$. Some variables can be independent from others, so this expression can often be simplified. Here, for instance, it is obvious that the state of the hole does not depend on the decisions of the agent, so the above conditional probability is equivalent to $P(E_3 = \text{dry} | R = \text{CIS})$. Probabilities on the branches outgoing the E_i 's therefore differ from one E_i to the other.

In addition to their capacity to model sequential decision making problems, decision trees can also be exploited to help agents making the best decisions. For this purpose, whatever the decision criterion chosen (EU, RDU, *etc.*), the idea is to look for an *optimal strategy*, i.e., in *every* decision node accessible given the set of all the decisions made previously, the choice of an alternative/decision among those possible at that node. Thus, a strategy considers all the states of nature possible. For instance, in Figure 9, the set of bold edges represents a strategy: when $T = \text{“yes”}$ is selected, as it is not possible to know which value R will take, we need to consider all the possible values for R and an alternative needs be selected for each node F_i . Note that, when the uncertainty within the chance nodes is modeled by probabil-

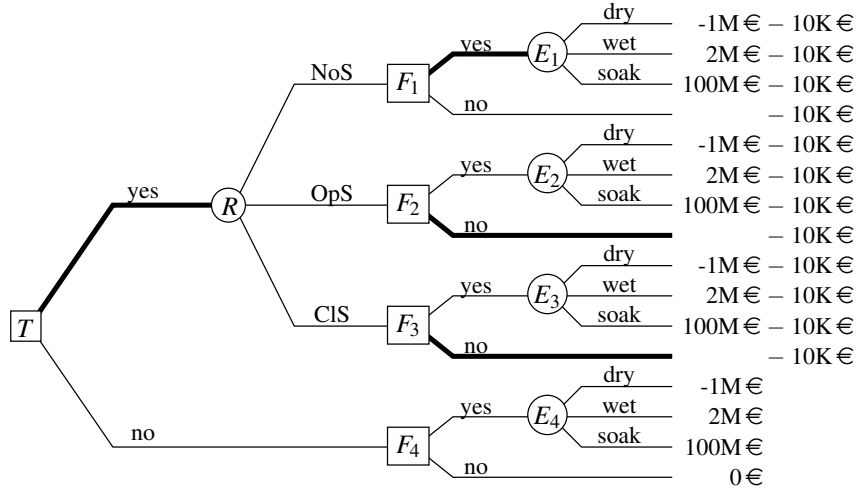


Fig. 9 Decision tree for the oil wildcatter problem.

ities, a strategy precisely corresponds to a lottery. Indeed, consider the strategy in bold edges in Figure 9. This one represent the fact that the agent will loose 10K € if $R = \text{“Ops”}$ or $R = \text{“ClS”}$ and that, if $R = \text{“NoS”}$, he will win 100M-10K € if $E_1 = \text{“soak”}$, 2M-10K € if $E_1 = \text{“wet”}$ and -1M-10K € if $E_1 = \text{“dry”}$. Therefore, this corresponds to lottery:

$$\langle -10\text{K} \text{€}, P(R = \text{Ops or Cls}) ; 100\text{M}-10\text{K} \text{€}, P(R = \text{NoS}, E_1 = \text{soak}) ; 2\text{M}-10\text{K} \text{€}, P(R = \text{NoS}, E_1 = \text{wet}) ; -1\text{M}-10\text{K} \text{€}, P(R = \text{NoS}, E_1 = \text{dry}) \rangle.$$

Therefore, finding the EU optimal strategy in a decision amounts to find the strategy whose corresponding lottery is optimal, i.e., it is maximal w.r.t. the EU criterion. Luckily, to determine it, it is not necessary to compute all the lotteries and to extract the best one. Actually, the above strategy can be rewritten as follows:

$$\langle \begin{array}{l} -10\text{K} \text{€}, P(R = \text{Ops}) ; \\ -10\text{K} \text{€}, P(R = \text{ClS}) ; \\ 100\text{M}-10\text{K} \text{€}, P(R = \text{NoS}) \times P(E_1 = \text{soak}|R = \text{NoS}) ; \\ 2\text{M}-10\text{K} \text{€}, P(R = \text{NoS}) \times P(E_1 = \text{wet}|R = \text{NoS}) ; \\ -1\text{M}-10\text{K} \text{€}, P(R = \text{NoS}) \times P(E_1 = \text{dry}|R = \text{NoS}) \end{array} \rangle. \quad (7)$$

