# Optimising and Satisficing under Partial Ignorance

Frans Voorbraak
Department of Mathematics and Computer Science
University of Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam, The Netherlands
fransv@wins.uva.nl

**Abstract**

In this paper, we study decision making in situations where the outcomes of the options are (in general) uncertain, without making the assumption that this uncertainty can be exactly quantified by means of a (single) probability measure. In such situations of partial ignorance, the traditional notion of optimising by maximising expected utility is in general rather weak, and we will discuss some proposals to augment or refine the criterion of maximal expected utility under partial ignorance. We will argue that one interesting possible refinement, which is related to the minimax regret rule originally proposed for decision under uncertainty, is best viewed as a rule for satisficing rather than optimising. To make this argument more precise, we first propose a formalisation of the notion of satisficing under risk.

Keywords: decision making, partial ignorance, satisficing, minimax regret rule, independence of irrelevant alternatives.

## 1 Introduction

In this paper, we study decision making in situations where the outcomes of the options are (in general) uncertain, without making the assumption that this uncertainty can be exactly quantified by means of a (single) probability measure. This assumption, which underlies Bayesian decision analysis under risk, is unrealistic for many practical applications, since often the available evidence is insufficient to determine a unique probability measure which represents the uncertainty.

Next to decisions under risk, there exists a traditional category of decision under strict uncertainty (or complete ignorance), where the decision maker is not able to quantify this uncertainty at all and can only list the possible states of nature. We position ourselves between these two extremes and allow the decision maker to be *partially* ignorant about the relevant uncertainties. Actually, the case of partial ignorance includes the two extremes of risk and complete ignorance as special cases.

When the necessary probabilities and utilities are known, the decision rule of maximising the expected utility seems hard to beat. Although this rule is not without its critics, we will assume it as a guiding ideal for optimal decision making under risk. We will see that for decisions under partial ignorance the rule is not as powerful as for decision under risk, and often alternative (or additional) criteria are required to achieve satisfactory behaviour.

We believe that also in many practical decisions under risk it may be wise to search for satisfactory, rather than optimal, options. For example, when buying a house, it is not very practical to go see *all* houses that meet your minimal criteria concerning price, location, etc. A practical solution is to view some of these houses and buy one that is satisfactory in the sense that it is sufficiently close to what you expect to be the best possible house given your minimal criteria. If you have realistic expectations concerning the quality of the houses meeting your minimal criteria, this solution boils down to choosing an action with a sufficiently high satisfaction level, as defined in section 3.

We essentially represent the belief state of a partially ignorant agent by a set of all probability measures satisfying the evidence. Actually, we propose a more structured representation which includes a set of constraints on probability measures (representing the hard, generic evidence ), possibly some assumptions, and a conditioning event (representing the specific evidence).

The problem of reasoning and deciding under partial ignorance has been considered by several authors. See, for example, [3, 4, 5, 6, 7, 8, 9, 21, 23]. In some of these studies, the belief states are assumed to be represented by *convex* sets of probability measures or by lower and upper probability measures. We will argue against such restrictions, and we will indicate how the explicit mention of the underlying evidential constraints in the belief states makes can help in making them more manageable than arbitrary sets of probability measures.

In [19, 20] we propose an extension of the minimax regret decision rule, originally proposed for decision under strict uncertainty, as an interesting rule for decision under partial ignorance. In these papers, the extended minimax regret rule is primarily introduced as a possible refinement of the rule of maximising the expected utility, which is a rather weak rule in the context of partial ignorance. But it was already pointed out that this extended minimax regret rule may be more appropriate for satisficing than for optimising.

In this paper, we attempt to formalise at least some aspects of the notion of satisficing, both in the context of decision under risk and in the context of decision under partial ignorance, and we consider in more detail the role the extended minimax regret rule can play.

The rest of this paper is organised as follows. In section 2, we list some reasons why in some situations satisficing may be preferred to optimising, and we discuss some formal properties of the notions of satisfactory alternative and satisfaction level. In section 3, we propose a formalisation of a relative notion of satisficing under risk.

Section 4 contains a brief description of several variants of the rationality condition known as “independence of irrelevant alternatives”, and section 5 contains a brief discussion of the notion of partial ignorance.

In sections 6 and 7, we formalise optimising, respectively satisficing, under partial ignorance, and we argue that the extended minimax regret rule, originally proposed for decision under uncertainty or total ignorance, is better suited for satisficing than for optimising.

# 2 Optimising versus Satisficing

There are several reasons why for some decision problems it may be advisable to settle for a satisfactory alternative rather than insist on choosing an optimal alternative. In particular, when the decision maker is not (assumed to be) ideally rational, there may be practical difficulties hindering optimal choices. Also, the available information may be insufficient to completely determine optimal behaviour. We will discuss some of the reasons for satisficing in more detail below.

That practical decision problems are sometimes hard to formalise as (simple) optimisation problems is not a new insight. Researchers like H. Simon and I.J. Good already pointed out that real (as oppposed to ideal) decision makers have to deal with practical difficulties such as deliberation costs when solving real-life decision problems. If one thing has become clear from the work on limited rational decision makers, it is that it will be extremely difficult to formalise all relevant aspects of decision making under limited rationality.

The implications of dropping the assumption of ideal rationality have also been investigated in other areas than decision theory, for example, in epistemic logic (see [2]) and belief revision (see [22]). In these papers, a distinction is made between the implicit beliefs of a bounded reasoner and his actively or explicitly held beliefs. A similar distinction will be made in this paper.

Ideally, a decision maker understands what his available alternatives are, he knows the (utility) value of each alternative for each possible state of nature, and he can assign probabilities to the possible states of nature. It is then possible to compute the expected utility of each alternative and to optimise behaviour by choosing an alternative with maximum expected utility.

