Rational Decision Making With Imprecise Probabilities

Jean-Yves Jaffray
LIP6 - Université Paris 6, 4, place Jussieu 75252 Paris cedex 05, France

Abstract

Decision criteria based on an imprecise probability representation of uncertainty have been criticized, from the normative point of view, on the grounds that they make the decision maker (DM) vulnerable to manipulations and, more generally, likely to take up a dominated strategy. This is indeed the case when the DM is both consequentialist (his choices in a subtree are not influenced by data concerning the rest of the tree) and sophisticated (his present choice, determined by backward recursion, is best given his future choices). Renouncing consequentialism, which, as shown by Machina, is a way out of the difficulty, may seem to increase excessively the complexity of the model. We revisit the whole question, and first argue that in sequential decision situations it is possible to separate preference from choice, without abandoning the Revealed Preference Creed; then, we propose to assume consequentialist preference and accept non-consequentialist behavior; building on these assumptions, we discuss McClennen's Resolute Choice model and its interpretation involving multiple Selves; finally, taking the decision aiding point of view, we suggest an implementation of Resolute Choice in which the consensus goal among the Selves is to select an non-dominated strategy.

Keywords. imprecise probabilities, rationality, dynamic decision making, resolute choice

1 Introduction

The recourse to a decision criterion involving a non-probabilistic representation of uncertainty -such as upper/lower probability intervals- has been criticized in the past, from the normative point of view, on the grounds that it makes the decision maker (DM) vulnerable to manipulations inducing losses and, more generally, makes him likely to take up dominated strategies in dynamic choice situations[1] (see e.g., Hammond 1988 [1]).

However, as remarked by McClennen (1990) [4] and Machina (1989) [3], all the arguments advanced rely heavily on the assumption that the DM is *consequentialist* (his choices in a subtree are not influenced by data concerning the rest of the tree). They further observe that renouncing consequentialism opens the door to alternative types of behavior , such as *Resolute Choice* (a term due to McClennen), in which the DM, by committing himself to a strategy chosen once for all, can get round the traps of sequential choice. However modeling Resolute Choice is not an easy task: a consequentialist DM has well defined local preferences which can be revealed (in the sense of Samuelson's Revealed Preference theory), so that one need then only a simple behavioral trait (myopia, sophistication, etc...) to be specified for predicting his choices in any dynamic decision problem; on the contrary, with a non-consequentialist DM, it is not clear what underlying preference elements could be revealed nor, if any, how a model of the psychological process leading to resolute choice should integrate them.

In this paper we adopt a decision aiding point of view which is close to that usually taken in multicriteria decision making (MCDM): there, the DM, feeling unable to make alone a synthesis of his conflicting objectives, calls in an Analyst who, on the basis of data the DM can provide at his request (such as local trade-offs between objective levels), proposes successive choices which hopefully converge to an acceptable compromise solution; here, we shall assume that the Analyst can determine what the DM would choose in single-decision problems and, moreover, how much he would be ready to deviate from it in specific circumstances. The normative goal of the Analyst will be to help

[1] The very same criticisms apply, in the case of risk (probabilistic uncertainty), to criteria deviating from EU theory

the DM select a strategy which is non-dominated (in some specific sense).

The paper is organized as follows: first, the revelation of preference from choice in the framework of dynamic decision making under uncertainty (DDMUU) is discussed; then, the basic notions of consequentialism, dynamic consistency and rationality are defined and illustrated by an example. next, McClennen's resolute choice concept is presented and discussed; finally proposals for implementing resolute choice are made.

2 Dynamic decision making under uncertainty (DDMUU)

N.B. *In the following, algebras of events and sets of decisions considered are all finite. Thus only basic concepts of Decision Analysis (decision tree, decision node, chance node, strategy, substrategy, etc...) are needed and we assume that the reader is familiar with them as well as with their standard graphical representation conventions. The set of consequences, interpreted as gains or losses, is always* $\mathfrak{R}$.

2.1 Choice and preference

The fact that DDMU involves multiple, sequential and conditional, choices makes it necessary to look closely at the relation between choice and preference. Economists consider that preference, which is an abstract concept, should be derived from choice, which can be observed: in Samuelson's (1948) [6] terms, preference has to be *revealed* by choice. This requirement cannot be satisfied *stricto sensu* in DDMUU models, where preference among strategies plays a fundamental role, since, obviously, observation cannot reveal a whole strategy but only the sequence of decisions along the path determined by the state of nature which has occurred. To get around this difficulty, one has to admit that even hypothetical choices bring meaningful information. Then, paths which would be taken if such of such state of nature would occur can also be revealed.