Remark that the last three lines correspond to $P(R = \text{NoS})$ times the following lottery:

$$\langle \begin{array}{l} 100\text{M}-10\text{K} \text{€}, P(E_1 = \text{soak}|R = \text{NoS}) ; \\ 2\text{M}-10\text{K} \text{€}, P(E_1 = \text{wet}|R = \text{NoS}) ; \\ -1\text{M}-10\text{K} \text{€}, P(E_1 = \text{dry}|R = \text{NoS}) \end{array} \rangle. \quad (8)$$

which is nothing else than the lottery resulting from the bold strategy in the subtree whose root is F_1 . If, in the bold strategy of Figure 9, Decision $F_1 = \text{“yes”}$ is substituted by $F_1 = \text{“no”}$, it is easy to see that the resulting lottery will differ from that of Equation (7) only by the last three lines of the latter that are substituted by $P(R = \text{NoS})$ times lottery $\langle -10K \text{€}, 1 \rangle$, which is nothing else than the lottery corresponding to the strategy of the subtree rooted at the lower branch of F_1 . Consequently, to compare according to the EU criterion two lotteries L_1, L_2 that differ only in a subtree of the decision tree, it is sufficient to compute their respective lotteries in this subtree and to select the one with the highest EU score. As a matter of fact, the expectations of the sub-lotteries in the other subtrees are identical for both L_1 and L_2 so, due to the linearity of EU, they are irrelevant to compare L_1 and L_2 . This justifies that the following dynamic programming-based algorithm by backward induction can determine the EU-optimal strategy in all the decision tree: first, select the decisions that maximize EU in all the subtrees rooted at the decision nodes that are the closest to the leaves of the decision tree (in Figure 9, this corresponds to the subtrees rooted at $F_i, i = 1, \dots, 4$, respectively); then substitute these subtrees by leaves whose utility values are the expectations of these decisions, and iterate this process until reaching the root of the decision tree. The decision selected at each step of this algorithm constitute the EU-optimal strategy.

The goal of this chapter is not to develop computational decisional algorithmics, so we will not detail further this backward induction mechanism. However, it was useful to mention it when considering features of the “new” decision models like RDU. Actually, for these nonlinear models, backward induction produces incorrect results, as we will show in the next example. Suppose that the probability transformation function of the agent is $\varphi(x) = e^{-\sqrt{-\ln(x)}}$, as suggested by Kahneman and Tversky, and that her utility function is $u(x) = x$. Now, consider the decision tree of Figure 10. On the arcs outgoing from the chance nodes are indicated the probabilities of occurrence of their respective events and, on the right of the leaves are displayed the utilities of the consequences of the decisions. Calculating the RDU values of the strategies in this decision tree, we have that:

$$\begin{aligned} RDU(a) &= 2 + (5 - 2)\varphi(0,73) + (30 - 5)\varphi(0,25) = 11,41 \\ RDU(bc) &= 5 + (10 - 5)\varphi(0,5) + (20 - 10)\varphi(0,25) = 10,26 \\ RDU(bd) &= 2 + (5 - 2)\varphi(0,75) + (30 - 5)\varphi(0,25) = 11,46 \\ RDU(c) &= 10 + (20 - 10)\varphi(0,5) = 14,35 \\ RDU(d) &= 2 + (30 - 2)\varphi(0,5) = 14,18. \end{aligned}$$

In other words, in the subtree rooted at F , Strategy c is preferable to d , but in the subtree rooted at E , the optimal strategy is bd rather than bc .