More formally, an idealised decision problem can be represented by a quadruple $\langle \Omega, P, A, U \rangle$, where $\Omega$ is a sample space representing the possible states of nature, $P$ is a probability measure over $\Omega$ representing the degrees of belief of the decision maker about the state of nature, $A$ is a set of functions on $\Omega$ representing the available actions among which the decision maker has to choose, and $U$ is a real-valued function on the set $\{a(\omega) : a \in A, \omega \in \Omega\}$ of outcomes of the actions. The expected utility (given $P$ and $U$) of an action $a$ is then defined to be

$$EU_{P,U}(a) = \sum_{\omega \in \Omega} U(a(\omega))P(\omega)$$

When no confusion is likely to occur, we will drop the explicit mention of the underlying probability function $P$, and we will simply write $U(a)$ instead of $EU_{P,U}(a)$. (As a result, the domain of the function $U$ is extended to include not only the possible outcomes of actions but also the actions themselves.)

We first give some examples of decision problems that may call for satisficing rather than optimising, and then discuss how they relate to the idealised decision problem described above.

**Example 1 (robot navigation)** *Consider a mobile robot moving around in an office-building. The robot has two manoeuvring modes, namely a normal mode and a smart mode. Manoeuvres in the smart mode are more precise, since they make heavy use of sensory feedback, but they can be time-consuming and they make it difficult to use the sensors for other tasks.*

*Suppose the next goal of the robot is to position itself in front of the next door on the left. Let us assume that the shorthest path to this goal position is achieved by moving forward 2.71 metres after making a 37 degrees left turn. Even if the robot is able to achieve this exactly by moving in the smart mode, it may be smarter to move in normal mode, because it is difficult to predict how much time the smart manoeuvres take, whereas the normal mode may be sufficiently precise to adequately approximate the shortest path.*

*Therefore, instead of insisting on taking the shortest, optimal path, it may be wise to choose a satisfactory solution, or nearly optimal path, if the latter can be achieved by simple, readily available means.*

**Example 2 (house searching)** *John is looking for a new house. A real-estate agent provides him with a (long) list of available houses meeting John's minimum requirements, concerning the maximum distance to John's office, the minimum-numbers of bedrooms, the maximum price, etc. The list contains some general information on the houses, but to really assess the value of a house, John has to visit it for inspection.*

*Since the list is too long for John to visit all listed houses, he has to decide after inspecting only some of them. In other words, instead of insisting on choosing the best of all available house, he has to settle for choosing among a subset of these houses, thereby perhaps missing the best one. Of course, John should keep on inspecting houses until he comes across one that is at least satisfactory, or until the cost of viviting new houses and postponing the decision becomes too high.*

In both examples one can say that the decision problem is originally presented in a way that abracts from practical considerations. The robot of example 1 cannot simply choose the shortest path to the goal, but it has to choose between possible actions of its actuators and sensors, where typically the results of these actions are not known exactly. In example 2, John cannot simply pick the best house from the list, since he first has to determine more exactly how good the houses are. In practice, John has to decide which houses to inspect first, how thorough such an inspection should be, and when to stop visiting new houses from the list.

Practical considerations, as those mentioned above, can, to some degree, be formalised in a decision-theoretic framework, but one will soon reach a point where the multitude of alternatives and relevant circumstances cannot longer be represented in a workable model. Also, the necessary probabilities and utilities will be hard to obtain. When applying decision theory, abstracting from practical intricacies is unavoidable.

One could maintain that decision theory should only be applied when the effects of the elements that are abstracted away are negligible, but we propose an approach where these effects are not assumed to be negligible, but taken into account in an abstract way. For example, in the case of the navigating robot (example 1), one can represent the difference in practical costs of the two kinds of manoeuvres as a qualitative difference instead of trying to formalise the exact costs.

Let us list some possible ways in which a decision problem can deviate from the idealised decision problem formalised in the beginning of this section.

1. (impractical options) Theoretically, all possible actions and all required probabilities and utilities are known, but some of the actions are practically more feasible than others. In example 1, the manoeuvres made in normal mode can be considered to be practically more feasible than the manoeuvres in smart mode.

2. (indeterminate utilities) All possible actions and all required probabilities are known, but some of the required utilities are not known exactly. In example 2, the utility of the uninspected houses is indeterminate.

3. (indeterminate probabilities) All possible actions and all required utilities are known, but some of the required probabilities are not known exactly. An example of this situation will be given later.

4. (missing alternatives) Some possible actions are not being considered. In example 2, the list provided by the real-estate agent may be incomplete.

In the next section, we will formalise decision problems in a such a way that we can capture all these cases except the case where the probabilities are indeterminate, which will be treated in a later section. An essential element of our proposed formalisation is the introduction of a set $A_0$ of practically feasible alternatives out of which a satisfactory element has to be chosen.

This set $A_0$ is closely related to what is called the set of "behavioural" alternatives by Herbert Simon in [15]. However, whereas in [15] the behaviourial alternatives are explained to be those alternatives of which the decision maker is aware, we allow other interpretation of these behaviourial alternatives as well. In spite of our somewhat liberal interpretation of the term "behaviourial alternative", we decided not to change the term, for lack of a better one.

One possible interpretation of behaviourial alternative is that of practically feasible option as in example 1. But also the inspected houses of example 2 can be considered to be behavioural alternatives, whereas the other houses on the list are merely theoretical alternatives. We will assume that the utility of a behavioural alternative is exactly known, whereas the utility of the other alternatives may be highly indeterminate.

Let us first discuss some reasonable minimal conditions on the notion of satisfactory alternative. We let $\geq$ denote the preference relation of the decision maker on the set of behavioural alternatives. In general, we allow this preference relation to be a partial ordering. However, in case the utilities of actions are determined, we assume that the preference relation mirrors the numerical comparison of the utilities, and is a preorder.

Satisfactory actions need not be maximal according to $\geq$, but the set of satisfactory actions should at least be compatible with respect to $\geq$ in the following sense.

