



HAL
open science

Beliefs and Dynamic Consistency,
Jean-Marc Tallon, Jean-Christophe Vergnaud

► **To cite this version:**

Jean-Marc Tallon, Jean-Christophe Vergnaud. Beliefs and Dynamic Consistency,. Knowledge, Beliefs and Economics, Edward Elgar, pp.137-154, 2006. hal-00306458

HAL Id: hal-00306458

<https://hal.science/hal-00306458>

Submitted on 25 Jul 2008

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.

Beliefs and Dynamic Consistency

Jean-Marc Tallon* & Jean-Christophe Vergnaud†

March 2003

1 Introduction

In this chapter, we adopt the decision theoretic approach to the representation and updating of beliefs. The standard bayesian approach, as developed by Savage (1954), provides the conditions under which a decision maker whose preferences satisfy a certain number of “rationality” axioms (among which the famous “sure thing principle”) must have probabilistic beliefs. Thus, under these axioms, any decision maker is able (or acts as if he were able) to come up with a probability distribution over the states of nature (i.e., the sources of uncertainty) and, furthermore, acts as if maximizing an expected utility with respect to this prior. This construction underlies most of economic theory. However, as early as 1921, Keynes (1921) and Knight (1921) were already concerned that probabilistic beliefs might miss an important aspect, namely the confidence that agents have in their own beliefs. Doubts about the fact that probabilistic beliefs could represent any type of uncertainty were also raised by Shackle (1952). These (theoretical) concerns were backed by experimental evidence showing that subjects can have problems in dealing with probabilistic information in a way predicted by the theory (Allais (1953)), and more, drastically, had difficulties to come up with a probabilistic prior when no such information is a priori at hand (Ellsberg (1961)). In more recent years, more general axiomatic models of decision under uncertainty have emerged, in which agents could well have non-additive beliefs (Schmeidler (1989)) or multiple prior beliefs (Gilboa and Schmeidler (1989)). These models are precisely meant to deal with “Ellsberg paradox” while retaining enough structure so as to have some normative appeal on top of their better descriptive ability.

This paradox has been central in the development of decision theory under uncertainty and we hence briefly recall it here. Consider an urn in which there are 90 balls. The decision maker is told that there are 30 red balls and that the

*CNRS-EUREQua, jmtallon@univ-paris1.fr

†CNRS-EUREQua, vergnaud@univ-paris1.fr

remaining 60 balls can be black or yellow, in unknown proportion. The decision maker is then asked to choose among the following “acts” (decisions):

	Red	Black	Yellow
f	100	0	0
g	0	100	0
f'	100	0	100
g'	0	100	100

Act f is the act “betting on drawing a red ball”, g is “betting on black”, f' is “betting on red or yellow” and g' is “betting on black or yellow”. What the experiment shows is that a large majority of subjects prefer f to g while also preferring g' to f' . If one were to impose probabilistic beliefs on the agent, the first preference suggests that subjects think there are less than 30 black balls, and hence think that there is a probability less than $1/3$ of winning for the bet f' . On the other hand, the second preference suggests that they think that there are more black and yellow balls than red and yellow balls, thus revealing (subtracting the number of yellow balls) they think that there are more black balls than red balls, i.e., that the probability of drawing a black ball is larger than $1/3$. Hence, the behavior exhibited in this experiment is incompatible with agent having probabilistic beliefs, even in such a simple setting. The main reason is of course that subjects lack any precise probabilistic information on the number of yellow and black balls. Bayesian theory implies that this lack of “objective” knowledge is immaterial in the sense that agents should reveal a probabilistic beliefs, regardless of the information they have. Ellsberg experiments suggest that subjects are reluctant to do so. Furthermore, it is not *a priori* obvious that such reluctance should be deemed “irrational” and hence discarded on normative grounds. On the contrary, one might argue that in this case, the only information at hand is of the form of a set of probability distributions and should be taken as such by the decision maker, thus possibly making a behavioral difference compared to a situation in which this set is reduced to a singleton.

What this discussion suggests is that there is some latitude in modelling beliefs in a one-shot, static decision problem. Indeed, models with non probabilistic beliefs exist and represent behavior which, although violating the sure thing principle, is difficult to consider as irrational.¹ Things are actually a bit more complicated when moving to dynamic decision problems, as argued in Hammond (1988) and (1989).² Hammond’s argument is as follows: in dynamic choice problems, “weak” rationality requirements imply that beliefs have to be probabilistic. If not, the decision maker would behave in an inconsistent way. One

¹See for instance Schmeidler (1989) and Gilboa and Schmeidler (1989). For a survey, see Karni and Schmeidler (1991) and Cohen and Tallon (2000).