This phenomenon is not restricted to the RDU criterion: it is general and appears as soon as the criterion departs from EU. In fact, to produce correct results, backward induction requires two properties: consequentialism and dynamic consistency. The first one states that, in each subtree, the optimal strategy depends only on this




Fig. 10 Decision tree and the RDU criterion.

subtree and not on the rest of the decision tree. The second property states that an optimal strategy in a subtree is an extension of optimal strategies in its own subtrees. As an example, if, in Figure 10, bd is an optimal strategy for the subtree rooted at E , then d must also be an optimal strategy in the subtree rooted at F . Unfortunately, consequentialism + dynamic consistency implies the “sure thing principle” (or at least a slightly weakened version) which leads to the EU criterion.

To complete our brief overview of sequential decision making, note that there exist compact representations of decision trees like, for instance, influence diagrams [Howard and Matheson, 1984; Shachter, 1986; Jensen et al., 1994]. The first key idea consists of considering decision trees as representations of “big” multivariate functions. The case of the decision trees with a symmetric structure simplifies the illustration of this idea: consider the trees of Figure 11. Instead of considering the utility values independently from one leaf to the other, consider the *set* of all these utility values as the result of a function depending on the values of D and O that led to the corresponding leaves. Similarly for the probabilities indicated on the branches of the decision tree, do not consider the values separately but as a whole as the probability distribution $P(O|D)$ depending on the values of D and O . The second key idea consists of exploiting the structural independences inherent to the decision problem. There is often a large number of such independences and those usually greatly simplify the “big” functions mentioned earlier. As an example, observe the 4 decision trees of Figure 11. At first sight, they seem quite similar. However, upon examining carefully the probabilities and the consequences/utilities displayed beside the branches of the tree, fundamental differences can be observed among these trees. In the first one, probabilities and utilities differ from one another on all the branches and, therefore, none of the functions $P(O|D)$ and $u(D, O)$ can be simplified. This is precisely what is represented by influence diagram 1 in Figure 12: circles represent chance nodes, to which are associated the conditional probabilities of these nodes given their parents in the graph (like in a Bayesian network [Pearl, 1988]); lozenges represent the utility multivariate functions and the variables they depend on correspond to those at the tails of their ingoing arcs. In tree 2 of Figure 11, it can be noticed that utility values depend on the branch outgoing from O where they are located but they do not depend on D . In other words, utility $u(D, O)$ can be summarized as $u(O)$ and this is precisely what influence diagram 2 of Figure 12 represents.

In Tree 3, utilities depend on D but not on O , hence influence diagram 3. Finally, in Tree 4, probabilities $P(O|D)$ do not depend on the value of D , which corresponds to influence diagram 4. To complete our description of influence diagrams, note that, although no function is associated with decision nodes, the latter can also have ingoing arcs. In this case, these arcs indicate the nodes (decisions and/or chances) whose values are known to the agent when she makes her decision.

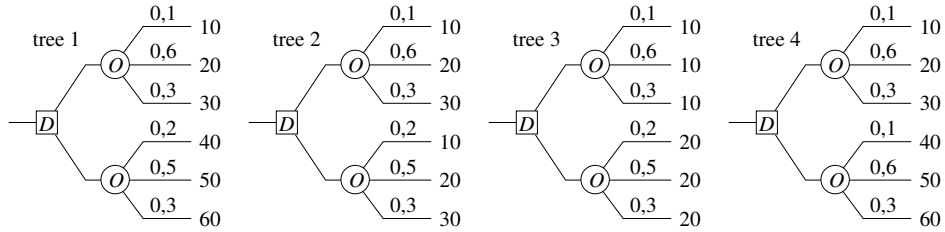


Fig. 11 Structural dependences in decision trees.

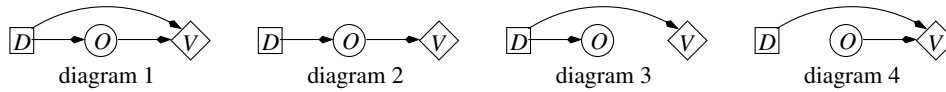


Fig. 12 Influence diagrams.