**$\geq$-compatibility** If $a$ is satisfactory and $b \geq a$, then $b$ is satisfactory.

We propose to introduce levels, or degrees, of satisfaction, and we assume that if $a$ is $\alpha$-satisfactory, then $a$ is $\beta$-satisfactory, for all $\beta < \alpha$. The satisfaction level $S(a)$ of $a$ is defined to be the maximal $\alpha$ such that $a$ is $\alpha$-satisfactory. By convention, we use the interval $[0, 1]$ as the set of satisfaction levels.

It seems natural to refine the $\geq$-compatibility condition as follows.

**level $\geq$-compatibility** If $a$ is $\alpha$-satisfactory and $b \geq a$, then $b$ is $\alpha$-satisfactory.

The following proposition lists some properties that can be derived from level $\geq$-compatibility.

**Proposition 1** *Assume level $\geq$-compatibility.*

1. *If $a$ and $b$ are equally preferred ($a \approx b$), then $a$ is $\alpha$-satisfactory iff $b$ is $\alpha$-satisfactory.*
2. *If $b \geq a$, then $S(b) \geq S(a)$.*
3. *If $a$ is maximally satisfactory, then $a$ is $\geq$-maximal.*

In general, we cannot expect the reverse of these last two implications to hold, since we allow $\geq$ to be a *partial* ordering. If $\geq$ is a partial ordering, $\geq$-maximal alternatives need not be equally satisfactory, and therefore it is also not the case that the maximally satisfactory actions always have satisfaction level 1.

The second property of proposition 1 does not guarantee that strict preference between alternatives results in different satisfaction levels. In other words, it may be the case that $b$ is strictly preferred to $a$, but still $S(b) = S(a)$. The following condition ensures that strict preference is mirrored in the comparison of the satisfaction levels.

**level $>$-compatibility** If $b > a$, then $S(b) > S(a)$.

We will assume both level $\geq$-compatibility and level $>$-compatibility, and we will call this *preference compatibility.*

**Proposition 2** *Assume preference compatibility and assume that $\geq$ is a total ordering.*

*1. $b \geq a$ iff $S(b) \geq S(a)$.*

*2. $a$ is maximally satisfactory iff $a$ is $\geq$-maximal.*

Starting from satisfaction levels, one can define the set of satisfactory alternatives to be the set of alternatives with a satisfaction level above (or equal to) a particular threshold. Notice that under this definition, the level $\geq$-compatibility condition implies $\geq$-compatibility.

The threshold mentioned above can be viewed as a normalised version of what is called "aspiration level" in [15]. We prefer a *normalised* version, since we favour a *relative* view on satisficing, where the level of satisfaction depends on what the possible alternatives are. For example, when a decision maker is faced with a dilemma and has to choose between two evils, he should be satisfied if he is able to identify and execute the choice for the lesser evil, even though the same choice might be considered totally unaccepable under more favourable circumstances.

Whenever the (expected) utility of an alternative $a$ is determined, we assume that its satisfaction level is a function of its (expected) utility. In other words, $S(a) = f(U(a))$, for some function $f$. From the assumption of preference compatibility it follows that $f$ should be monotonically increasing.

In the case of determinate utilities, it is natural to view the function $S$ as a normalised version of the utility function $U$. This can be made more precise if one assumes the following. Let $\alpha a + (1-\alpha)b$ denote the combined action of choosing either $a$ or $b$, where $\alpha$ is the probability that $a$ is chosen.

**mixture compatibility** $S(\alpha a + (1-\alpha)b) = \alpha S(a) + (1-\alpha)S(b)$.

As is well known from [11], preference compatibility and mixture compatibility together characterise (up to a positive linear transformation) the utility function in the case of a total preference order. That is, under relative weak assumptions on the ordering $>$ and the mixture function, preference compatibility and mixture compatibility together imply that there is utility function such that $>$ exactly mirrors the comparison of expected utility. Moreover, we have the following.

**Proposition 3 (Von Neumann and Morgenstern)** *If $f$ and $g$ are two real-valued functions compatible with a total preference order and with mixture, then $f = \lambda g + \mu$, for some $\lambda > 0$ and $\mu$.*

It follows that in the case of determinate utilities $S$ is a positive linear transformation of $U$. In general, a slightly more complex relation between utilities and satisfaction levels has to be found in case the utilities are allowed to be indeterminate.

# 3 Satisficing under Risk

It follows from the previous section that in the case of a decision problem under risk $\langle \Omega, P, A, U \rangle$ the satisfaction level $S_{P,U}(a)$ of an action $a$ is a positive linear transformation of its (expected) utility. To completely determine the satisfaction levels, it is therefore sufficient to fix two values of this transformation.

Given that we have opted for a *relative* notion of satisficing, one obvious choice is to set $S_{P,U}(a) = 1$ whenever alternative $a$ has the maximum expected utility of all available alternatives, i.e., $EU_{P,U}(a) = max\{EU_{P,U}(b) : b \in A\}$. Let us write $M_{P,U}$ to denote $max\{EU_{P,U}(b) : b \in A\}$.

One could set $S_{P,U}(a) = 0$ whenever alternative $a$ has the minimum expected utility of all available alternatives, but this would have some undesirable consequences. For example, it would effectively mean that there do not exist decision problems under risk where every alternative is satisfactory.

**Example 3** *Suppose John has to choose between two houses with approximately the same (expected) utility. Then it is reasonable to expect him to be satisfied with any one of the two possible choices. Therefore, it does not seem correct to set the satisfaction level of the house with the (slightly) smaller expected utility to 0.*

In answer to the problem above, we propose a transformation which assigns satisfaction level 0 to some utility strictly smaller than the expected utility of each available alternative. Let us write $m_{P,U}$ to denote this value. In many cases, one can, for example, define $m_{P,U}$ to be $min\{U(b(\omega))) : b \in A, \omega \in \Omega\}$, i.e., the utility of the worst possible outcome of the available alternatives. We then arrive at the following formalisation of satisfaction levels for decision making under risk.