The sequential character of the decision process is the source of a more serious difficulty. At any decision node N, the DM only controls the next decision; subsequent decisions are controlled at subsequent decision nodes. Thus it would be preposterous to infer that the substrategy with root N is the substrategy that the DM prefers at N. Does then the very concept of revealable preference at N make sense? A natural idea is that the substrategy the DM prefers at N is the strategy he would enforce if he had the possibility to do so. Such a situation can be represented by modifying the subtree with root N, T(N), and replacing it by a new subtree, T'(N) , build as follows : its root, N, is the unique decision node; each decision issued from it corresponds to a substrategy of the initial subtree; and subsequent chance nodes describe the progressive determination of the outcome of the associated substrategy (T'(N) is a kind of normal form of T(N)). This however is not enough, since choices at different nodes of a tree can be interdependent (and, in fact, will be interdependent in our model), in which case preference at N, when derived from the choice at N in this modified tree, is not well defined (Wakker 1997 [7]). Thus one is led to try and derive preference at N from choices in a tree where N is the unique decision node. One can for instance consider T'(N) in isolation; or further modify the initial tree by assessing the decisions at all decision nodes outside T'(N). Different definitions may lead to different preference relations.

To sum up, preference at a node of a decision tree can always be linked to choice as required by Revealed Preference theory. However the choice used to derive this preference may have to be an hypothetical choice in a different tree and remains somewhat arbitrary. The underlying idea is that choice observed in multiple decision situations is the result of a complex elaboration process involving "intrinsic" preference which can (only) be revealed by isolated choice.

2.2 Consequentialism and dynamic consistency

When choice at a decision node and preference at the same node are clearly differentiated, the fundamental notions of consequentialism and dynamic consistency can be defined with precision.

The DM displays a *consequentialist* behavior when the substrategy he chooses in a subtree does not depend on the rest of the decision tree. This is true in particular of the derived tree where he becomes a dictator at node N, so that his preference at node N is well defined and moreover independent of the rest of the tree; thus a consequentialist DM has also consequentialist preference. A consequentialist DM can anticipate his future choices and construct recursively, by backward induction, a strategy which at any node induces the best substrategy (for the local preference) in the restricted set of available substrategies.This is called *sophisticated* behavior.

A DM has *dynamically consistent* preference whenever substrategies induced at future nodes by strategies judged optimal according to present preference are themselves optimal according to future preference.

When a DM is both consequentialist and dynamically consistent, the sophisticated strategy is an opti-

Figure 1: Example.

mal strategy for the initial preference and all its substrategies are optimal for the preferences at their root nodes. This is in particular the case of subjective expected utility (SEU) maximizers.

As an illustration consider the following example.

Example 1 *Denote by MAX+MIN the criterion which ranks decisions according to the sum of their highest and lowest possible consequences (this criterion can be defended in a situation of complete ignorance). Consider the following decision tree in Figure 1:*

i) A first DM has consequentialist preference,expressed by criterion MAX+MIN at both nodes N_0 and N_1. At node N_0 he prefers strategy (D, d') to strategy (D, d), whereas at N_1he prefers d to d': he is not dynamically consistent. The sophisticated strategy is D', which is strictly dominated by (D, d'), i.e., whatever happens he would be better off with this last strategy.

ii) A second DM has also preference criterion MAX+MIN at N_0 and is dynamically consistent: his preference ranking at N_1 is induced by his preference ranking at N_0: since he prefers (D, d') at N_0, he prefers d' at N_1.However, he cannot be consequentialist since, in the modified tree where D yields 10 if E^c occurs, he would prefer (D, d) at N ,hence d at N_1: his preference at N_1 depends on the value of a gain outside of the subtree with root N_1.

2.3 Rational behavior

As we have seen, in the above example the first DM chooses a strictly dominated strategy. Such behavior is generally considered as irrational, that is to say, unacceptable from a normative point of view.

Rationality requirements concerning DDMUU are best expressed on acts. Every strategy generates an act a, i.e., a mapping *state of nature* $\longmapsto$ *consequence:* $\omega \mapsto a(\omega)$. Then, act a *weakly dominates* act b when $a(\omega) \geq b(\omega)$ for all ω and $a(\omega_0) > b(\omega_0)$ for some ω_0, whereas act a *strictly dominates* act b when $a(\omega) > b(\omega)$ for all ω. Whereas the respect of strict dominance is an undisputed rationality requirement, the respect of weak dominance is not, at least in its general formulation (ω_0 might be negligible). Anyhow, in particular situations of uncertainty specific dominance requirements are generally added. For instance, under *risk*, a situation characterized by the existence of a probability P on the events, the chosen act, say b, should not be *stochastically dominated*, i.e., there should be no other feasible act a satisfying : $P\{\omega : a(\omega) \leq c\} \leq P\{\omega : b(\omega) \leq c\}$ for all c and $P\{\omega : a(\omega) \leq c_0\} < P\{\omega : b(\omega) \leq c_0\}$ for some c_0.