To conclude this section, note that models for representing sequential decision making problems are not restricted to decision trees and their compact representations (e.g., influence diagrams). Other formalisms do exist, which can be better suited for particular tasks. For instance, we can cite Markov decision processes (MDP) [Bellman, 1957; Howard, 1960; Puterman, 1994] or partially observed MDPs [Sondik, 1971; Monahan, 1982], which are especially useful in planning. Although these models have been based initially on probabilities, their possibilistic counterparts have been proposed in the literature [Fargier et al., 1998; Sabbadin, 2001]. In this chapter, we will not develop further these models since Chapter 10 of Volume 2 is devoted to them.

6 Conclusion

This chapter has provided a brief and non exhaustive overview of the theory of decision making under uncertainty. As we have seen, justifying mathematically the

proposed decision making models, relying on simple axioms reflecting common-sense features that are expected to be satisfied by any “*rational*” agent, has been one of the major concerns in the decision theoretic community. These axioms enable to justify to users these models and, more importantly, the recommendations they provide. This is a key point to make human agents/decision makers accept these models. Currently, the main research topics of the field are threefold and are focused on: i) the elicitation of preferences; ii) the models of uncertainty and their learning; and iii) the recommendation algorithms based on these models. Researches on preference elicitation focus on the minimization of the number of questions to ask to the agent to capture her preferences, but also on how to focus questions in order to elicit only the parts of the utility function that are needed to make “good” recommendations [Wang and Boutilier, 2003; Gonzales and Perny, 2004; Boutilier et al., 2010; Lu and Boutilier, 2011]. As for the uncertainties, new compact graphical models have been introduced recently (*Probabilistic Relational Models, Markov Logic Networks, Multi-Entity Bayesian networks*, etc), which notably enable learning from relational databases probability distributions defined over high-dimensional spaces, taking into account generic domain knowledge [Getoor and Taskar, 2007; Kok and Domingos, 2009; Khosravi et al., 2010]. Finally, recommendation algorithms have to address problems over combinatorial spaces of ever increasing sizes [de Salvo Braz et al., 2005; Regan and Boutilier, 2011].

For many years, in artificial intelligence, expected utility (EU) has been considered as the only reasonable model for decision under uncertainty. However, these last years, new decision theoretic models like RDU or Choquet have been introduced in the major AI conferences and their place shall increase in the next years. Indeed, they are not only capable to model faithfully the behaviors of agents facing uncertainty and ambiguity, but they also proved to be very useful for modeling fair and robust decision making problems. Finally, their expressive power should make them the models of choice for preference elicitation for high stakes strategic decision problems. However, exploiting such models requires a high level of information about the preferences of the agents as well as about the likelihoods of the events that may occur. Unfortunately, in some AI decision problems (like planning in partially known environments, preference elicitation and recommendation), the information available does not usually allow to precisely quantify the utility of an action or the probability of an event. In such situations, by relying on an ordinal representation of preferences and uncertainties, the qualitative models presented in the preceding sections prove to be better suited. To a large extent, these models are still unknown outside the academic world but, in the near future, their exploitation in industrial applications should increase significantly.

References

- Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école américaine. *Econometrica*, 21:503–546.