**Definition 1 (satisficing under risk)** *Let $\langle \Omega, P, A, U \rangle$ be a decision problem under risk. The* satisfaction level $S_{P,U}(a)$ *of an action $a$ is defined as follows.*

$$S_{P,U}(a) = \frac{EU_{P,U}(a) - m_{P,U}}{M_{P,U} - m_{P,U}}.$$

This definition is rather abitrary, in particular concerning the assignment of satisfaction level 0. However, the exact parameters of the transformation of utility do not seem to matter much, as long as the transformation forces the satisfaction levels of available alternatives to lie in the interval $[0, 1]$ and it is understood that the choice of the threshold for satisficing has to take into account which particular transformation is used.

Since utilities are determined up to a positive linear transformation, one can, before defining satisfaction levels, first transform $U$ to a function $U'$ such that $m_{P,U}$ is transformed to $m_{P,U'} = 0$. Then the expression for the satisfaction level $S_{P,U'}(a)$ simplifies to $\frac{EU_{P,U'}(a)}{M_{P,U'}}$.

Notice that in the case of satisficing under risk, the reverse of the third part of proposition 1 does hold, and the actions which maximise expected utility all have satisfaction level 1. In other words, optimal actions can be defined as actions with satisfaction level 1.

Let us now consider a generalisation of decision under risk, where we allow the utilities to be indeterminate, and where we distinguish a subset of the (theoretically) available alternatives as the set of readily available, or practically feasible, or behavioural alternatives.

**Definition 2 (generalised decision under risk)** *A generalised decision problem under risk is a tuple* $\langle \Omega, P, A, A_0, \Upsilon \rangle$*, where* $\Omega, P,$ *and* $A$ *are as before and where* $A_0 \subseteq A$ *represents the set of behavioural alternatives, and* $\Upsilon$ *is a set of utility functions.*

If $A_0 = A$ and $\Upsilon$ is a singleton, then the generalised problem reduces to an ordinary decision problem under risk. The behavioural alternatives of $A_0$ represents actions that can be chosen in practice, whereas all other members from $A$ represent theoretically possible actions that are (at least at the moment) considered impractical.

In example 1 (robot navigation) the possible manoeuvres in normal mode can be interpreted as behavioural alternatives, and the smart mode manoeuvres are the theoretical possible actions. In our view, the satisfaction level of a normal mode manoeuvre depends on the maximal attainable expected utility, where also the smart mode manoeuvres are taken into account. Therefore, the decision maker is not only required to be aware of the behavioural alternatives, but he also has to be aware (to some degree) of the other possibilities. However, since these other possibilities are not considered to be practical options, their utilities may be left indeterminate.

The situation of example 2 (house searching) can be formalised as a generalised decision problem by defining $A$ as the set of choices for a house from the list provided by the real-estate agent, and $A_0$ as the subset of $A$ obtained by restricting to the inspected houses. The utility of an unexamined house is highly indeterminate, whereas the value of the inspected houses can be assumed to be known exactly. (It is not difficult to drop this last assumption, and allow also some indeterminacy in the utility of the inspected houses.)

Generalised decision problems can also model "missing alternatives", i.e. possible actions that are not being (explicitly) considered. For example, a possibly suitable house not on the list provided by the real-estate agent can simply be represented by a generic "other house not on the list". Since very little is known about such a house, its utility will typically be highly indeterminate. Using this representation, finding such a house would not have to result in a *revision* of the belief state of the decision maker, but can be accomodated by a kind of *refinement* of the set of alternatives.

The definition of satisficing under risk can be generalised to generalised decision problems as follows.

**Definition 3 (generalised satisficing under risk)** *Let* $\langle \Omega, P, A, A_0, \Upsilon \rangle$ *be a generalised decision problem under risk. The* satisfaction level $S_{P,\Upsilon}(a)$ *of a behavioural alternative* $a$ *is defined as follows.*

$$S_{P,\Upsilon}(a) = \inf\{S_{P,U}(a) : U \in \Upsilon\}.$$

For a generalised decision problem $\langle \Omega, P, A, A_0, \Upsilon \rangle$, the maximal attainable expected utility can be defined to be $M_{P,\Upsilon} = sup\{EU_{P,U}(a) : a \in A, U \in \Upsilon\}$. If we write $m_{P,\Upsilon}$ to denote the value $min\{U(b(\omega))) : b \in A, \omega \in \Omega, U \in \Upsilon\}$, we can define the following approximation of satisfaction levels for generalised decision under risk.

**Definition 4 (approximated satisficing under risk)** *Let $\langle \Omega, P, A, A_0, \Upsilon \rangle$ be a generalised decision problem under risk. Let $U$ be some utility function from the set $\Upsilon$. The* approximated satisfaction level *$s_{P,\Upsilon}(a)$ of a behavioural alternative $a$ is defined as follows.*

$$s_{P,\Upsilon}(a) = \frac{EU_{P,U}(a) - m_{P,\Upsilon}}{M_{P,\Upsilon} - m_{P,\Upsilon}}.$$

An advantage of the approximated satisfaction levels is that they are normalised utilities, and therefore they are relatively well understood. The function $s_{P,\Upsilon}$ satisfies mixture compatibility, whereas $S_{P,\Upsilon}$ does not. Of course, both functions satisfy preference compatibility. An advantage of definition 3 is that it can straightforwardly be adapted to the case where the behaviourial alternatives are also allowed to have indeterminate utilities.

We propose satisficing, in the context of a generalised decision problem under risk, to consist of choosing among the behavioural alternatives $A_0$ an action with a satisfaction level at least equal to some threshold. One can define $\alpha$-satisficing to be satisficing with threshold $\alpha$.

It might seem reasonable to require that if there are several sufficiently satisfactory actions among $A_0$, then one should choose an action with maximal satisfaction level. Notice, however, that this procedure still does not reduce to ordinary optimising restricted to $A_0$. This is due to the fact that it is not guaranteed that a maximally satisfactory alternative will be chosen, since the expected utility of each behavioural alternative may lie below the threshold. Also, we will argue in section 7 that in some cases optimising using satisfaction levels may result in dependence on irrelevant alternatives.