A natural extension of stochastic dominance (SD) to the case of *imprecise probability*, in the interpretation which postulates the existence of a true, but unknown probability, which can only be located in a set of probabilities $\mathcal{P}$, is *imprecise stochastic dominance* (ISD), defined as follows: a ISD b when $P\{\omega : a(\omega) \leq c\} \leq P\{\omega : b(\omega) \leq c\}$ for all c and for all $P \in \mathcal{P}$, and $P_0\{\omega : a(\omega) \leq c_0\} < P_0\{\omega : b(\omega) \leq c_0\}$ for some c_0 and some $P_0 \in \mathcal{P}$ .Different extensions of SD might fit better other semantics of imprecise risk, for instance: a ISD' b when $\sup_{P \in \mathcal{P}} P\{\omega : a(\omega) \leq c\} \leq \inf_{P \in \mathcal{P}} P\{\omega : b(\omega) \leq c\}$ for all c, with $<$ instead of $\leq$ for some c_0.

NB. *By extension, we attribute from now on to strategies the properties of the corresponding acts.*

Strategies which are not weakly dominated (hence not strictly dominated either) are easily generated since they can be obtained as optimal strategies in the maximization of expected gain for any positive probabilities. It is also well known that, under risk, strategies which are not stochastically dominated can be found by maximizing expected utility for any increasing utility functions $u(.)$. Under imprecise risk, a strategy which uniquely maximizes expected utility for some $u(.)$ and some $P \in \mathcal{P}$ cannot be ISP-dominated. Other methods (used in Multicriteria Decision Making) exist which allow to generate non-dominated strategies which are not accessible by the previous methods.

In the rest of the paper, "dominance" will refer indifferently to any of the above relations which is relevant to the situation considered. In particular, in the case of imprecise probability, the relation considered could be strict dominance as well as ISD or ISD'.

3 Resolute choice

3.1 McClennen's resolute choice

Example 2 *(continued) In Example1, the behavior of the second DM may receive the following explanation: the MAX+MIN criterion dictates his choice whenever he has to take a unique decision. However, when examining the decision tree, he becomes aware of the fact that if he applies this criterion at each decision node he will end up selecting a dominated strategy (as does the first DM), which he considers to be a waste. Therefore he takes the resolution to continue at node N_1 the strategy he had judged best at node N_0, thus managing to select a non-dominated strategy.*

Submission of later choice to initial preference is of course only a special case of more general forms of compromise between present and future wishes. We shall refer to any such compromise by the name of *resolute choice* in accordance with McClennen's (1990) [4], p.260, definition : "*the theory of resolute choice is predicated on the notion that the single agent who is faced with making decisions over time can achieve a cooperative arrangement between his present self and his relevant future selves that satisfies the principle of intrapersonal optimality*".

3.2 Discussion

A fundamental problem with resolute choice is the question of its psychological feasibility. This problem is relevant not only from a descriptive point of view, but also from the prescriptive, decision aiding, point of view adopted in this paper. The solution proposed by the Analyst to the DM has to be carried out by the latter, who should not reject it, as psychologically unacceptable, at some point.

Can the Selves co-operate efficiently? Can they cooperate at all? A preliminary question is: what exactly is a Self, and is it an appropriate concept?

A Self is attached to a decision node N, but cannot be identified with the future (and contingent) DM at N, since it must pre-exist at the moment (say date 0) where the conference takes place. Thus it is just a name for that part of the present DM's mind which manages the interests of the future node N-DM (we shall nonetheless write "the Self's preference" for "the DM's preference at N"; etc...)

Karni and Safra (1989) [2] identify the Selves with the players of a game and predict, in accordance with a standard result of Game Theory, that the non-cooperative solution - sophisticated choice, a subgame perfect Nash equilibrium (Rasmusen 1989 [5]), will take place. However, there is a fundamental aspect in which Selves differ from non-cooperative game players, which casts some doubt on the relevance of this analysis: as parts of the same mind, the Selves cannot hide one another their intentions, which leaves no room for bluff, betrayal, etc...This can only facilitate cooperation.

On the other hand, the Selves will not cooperate unless they have consensus goals, such as achieving a non-dominated choice (of some kind). It is not obvious that the Self who is attached to node N should care about what happens outside the subtree with root N. However, any coalition of Selves, regrouping all the Selves likely to exist at date t if a given strategy is chosen between 0 and t, should collectively reject dominated strategies. This may justify a collective effort.