- Anand, P. (1993). The philosophy of intransitive preference. *The Economic Journal*, 103(417):337–346.
- Anscombe, F. and Aumann, R. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205.
- Argall, B. D., Chernova, S., Veloso, M., and Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57:469–483.
- Arrow, K. J. (1965). *Aspects of the Theory of Risk Bearing*, chapter The theory of risk aversion, pages 90–120. Yrjo Jahnsson Foundation.
- Bellman, R. (1957). *Dynamic Programming*. Princeton University Press.
- Bernoulli, D. (1738). Specimen theoriae novae de mensura sortis. *Commentarii academiae scientiarum imperialis Petropolitanae*, 5:175–192.
- Bleichrodt, H. (1996). *Applications of Utility Theory in the Economic Evaluation of Health Care*. PhD thesis, Erasmus University, Rotterdam, the Netherlands.
- Boutilier, C. (1994). Towards a logic for qualitative decision theory. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'94)*, pages 75–56.
- Boutilier, C. (2002). A POMDP formulation of preference elicitation problems. In *Proc. National Conf. on Artificial Intelligence (AAAI'02)*, pages 239–246.
- Boutilier, C., Regan, K., and Viappiani, P. (2010). Simultaneous elicitation of preference features and utility. In *Proc. National Conf. on Artificial Intelligence (AAAI'10)*, pages 1160–1167.
- Brafman, R. I. and Tennenholtz, M. (1996). On the foundation of qualitative decision theory. In *Proc. National Conf. on Artificial Intelligence (AAAI'96)*, pages 1291–1296.
- Chajewska, U., Koller, D., and Parr, R. (2000). Making rational decisions using adaptive utility elicitation. In *Proc. National Conf. on Artificial Intelligence (AAAI'00)*, pages 363–369.
- Chateauneuf, A. (1999). Comonotonicity axioms and RDU theory for arbitrary consequences. *J. of Mathematical Economics*, 32:21–45.
- Chateauneuf, A. and Cohen, M. (1994). Risk-seeking with diminishing marginal utility in a non-expected utility model. *J. of Risk and Uncertainty*, 9:77–91.
- Chateauneuf, A., Cohen, M., and Meilijson, I. (2004). Four notions of mean-preserving increase in risk, risk attitudes and applications to the rank-dependent expected utility model. *J. of Mathematical Economics*, 40(6):547–571.
- Chew, S., Karni, E., and Safra, Z. (1987). Risk aversion in the theory of expected utility with rank dependent preferences. *J. of Economic Theory*, 42:370–381.
- Chew, S. and Wakker, P. P. (1996). The Comonotonic Sure Thing Principle. *J. of Risk and Uncertainty*, 12:5–27.
- Conati, C., Gertner, A. S., VanLehn, K., and Drudzel, M. J. (1997). On-line student modeling for coached problem solving using Bayesian networks. In *Proc. Int. Conf. on User Modeling (UM'97)*.
- Dasgupta, P. (2006). Distributed automatic target recognition using multiagent UAV swarms. In *Proc. Int. Conf. on Autonomous Agents and Multiagent Systems (AA-MAS'06)*, pages 479–481.

- de Salvo Braz, R., Amir, E., and Roth, D. (2005). Lifted first-order probabilistic inference. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'05)*, pages 1319–1325.
- Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38:325–339.
- Doucet, A. and Johansen, A. (2011). *The Oxford Handbook of Nonlinear Filtering*, chapter A Tutorial on Particle Filtering and Smoothing: Fifteen years Later, pages 656–704. Oxford University Press.
- Dubois, D., Fargier, H., and Perny, P. (2003). Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. *Artificial Intelligence Journal*, 148(1):219–260.
- Dubois, D., Fargier, H., Perny, P., and Prade, H. (2002). Qualitative decision theory: from Savage’s axioms to nonmonotonic reasoning. *Int. J. of the Association of Computer Machinery*, 49(4):455–495.
- Dubois, D., Fargier, H., and Prade, H. (1997). Decision-making under ordinal preferences and uncertainty. In *Proc. Conf. on Uncertainty in Artificial Intelligence (UAI'97)*, pages 157–164.
- Dubois, D., Le Berre, D., Prade, H., and Sabbadin, R. (1999). Using possibilistic logic for modeling qualitative decision: ATMS-based algorithms. *Fundam. Inform.*, 37(1-2):1–30.
- Dubois, D. and Prade, H. (1995). Possibility theory as a basis of qualitative decision theory. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'95)*, pages 1924–1930.
- Dubois, D., Prade, H., and Sabbadin, R. (1998). Qualitative decision theory with Sugeno integrals. In *Proc. Conf. on Uncertainty in Artificial Intelligence (UAI'98)*, pages 121–128.
- Ellsberg, D. (1961). Risk, ambiguity and the Savage axioms. *The Quarterly J. of Economics*, 75:643–669.
- Fargier, H., Lang, J., and Sabbadin, R. (1998). Towards qualitative approaches to multistage decision making. *Int. J. of Approximate Reasoning*, 19:441–471.
- Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.
- Fishburn, P. C. (1982). *The foundations of expected utility*. Kluwer.
- Fishburn, P. C. and Roberts, F. S. (1978). Mixture axioms in linear and multilinear utility theories. *Theory and Decision*, 9:161–171.
- Franklin, R., Spiegelhalter, D., Macartney, F., and Bull, K. (1991). Evaluation of an algorithm for neonates. *British Medical Journal*, 302:935–939.
- Getoor, L. and Taskar, B. (2007). *Introduction to Statistical Relational Learning*. MIT Press.
- Gilboa, I. (1987). Expected utility with purely subjective non-additive probabilities. *J. of Mathematical Economics*, 16:65–88.
- Gilboa, I. (2009). *Theory of Decision under Uncertainty*. Econometric Society Monographs. Cambridge University Press.
- Gonzales, C. and Perny, P. (2004). GAI networks for utility elicitation. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'04)*, pages 224–234.