If no sufficiently satisfactory (behavioural) alternative can be found, then the decision maker has to do some work to come to a decision. It is far from obvious what exactly should be done by the decision maker in such a case.

For example, if at a certain moment John considers all the previously inspected houses unsatisfactory, he is not yet ready to make a decision and buy some house. Then he can visit some more houses from the list (in order to reduce the indeterminacy of the utilities), he can try to get some information about the actual state of the world (adjusting the probability function), or he can adjust his expectations and decide to require a lower threshold for satisficing.

In other words, the problem of determining the value of (additional) information, which is elegantly solved in the context of optimising under risk, is much more complex in the case of satisficing under indeterminate utilities. As discussed in [19, 20], the presence of partial ignorance, or indeterminate probabilities, can further complicate the issue. The complicating factors are that generalising decision under risk results in several different relevant types of information, and that the notion of satisficing includes the possibility of "lowering one's standard" by lowering the threshold of satisfaction, or the aspiration level.

Perhaps it will be impossible to solve the matter without introducing a separate metareasoning level, as proposed in [13]. However, it seems that the present formalisation at least gives some indications for the best non-behavioural alternatives to explore next. For example, alternatives with a utility equal to the maximal attainable expected utility seem promising.

# 4 Independence of Irrelevant Alternatives

In [12] it is argued that optimising is not the only rational decision procedure and the condition of "independence of irrelevant alternatives" (actually called "consistency condition" in [12]) is proposed as a test of the rationality of a choice function.

In fact, this condition of independence of irrelevant alternatives is perhaps better known as a condition on (rational) preference orderings.

**Definition 5 (order independence)** *Let $\geq_A$ denote the preference ordering of a decision maker on the set of alternatives $A$. The preferences of the decision maker are said to satisfy independence of irrelevant alternatives iff for all sets of alternatives $A', A''$ including both $a_1$ and $a_2$, $a_1 \geq_{A'} a_2$ implies $a_1 \geq_{A''} a_2$*

In the words of Rubinstein [12]: "if the decision maker ranks one alternative above another when facing a set of options that includes both, he will rank them identically when encountering any other decison problem in which these two alternatives are available."

The ordering on actions induced by comparing their (expected) utility satisfies independence of irrelevant alternatives. It follows that in the case of decision under risk also the ordering induced by satisfaction levels satisfies independence of irrelevant alternatives.

**Definition 6 (choice function)** *A* choice function *on a set of alternatives $A$ is a partial function $C : 2^A \to A$ such that $C(A') \in A'$ whenever $C(A')$ is defined.*

**Definition 7 (choice independence)** *Let $A$ be a set of alternatives. A choice function $C$ on $A$ is said to satisfy independence of irrelevant alternatives iff for all $A' \subseteq A'' \subseteq A$, $C(A'') \in A'$ implies $C(A') = C(A'')$.*

**Proposition 4** *If the choice function is defined from a total, asymmetrical preference ranking by choosing the maximally preferred option, then order independence implies choice independence.*

It follows that if all alternatives have different (expected) utilities, then maximising expected utility is choice function under risk satisfying independence of irrelevant alternatives. In case some alternatives have the same expected utility (and the induced preference ordering is not asymmetrical) a suitable tie-breaking rule is necessary to guarantee choice independence. For example, one could use any fixed total, asymmetrical ordering $O$ of the alternatives to choose between equally valued alternatives.

Optimising expected utility is not the only procedure satisfying choice independence. For example, choosing the first (according to a fixed total, asymmetrical ordering $O$) alternative with a utility above (or equal to) a fixed threshold also satisfies choice independence. This latter procedure can be considered to be a form of satisficing, although, based on our normalised interpretation of satisfaction levels, we propose a slightly different satisficing procedure.

It is possible to abstract from the considerations concerning "suitable" tie-breaking rules, by generalising the notion of choice function.

**Definition 8 (generalised choice function)** *A* generalised choice function *on a set of alternatives $A$ is a function $C : 2^A \to 2^A$ such that $C(A') \subseteq A'$.*

A generalised choice function does not necessarily choose a unique alternative, but only selects a set of "acceptable" alternatives.

**Definition 9 (generalised choice independence)** *Let $A$ be a set of alternatives. A generalised choice function $C$ on $A$ is said to satisfy generalised independence of irrelevant alternatives iff for all $A' \subseteq A'' \subseteq A$, $C(A'') \cap A' \subseteq C(A')$.*

Notice that for choice functions, generalised choice independence and choice independence are equivalent. In section 7, we will see that a satisficing procedure can satisfy generalised choice independence, whereas the underlying preference ordering does not satisfy order independence.

# 5 Partial Ignorance

In [17, 18] we introduce a generalization of Bayesian probability theory called partial probability theory (PPT). In PPT, a belief state is represented by constraints on probability measures representing generic information about the situation plus a set $C$ representing the available specific evidence on which one has to condition the probability measures satisfying the constraints. More precisely, we have the following.

**Definition 10** *A* (partial) probabilistic belief state with assumptions (pbsa) *is a tuple $\langle \Omega, \mathcal{B}, \mathcal{A}, C \rangle$, where $\Omega$ is a sample space (which we assume to be finite), $\mathcal{B}$ is a set of (belief) constraints on probability measures over $\Omega$, $\mathcal{A}$ contains constraints representing assumptions, and $C$ is a subset of $\Omega$ on which the probability measures are to be conditioned.*

The constraints of $\mathcal{A}$ represent reasonable assumptions for which strictly speaking no evidence is available. The notions of generic and specific information are explained in [1] and [17]. The following example illustrates the different ingredients in a pbsa.