²For a recent review, see Hammond (1998a) and (1998b).

of the requirement is that the choice behavior should be *dynamically consistent*, in a sense to be made precise below, but which loosely speaking, requires the agent to be able to implement dynamically the choice made *ex ante* prior to any reception of information: if a certain course of action was judged *ex ante* to be optimal, the same course of action should be implemented after arrival of information. Hammond’s conclusion is then that the expected utility model of Savage (1954) together with Bayes’ rule of updating beliefs satisfies his rationality requirement.³ Actually, the result is even stronger: any model that represents behavior that does not respect the sure thing principle will fall under the criticism that it is not dynamically consistent.

In this chapter we take up this issue and propose a reconsideration of Hammond’s argument. After reviewing the argument more formally, we propose a weaker notion of dynamic consistency. We observe that this notion does not imply the full fledged sure thing principle thus leaving some room for models that are not based on expected utility maximization. However, these models still do not account for “imprecision averse” behavior such as the one exhibited in Ellsberg experiment and that is captured by non-Bayesian models such as the multiple prior model of Gilboa and Schmeidler (1989). We therefore go on with the argument and establish that such non-Bayesian models possess the weak form of dynamic consistency when the information considered consists of a reduction in imprecision (in the Ellsberg example, some information about the proportion of Black and Yellow balls).

2 Preliminary example: on the correct use of Bayes updating rule

Our aim in this section is to alert the reader, *via* a simple example, that the mere description of sources of uncertainty is as important as the fact of being able to put probability weights on them. Indeed, “well-behaved” bayesian decision makers might act in a non (Bayes-)rational way if they do not represent correctly the states of nature describing the uncertainty as well as the arrival of information.

The example goes as follows (there are many versions of it, we chose a “dramatic” one). There are three prisoners, a , b , and c in a cell; one of them has been sentenced to death. However, they don’t know which of the three has been sentenced. Prisoner a has managed to ask a warden which of the other two prisoners has not been sentenced. Suppose that the warden tells a that b has been recognized to be innocent. How does this information affect a ’s prior beliefs (that will be assumed realistically to be uniform $(1/3, 1/3, 1/3)$) on who is the

³For a somewhat different argument, see Epstein and LeBreton (1993).

sentenced prisoner?

A quick (and bayesian-like) reasoning would be to update beliefs by saying that the probability is now $1/2$ that a is the convicted prisoner. However, this reasoning is based on a misrepresentation of what the states of nature (i.e., the full description of all sources of uncertainty) are. In this example, the prisoner fails to understand that the messages that the warden can give are actually constrained. Indeed, a state of nature in this problem has to be given by a couple (sentenced prisoner, announcement of the warden). Thus, the possible states are $\{(a, b), (a, c), (b, c), (c, b)\}$. Prior beliefs of a over this set are $(1/6, 1/6, 1/3, 1/3)$ (assuming that he believes that the warden has equal chance to say b or c if he (a) is the one sentenced to death). If the warden says b is innocent, i.e., that the state is not (a, c) or (b, c) , then bayesian revision yields a 's new beliefs: with probability $1/3$ he is the sentenced prisoner, with probability $2/3$, c is the sentenced prisoner. With this treatment, the information revealed by the warden is not informative as to a 's probability of being convicted but teaches a lot about c 's probability.

The upshot of this example is simple: the description of what uncertainty really is, that is, the description of the space of states of nature, is crucial. While deceptively simple, this remark points to the fact that, in decision problem involving the potential arrival of new information, the description of a state must include the possible messages.

Now, this example can also be turned into a decision problem. Another version of the same example is that of a TV game-show. The announcer shows three doors to a player. There is, say, a car behind one of these doors. He then asks the candidate to pick one door. If the car is behind it, he wins it, otherwise he loses. After the candidate has announced which door he picked (say door a), the announcer tells him that he is going to open one of the remaining two doors, behind which there is no car, say door b . He then asks the candidate if he wants to change his choice and now pick door c , or if he wants to stick to his initial choice. Bayesian revision, as presented above, requires the candidate to now pick door c rather than sticking to a . Furthermore, the candidate should be able *ex ante* to actually announce all his strategies, as a function of which door is opened, i.e., he is able to plan in advance that if the announcer opens door b then he should switch to door c and *vice versa*. An important property of probabilistic beliefs, coupled with Bayes' revision rule, is that the *ex ante* plan and the choice that the decision maker actually makes after hearing the information dictate the same thing. This, loosely speaking, is the property of dynamic consistency that was put forward by Hammond (1988) as a justification for expected utility theory. We now turn to a more formal discussion of this principle.

3 Strong dynamic consistency implies Bayesianism

In this section, we examine more systematically the dynamic consistency argument in favor of Bayesian decision theory, using the conditional preferences to represent dynamic choice situations. We will illustrate the abstract discussion by the example of the value of information, which can be explained through a dynamic representation of choices (*via* decision trees) and an “equivalent” representation, casts in the language of static preferences over acts, associated to one-shot decisions.