- Herstein, I. N. and Milnor, J. (1953). An axiomatic approach to measurable utility. *Econometrica*, 21:291–297.
- Horvitz, E. and Barry, M. (1995). Display of information for time-critical decision making. In *Proc. Conf. on Uncertainty in Artificial Intelligence (UAI'95)*, pages 296–305.
- Horvitz, E., Breese, J., Heckerman, D., Hovel, D., and Rommelse, K. (1998). The Lumière project: Bayesian user modeling for inferring the goals and needs of software users. In *Proc. Conf. on Uncertainty in Artificial Intelligence (UAI'98)*, pages 256–265.
- Howard, R. A. (1960). *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, Massachusetts.
- Howard, R. A. and Matheson, J. E. (1984). Influence diagrams. In Howard, R. and Matheson, J., editors, *Readings on the Principles and Applications of Decision Analysis*, volume 2, pages 719–762. Strategic Decision Group, Menlo Park, CA.
- Hurwicz, L. (1951). *Optimality Criteria for Decision Making Under Ignorance*, volume 370. Cowles Commission discussion paper, Statistics.
- Ingersoll, J. (1987). *Theory of Financial Decision Making*. Rowman and Littlefield.
- Jaffray, J.-Y. (1988). Choice under risk and the security factor: an axiomatic model. *Theory and Decision*, 24(2):169–200.
- Jaffray, J.-Y. (1989). Linear utility theory for belief functions. *Operations Research Letters*, 8:107–112.
- Jensen, F., Jensen, F. V., and Dittmer, S. L. (1994). From influence diagrams to junction trees. In *Proc. Conf. on Uncertainty in Artificial Intelligence (UAI'94)*.
- Jensen, F. V., Kjærulff, U., Kristiansen, B., Langseth, H., Skaanning, C., Vomlel, J., and Vomlelova, M. (2001). The SACSO methodology for troubleshooting complex systems.
- Jensen, N. E. (1967). An introduction to Bernoullian utility theory. I: Utility functions. *Swedish j. of Economics*, 69:163–183.
- Kahneman, D. and Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3:430–454.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, 47:263–291.
- Keeney, R. L. and Raiffa, H. (1993). *Decisions with Multiple Objectives - Preferences and Value Tradeoffs*. Cambridge University Press. (Version originale en 1976 chez Wiley).
- Khosravi, H., Schulte, O., Man, T., Xu, X., and Bina, B. (2010). Structure learning for Markov logic networks with many descriptive attributes. In *Proc. National Conf. on Artificial Intelligence (AAAI'10)*.
- Knight, F. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin.
- Kok, S. and Domingos, P. (2009). Learning Markov logic network structure via hypergraph lifting. In *Proc. Int. Conf. on Machine Learning (ICML'09)*.
- Kraft, C. H., Pratt, J. W., and Seidenberg, A. (1959). Intuitive probability on finite sets. *Annals of mathematical statistics*, 30:408–419.
- Lehmann, D. (1996). Generalized qualitative probability: Savage revisited. In *Proc. Conf. on Uncertainty in Artificial Intelligence (UAI'96)*, pages 381–388.