**Example 4 (battery recharge)** *Consider a robot which has to recharge its battery. This can be done in two rooms, let us call them room 1 and 2. We will write $a_i$ for the choice for room $i$. An example of generic information might be: "the door of room 1 is at least 40% of the time open". Suppose there is no other information available, and let $p_i$ denote the probability of door $i$ being open. Then $\mathcal{B} = \{p_1 \geq 0.4\}$. Since doors are typically sometimes open and sometimes closed, it might be reasonable to include $0 < p_i < 1$ in $\mathcal{A}$. However, such additional assumptions should be invoked only when they are necessary, for example, in case no reasonable decision can be made without assumptions. A good example of specific evidence is information about the state of the doors obtained by the sensors of the robot.*

Any pbsa $\langle \Omega, \mathcal{B}, \mathcal{A}, C \rangle$ naturally induces the set $\mathcal{B}(\Omega)$ of probability measures satisfying the constraints, i.e., $\mathcal{B}(\Omega) = \{P : P \text{ satisfies } \mathcal{B}\}$. The conditioning information can be incorporated by means of extended Bayesian conditioning: $\mathcal{B}(\Omega)_C = \{P_C : P \text{ satisfies } \mathcal{B}, P(C) > 0\}$.

Since such sets of probability measures may be rather complex entities, several authors propose to restrict attention to closed convex sets, and other authors propose not to work with the sets of measures themselves, but with their lower and upper envelopes. Since different sets of probability measures can have the same lower and upper envelopes, the latter proposal involves some loss of information, which may be quite serious. In our view, even the former, less radical, proposal unnecessarily restricts the range of applicability of the theory. See example 6 in section 6.

The proposed PPT belief states are in some sense richer than sets of probability measures, since they distinguish the set of probability measures satisfying the (hard) evidence and the (possibly smaller) set of probability measures additionally satisfying some reasonable assumptions. For example, in PPT, the maximum entropy principle can in principle be represented as an additional assumption on top of the hard evidence represented by the constraints in $\mathcal{B}$.

However, PPT belief states are in some sense also less general than arbitrary sets of probability measures, since the sets induced by a pbsa are always definable by the constraints. Therefore, one can use the constraints to reason with the represented information, and if the language in which the constraints are formulated is sufficiently restricted this may be far less complex than reasoning with arbitrary sets of probability measures. Of course, a restricted language does not allow the representation of all arbitrary sets of probability measures.

For the purpose of this paper, we will abstract from the details of PPT belief states and we will simply represent partial ignorance by a set of probability functions.

**Definition 11 (decision under partial ignorance)** *A decision problem under partial ignorance is a tuple $\langle \Omega, \Pi, A, U \rangle$, where $\Omega, A$, and $U$ are as before, and where $\Pi$ is a set of probability functions over $\Omega$.*

In this paper, we do not distinguish between behavioural alternatives and theoretical possibilities in the context of deciding under partial ignorance. Also, we do not study decision problems where both probabilities and utilities are allowed to be indeterminate. Although combining both kinds of indeterminacy does not present any insurmountable difficulties, we feel it would unnecessarily complicate the presentation.

# 6 Optimising under Partial Ignorance

The maximum expected utility decision criterion for deciding under risk can be naturally extended to situations where the uncertainty is only partially quantified.

**Definition 12 (extended MEU)** *Let $\langle \Omega, \Pi, A, U \rangle$ be a decision problem under partial ignorance. Define a preference relation between actions as follows.*

$$a \geq_{\Pi} b \text{ iff } \forall P \in \Pi(EU_{P,U}(a) \geq EU_{P,U}(b)).$$

*According to the extended maximum expected utility rule one should choose an action which is maximally preferred according to the above relation $\geq_{\Pi}$.*

One might consider replacing the given definition of the extended MEU rule with a criterion that only refers to upper and lower probability and expected utility. (For example, action $a$ is weakly preferred over action $b$ iff the lower expected utility of $a$ is at least the upper expected utility of $b$.) The following example indicates a problem with this approach.

**Example 5** *Consider again example 4, but now suppose that all that is known about the state of the doors is that it is more probable that door 1 is open than that door 2 is open. That is, $\mathcal{B} = \{p_1 > p_2\}$. The choice for $a_1$ is still supported by the extended MEU rule, but actions $a_1$ and $a_2$ have the same lower an upper expected utility.*

In case the available evidence completely determines a probability measure quantifying the uncertainty about the state of nature, the extended MEU criterion is exactly the ordinary maximum expected utility criterion, and the chosen action can be said to be (weakly) preferred to every other available action.

However, in general, the preference relation of definition 12 is only a *partial* ordering. This means that actions may be incomparable and that the extended MEU criterion does not necessarily point to a unique set of equivalent actions. For example, given the information about the state of the doors in the example 4 (battery recharge), the extended MEU rule does not decide between actions $a_1$ and $a_2$.

In [4], Giron and Rios give an axiomatic characterisation of $\geq_{\Pi}$ which only differs from a characterisation of the usual maximum utility preference ordering by allowing the ordering to be partial. This shows that in a sense the partiality of the preference is a characteristic property of the extended MEU criterion.

This result in [4] is actually given for those $\geq_{\Pi}$ where $\Pi$ is assumed to be a non-empty closed convex set of probability measures. The following result shows that this assumption is not essential for the characterization.

**Proposition 5** *Let $conv(\Pi)$ be the closed convex hull of $\Pi$. Then $\geq_{conv(\Pi)} = \geq_{\Pi}$.*

The above result might be (and has been) interpreted as an argument in favour of the restriction to closed convex sets of probability measures. However, the following example shows the importance of having available non-convex sets at the credal level.

**Example 6** *Consider again example 4, but now suppose that according to one source one of the doors is twice as likely to be open as the other one, whereas according to a second source it is more probable that door 1 is open than that door 2 is open. That is, $\mathcal{B}_1 = \{p_1 = 2p_2 \vee p_2 = 2p_1\}$ and $\mathcal{B}_2 = \{p_1 > p_2\}$. Let $\Pi_i$ be the set of probability measures compatible with $\mathcal{B}_1$. Then $(\Pi_1 \cap \Pi_2)$ supports the (intuitively correct) conclusion that door 1 is twice as likely to be open as door 2 ($p_1 = 2p_2$). But this conclusion is not supported by $conv(\Pi_1) \cap conv(\Pi_2)$.*

Although the extended maximum expected utility criterion is quite natural, it has an obvious drawback. The criterion is rather weak, particularly if there is much ignorance, in the sense that the relevant probabilities are highly indeterminate.