Suppose now that each Self knows himself perfectly, i.e., not only knows his preference but also knows how much he is willing to give up, by deviating from his best feasible choice to favor the achievement of a consensus goal. Suppose moreover that, as parts of the same mind, he also knows perfectly all the the other Selves. Then, it is not difficult to admit that some conference should take place and attempts at achieving a satisfying cooperative solution should be made. These attempts might of course fail if the willingness to cooperate of the Selves is not sufficient.

The arrangement processes taking place spontaneously in people's mind are not known. Anyhow, decision aiding methods do not have necessarily to mimic them. Since sequential choice is involved, it is natural to look for processes relying on backward induction. We shall propose two closely related solutions.

3.3 Two cooperative decision processes

FRAMEWORK. Imprecise probability situation, in which the DM has a separable preference criterion (the same at each decision node). The value of strategy or substrategy S (in fact the value of the act it generates) is denoted by $V(S)$. He is willing to cooperate, but refuses to lose a value superior to θ with respect to its best feasible choice. The consensus goal is to achieve a non-dominated strategy (for a specific dominance relation).

PRELIMINARY STEP (common to the two processes). A set $\mathcal{S}$ of non-dominated strategies is generated. Each one receives a label. The initial decision tree is pruned by cutting off all decisions which do not belong to at least one of the strategies in $\mathcal{S}$. The new tree has each of its decision edges labeled with the labels of all the strategies it is part of.

3.3.1 The simple veto process

RECURSIVE STEP. The current period is t. At each decision node N of period t, the DM evaluates the best labeled substrategy at N among the available labeled substrategies. Let v be its value. He then vetoes all the presently available labeled strategies that induce a substrategy at N with value inferior to $v - \theta$. Any labeled strategy vetoed by at least one of the nodes N of period t becomes unavailable at the earlier periods.

IF there is no labeled strategy left, *no consensus has been achieved.* END. ELSE (t-1) becomes the current period.

IF t=1, the root node is the unique N . The DM chooses the best strategy among the available labeled strategies. ELSE one goes back to the recursive step.

Thus the process can fail to select a non-dominated strategy, by lack of flexibility (one may perhaps still hope to find one, though, by increasing the number of non-dominated strategies proposed initially). Note that, if the DM is totally flexible, the strategy selected is simply the one judged best at the root.

3.3.2 The strong veto process

RECURSIVE STEP. The current period is t. At each decision node N of period t, for each decision d at that node, the DM evaluates the best labeled substrategy at N among the available labeled substrategies. Let S_d be this substrategy and $V(S_d)$ its value. Let $v = \sup_d V(S_d)$. He then vetoes all the presently available labeled strategies that induce a substrategy at N except those of the S_d 's whose value is superior to $v - \theta$. Any labeled strategy vetoed by at least one of the nodes N of period t becomes unavailable at the earlier periods.

IF there is no labeled strategy left, *no consensus has been achieved.* END. ELSE (t-1) becomes the current period.

IF t=1, the root node is the unique N .The DM chooses the best strategy among the available labeled strategies. ELSE one goes back to the recursive step.

This process requires less computations then the first one, since a more drastic selection is operated at each step.

4 Conclusion

A decade ago, Machina and McClennen made an important theoretical point in showing that renouncing consequentialism opened the way to forms of rational behavior which did not conform to SEU theory. In particular, this allows for the use of an imprecise probability representation of uncertainty in decision models. How to put in practice these models is still an open question. In this paper, we have made a proposal, based on McClennen's resolute choice approach, to help the DM choose non-dominated strategies in dynamic decision problems. Although this proposal is only sketched and should be considered as tentative, it demonstrates that even with non-consequentialist behavior some rolling-back of decision trees can be helpful and make the model operational.

References

[1] Peter Hammond (1988). Consequentialist Foundations for Expected Utility. *Theory and Decision* 25:25-78.

[2] Edi Karni and Zvi Safra (1989). Ascending bid auctions with behaviorally consistent bidders, *Annals of Operations Research* 19:435-446.

[3] Mark J. Machina (1989). Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty", *Journal of Economic Literature* 27:1622-1688.

[4] Edward F. McClennen (1990). *Rationality and Dynamic Choice : Foundational Explorations.* Cambridge University Press, Cambridge.

[5] Eric Rasmusen (1989). *Games and Information: An introduction to Game Theory.* Blackwell, Cambridge.

[6] Paul Samuelson (1948). Consumption Theory in Terms of Revealed Preference, *Econometrica* 15:243-253.

[7] Peter Paul Wakker (1997). Are Counterfactual Decisions relevant for Dynamically Consistent Updating under Nonexpexted Utility. *Working paper,* Center, Tilburg University, The Netherlands.