- Lu, T. and Boutilier, C. (2011). Robust approximation and incremental elicitation in voting protocols. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'11)*, pages 287–293.
- Machina, M. (1982). Expected utility analysis without the independence axiom. *Econometrica*, 50:277–323.
- Monahan, G. E. (1982). A survey of partially observable Markov decision processes: theory, models and algorithms. *Management Science*, 28:1–16.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman Publishers, inc.
- Perny, P. and Rolland, A. (2006). Reference-dependent qualitative models for decision making under uncertainty. In *Proc. European Conf. on Artificial Intelligence (ECAI'06)*, pages 422–426.
- Pratt, J. (1964). Risk aversion in the small and in the large. *Econometrica*, 32:122–136.
- Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley and Sons, New York.
- Quiggin, J. (1982). A theory of anticipated utility. *J. of Economic Behavior and Organisation*, 3:323–343.
- Quiggin, J. (1992). Increasing risk: another definition. In Chikani, A., editor, *Progress in Decision, Utility and Risk Theory*. Kluwer, Dordrecht.
- Quiggin, J. (1993). *Generalized Expected Utility Theory: The Rank-Dependent Model*. Springer.
- Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Addison-Wesley, Reading, MA.
- Ramsey, F. P. (1931). Truth and probability. In Ramsey, F., editor, *The Foundations of Mathematics and Other Logical Essays*. Harcourt, Brace and co.
- Regan, K. and Boutilier, C. (2011). Robust online optimization of reward-uncertain MDPs. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'11)*, pages 2165–2171.
- Rotschild, M. and Stiglitz, J. (1970). Increasing risk I: A definition. *J. of Economic Theory*, 2:225–243.
- Rotschild, M. and Stiglitz, J. (1971). Increasing risk II: Its economic consequences. *J. of Economic Theory*, 3:66–84.
- Sabbadin, R. (1998). *Une Approche Ordinale de la Décision dans l'Incertain: Axiomatisation, Représentation Logique et Application à la Décision Séquentielle*. Thèse de doctorat, Université Paul Sabatier, Toulouse, France.
- Sabbadin, R. (2001). Possibilistic Markov decision processes. *Engineering Application of Artificial Intelligence*, 14:287–300.
- Savage, L. J. (1954). *The Foundations of Statistics*. Dover.
- Schmeidler, D. (1986). Integral representation without additivity. *Proc. of the American Mathematical Society (AMS)*, 97:255–261.
- Shachter, R. (1986). Evaluating influence diagrams. *Operations Research*, 34:871–882.
- Shafer, G. (1976). *Mathematical Theory of Evidence*. Princeton University Press.

- Sondik, E. (1971). *The optimal control of partially observable Markov processes*. PhD thesis, Stanford university.
- Sordoni, A., Briot, J.-P., Alvarez, I., Vasconcelos, E., Irving, M., and Melo, G. (2010). Design of a participatory decision making agent architecture based on argumentation and influence function: application to a serious game about biodiversity conservation. *RAIRO Operations Research*, 44(4):269–284.
- Tan, S. and Pearl, J. (1994). Qualitative decision theory. In *Proc. National Conf. on Artificial Intelligence (AAAI'94)*, pages 928–933.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behaviour*. Princetown University Press, Princetown, New Jersey.
- Wakker, P. P. (1990). Under stochastic dominance Choquet expected utility and anticipated utility are identical. *Theory and Decision*, 29:119–132.
- Wakker, P. P. (1994). Separating marginal utility and risk aversion. *Theory and Decision*, 36:1–44.
- Wald, A. (1950). *Statistical Decision Functions*. Wiley and Sons, New York.
- Wang, T. and Boutilier, C. (2003). Incremental utility elicitation with the minimax regret decision criterion. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJ-CAI'03)*, pages 309–316.