The question what to do in case two alternatives $a$ and $b$ are incomparable according to $\geq_{\Pi}$ is rather controversial. Some authors propose that one should base one's decisions on a particular probability measure compatible with the available evidence. This probability measure is sometimes allowed to be aribitrarily chosen from the candidates, and sometimes it is assumed to be in some sense least informative, such as the measure with the maximum entropy. In [10] it is argued that if one decides on the basis of a particular probability measure compatible with the available evidence, then one could just as well use the strict Bayesian approach (and not allow indeterminate probabilities) from the start.

Other authors allow decision makers to be indecisive between incomparable maximally preferred actions. See, for example, [21]. In principle, we prefer to allow indecisiveness between incomparable maximally preferred actions, but we believe that more refined decision rules than extended MEU should be considered in order to assure that the remaining candidate actions are really acceptable.

For example, the extended MEU rule does not decide which door the robot of example 4 should choose in case $p_1 \geq 0.4$ is the only constraint. However, in this case, many people will intuitively prefer room 1.

In [19, 20] we propose to refine the extended MEU criterion by introducing some additional criteria obtained from decision rules originally proposed for decision under complete ignorance. These refinements can support intuitive preferences like the one mentioned above. In particular, we propose the following extended minimax regret rule.

**Definition 13 (extended minimax regret)** *Let $\langle \Omega, \Pi, A, U \rangle$ be a decision problem under partial ignorance. Define the $P, U$-regret $R_{P,U}(a)$ and the maxi-*

*mal regret $R_{\Pi,U}(a)$ of an action $a$ as follows.*

$$R_{P,U}(a) = M_{P,U} - EU_{P,U}(a).$$

$$R_{\Pi,U}(a) = \sup\{R_{P,U}(a) : P \in \Pi\}.$$

*Here, as before, $M_{P,U} = max\{EU_{P,U}(b) : b \in A\}$. According to the extended minimax regret rule one should choose an action with minimal maximal regret.*

The following proposition summarises some nice properties of this decision rule.

**Proposition 6** *The extended minimax regret rule generalises both Savage's minimax regret rule [14] and the ordinary MEU rule, and it refines the extended MEU rule. More precisely,*

1. *In the case of complete ignorance, the extended minimax regret rule reduces to Savage's minimax regret rule.*
2. *In the case of decision under risk, the extended minimax regret rule reduces to the ordinary MEU rule.*
3. *The extended minimax regret rule refines extended MEU in the sense that $a \geq_\Pi b$ implies $R_{\Pi,U}(a) \leq R_{\Pi,U}(b)$.*

However, there are some problems with using the extended minimax regret rule. For example, in [19, 20] we also consider refining extended MEU with an extension of Wald's maximin return rule and since the extended maximin return rule and the extended minimax regret rule do not necessarily agree, one has to choose between the two rules (and other related rules).

A perhaps even more serious (and well-known) problem is that the ordering obtained by comparing maximal regret does not satisfy order independence, i.e. a choice between two alternatives based on the (extended) minimax regret rule can be influenced by the availability of a third alternative.

# 7 Satisficing under Partial Ignorance

The definition of satisficing under risk can be extended to the case of satisficing under partial ignorance, in a way that is similar to our previous extension to the case of indeterminate utilities.

**Definition 14 (satisficing under partial ignorance)** *Let $\langle \Omega, \Pi, A, U \rangle$ be a decision problem under partial ignorance. The* satisfaction level *$S_{\Pi,U}(a)$ of an alternative $a$ is defined as follows.*

$$S_{\Pi,U}(a) = \inf\{S_{P,U}(a) : p \in \Pi\}.$$

Notice that the satisfaction level of an alternative is closely related to (a normalised version of) its minimal maximal regret. More precisely, if we write $R'_{P,U}(a)$ for $\frac{R_{P,U}(a)}{M_{P,U}-m_{P,U}}$, and $R'_{\Pi,U}(a)$ for $\sup\{R'_{P,U}(a) : P \in \Pi\}$, then $S_{\Pi,U}(a) = 1 - R'_{\Pi,U}(a)$. Therefore, maximising the satisfaction level can also be formulated as minimising the maximal normalised regret.

In view of this close relation between satisfaction level and maximal regret, it is no surprise that the ordering induced by comparing satisfaction levels (under partial ignorance) does not satisfy order independence.

**Example 7** *Consider the decision problem $\langle \Omega, \Pi, A, U\rangle$, where $\Omega = \{\omega_1, \omega_2\}$, $\Pi = \{P_1, P_2\}$, $A = \{a_1, a_2, a_3\}$, and the functions $P_1$, $P_2$, and $U$ are defined as follows. $P_1(\omega_1) = P_2(\omega_2) = 0.7$, $U(a_1(\omega_1)) = 100$, $U(a_1(\omega_2)) = 0$, $U(a_2(\omega_1)) = U(a_2(\omega_2)) = 40$, $U(a_3(\omega_1)) = 0$, and $U(a_3(\omega_2)) = 100$.*

*Then we have $EU_{P_1,U}(a_1) = 70$, $EU_{P_2,U}(a_1) = 30$, $EU_{P_1,U}(a_2) = EU_{P_2,U}(a_2) = 40$, $EU_{P_1,U}(a_3) = 30$, and $EU_{P_2,U}(a_3) = 70$. It follows that $a_2$ has the highest satisfaction level ($\frac{4}{7}$) of the three alternatives. However, if alternative $a_3$ is dropped from the considerations, then the satisfaction level of $a_1 = \frac{3}{4} > \frac{4}{7} =$ the satisfaction level of $a_2$.*

It follows that the (optimising) generalised choice function which selects the set of alternatives with maximal satisfaction level (or minimal maximal normalised regret) does not generalised choice independence. However, the (satisficing) generalised choice function that selects the set of alternatives with satisfaction above (or equal to) a fixed threshold does satisfy generalised *choice* independence. More precisely, we have the following.

**Proposition 7** *Let $\langle \Omega, \Pi, A, U\rangle$ be a decision problem under partial ignorance. Define the generalised choice function $C$ on $A$ as follows. For every $B \subseteq A$, $C(B) = \{a \in B : S_{\Pi,U}(a) \geq \alpha\}$, where alpha is some fixed threshold. Then $C$ satisfies generalised choice independence.*

The above choice function can also be defined in terms of selecting alternatives with extended maximal normalised regret below (or equal to) a fixed threshold. Therefore, we conclude that the notion of extended maximal regret is better suited for satisficing than for optimising.

# 8 Conclusion

We have presented a formalisation of a (relative) notion of satisficing, both in the context of decision under risk, and in the context of decision under partial ignorance. We discussed several reasons why a decision maker may opt for satisficing rather than optimising.

In particular, in the case of decision making under partial ignorance, the natural extension of the ordinary MEU criterion may be rather weak, especially in the presence of highly indeterminate relevant probabilities. An extension of the minimax regret criterion to partial ignorance can be used the strengthen or refine the extended MEU criterion. However, optimising based on comparing extended maximal regret is shown not to be independent of irrelevant alternatives, whereas satisficing is. Moreover, the extended maximal regret of an alternative is shown to be closely related to the satisfaction level of the alternative.

We conclude that the extended minimax regret criterion, previously studied as a criterion for optimising, is actually better suited for satisficing than for optimising. We do not claim that by applying the criterion based on extended maximal regret, or the associated satisfaction level, one obtaines a full-fledged formalisation of satisficing. This requires much more insight into questions concerning the value of information. These question are much harder in the context of satisficing and decision under partial ignorance than in the context of optimising under risk.

## Acknowledgments

The investigations were carried out as part of the PIONIER-project Reasoning with Uncertainty, subsidized by the Netherlands Organization of Scientific Research (NWO), under grant pgs-22-262.

## References

[1] D. Dubois, H. Prade and P. Smets, Representing partial ignorance, *IEEE Trans. on Systems, Man, and Cybernetics*, 26 (1996) 361-378.

[2] R. Fagin and J.Y. Halpern, Belief, awareness, and limited reasoning, *Artificial Intelligence* 34 (1988) 39-76.

[3] P.C. Fishburn. *Decision and Value Theory*, Wiley, New York, 1964.

[4] F.J. Giron and S. Rios. Quasi-Bayesian Behaviour: A more realistic approach to decision making?, in *Bayesian Statistics*, J.M. Bernardo et al. (eds.), University Press, Valencia, 1980, 17-38.

[5] Z.W. Kmietowicz and A.D. Pearman, *Decision Theory and Incomplete Knowledge*, Gower, 1981.

[6] H. Kyburg, *The Logical Foundations of Statistical Inference*, Reidel, Dordrecht, 1974.

[7] I. Levi, *The Enterprise of Knowledge*, MIT Press, Cambridge Mass.,1980.

[8] R. Loui, Decisions with indeterminate probabilities, *Theory and Decision* 21 (1986) 283-309.

[9] R. Luce and H. Raiffa, *Games and Decisions*, Wiley, New York, 1957. Reprinted by Dover Publications (New York) in 1989.

[10] P. Maher, What is wrong with strict bayesianism? *PSA 1986 Vol. 1*, A. Fine and P. Machamer (eds.), (1986) 450-457.

[11] J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behaviour*, Princeton University Press, Princeton, 1944.

[12] A. Rubinstein, *Modeling Bounded Rationality*, MIT Press, Cambridge Mass., 1998.

[13] S. Russell and E. Wefald, *Do the Right Thing. Sudies in Limited Rationality*, MIT Press, Cambridge Mass., 1991.

[14] L.J. Savage, *The Foundations of Statistics* (second revised edition), Dover, New York, 1972.

[15] H. A. Simon, A behavioral model of rational choice, *Quarterly Journal of Economics*, 69 (1955) 99-118. Reprinted in H. A. Simon, *Models of Bounded Rationality. Vol. 2*, MIT Press, Cambridge Mass. (1982) 239-258.

[16] H. A. Simon, Theories of bounded rationality, *Decision and Organization*, C. B. Radner and R. Radner (eds.), Amsterdam: North-Holland (1972) 161-176. Reprinted in H. A. Simon, *Models of Bounded Rationality. Vol. 2*, MIT Press, Cambridge Mass. (1982) 408-423.

[17] F. Voorbraak, Probabilistic belief expansion and conditioning. *Technical report LP-1996-07*, ILLC, University of Amsterdam, 1996.

[18] F. Voorbraak, Combining evidence under partial ignorance. *Qualitative and Quantitative Practical Reasoning. Proc. ECSQARU-FAPR'97*, D. Gabbay et al. (eds.), Springer, Berlin (1997) 574-588.

[19] F. Voorbraak, Decision analysis using partial probability theory, in: J. Doyle et al. (eds.), *Qualitative Preferences in Deliberation and Practical Reasoning (AAAI-97 Spring Symposium)*, Stanford University, 1997, 113-120.

[20] F. Voorbraak, Deciding under partial ignorance, in: *EUROBOT'97: Second EUROMICRO Workshop on Advanced Mobile Robots*, IEEE/CS Press, 1997, 66-72.

[21] P. Walley. *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, London, 1991.

[22] R. Wassermann, Resource bounded belief revision. *Technical report LP-1998-04*, ILLC, University of Amsterdam, 1998.

[23] M. Wolfensen and T.L. Fine, Bayes-like decision making with upper and lower probabilities, *Journal of the American Statistical Association* 77 (1982) 80-88.