We start with some notation. Let us consider the usual framework of Savage (1954), where Ω is the set of states of nature, \mathcal{A} is the set of acts or decisions which are functions from Ω into X , the set of consequences. An act hence specifies the “gain” of a decision in each state of nature: act f has payoff $f(\omega)$ in state $\omega \in \Omega$.

In this setting, information is captured by the fact that the agent might be told that a specific *event* has occurred. An event is a set of states of nature. Hence, the knowledge of an event enables the decision maker to disregard the complement event as being impossible. The preference-based approach to belief revision postulates that for any event $E \subseteq \Omega$, the decision maker has a preference relation conditional on this event.⁴ Such a conditional preference relation is denoted \succeq_E , interpreted as follows: $f \succeq_E g$ means that if the event E were true, then the agent would prefer act f to act g . What we are interested in is \succeq_Ω , which is the *ex ante* preference relation for which we want to examine whether rationality arguments force an expected utility type of preference. The relation \succeq_E can be seen as a potential *ex post* preference relation in the case where it happens that the event E is true. Denoting by E^c the complement of E , that is, the states that are not in E , one can also define a preference relation conditional on that event, \succeq_{E^c} . The following axiom captures some consistency requirement about how the *ex ante* preference relation relates to possible *ex post* preference relations, that is the agent’s preferences after he gets some information. It is at the heart of Hammond’s defense of expected utility based on dynamic principle.

Axiom (Strong dynamic consistency of preference).

$$\forall E \subseteq \Omega, \forall f, g \in \mathcal{A},$$

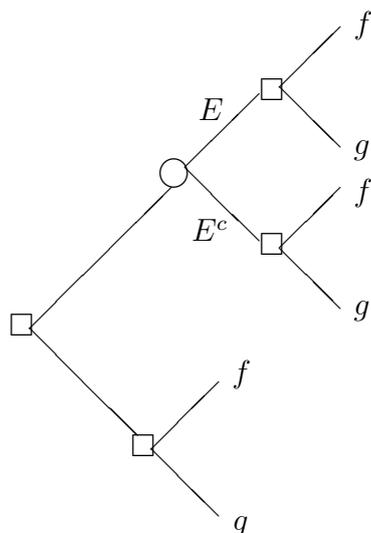
- if $f \succeq_E g$ and $f \succeq_{E^c} g$ then $f \succeq_\Omega g$,
- if $f \succeq_E g$ and $f \succeq_{E^c} g$ and at least one preference is strict (*i.e.*, $f \succ_E g$ or $f \succ_{E^c} g$) then we have $f \succ_\Omega g$.

⁴For a general approach to beliefs updating based on conditional preferences, see Wang (2002).

This axiom relates the *ex ante* preferences \succeq_{Ω} to possible *ex post* preferences preorders (conditional on E and E^c). It says that if f is preferred to g conditionally both to E and E^c , it should be also the case unconditionally. This axiom can be seen as a dominance axiom, preventing situations where we could have $f \succ_E g$ and $f \succ_{E^c} g$ while $f \prec_{\Omega} g$. Indeed, in such a case, there will be a clear contradiction between the *ex ante* preference and the *ex post* preferences leading potentially to problems in dynamic choice: the agent may have difficulty to carry on a decision plan if he suffers of contradicting interest between his initial preferences and the final ones.

Let us now examine more precisely the problem of dynamic choice for a specific case, namely that of information acquisition. Consider the situation where an agent is offered the possibility of postponing his choice after he gets some free information. It seems intuitive that in a pure problem of decision, the agent should not refuse such information, thus showing a “positive value of information”. To formalize this intuition, consider the following decision tree (figure 3) which represents the situation where an agent has to choose between f and g and can learn whether E or E^c is true before he chooses (as usual, squares represent decision nodes while circles represent chance nodes). If the decision maker chooses to go up at the beginning of the tree, this means he will first learn whether E or E^c occurred and then make a choice between f and g . On the other hand, if he chooses to go down, this means that he will choose between f and g without having any information on the event E .

Figure 1: Choosing to get informed



Observe that the two branches of the tree of figure 3 that start from the origin are equivalent to the two trees in figure 3. In the left part of figure 3, the simple choice between f and g without any information is shown. In the right

part, we represent the equivalent strategic form of a tree in which the decision maker has the possibility to get informed, and more specifically to learn whether E occurred or not. The act f_{EG} should be read as follows: if the true states is in E , then get $f(\omega)$, while if it is in E^c , then get $g(\omega)$. Hence, f_{EG} is the act that yields f if E and g if E^c .

Figure 2:



The obvious advantage of getting this information is that the decision maker can condition his choice on the realization of E . In other words, when comparing the two decision trees of figure 3, the decision maker will always prefer to act in the right hand side tree in which the available strategies include the one of the left hand side tree.

In the choice situation just described, an agent satisfying the *Strong dynamic consistency axiom*, will be willing to accept the information, which, as argued above is a sensible choice (if not the only sensible choice). Indeed, consider the case where $f \succ_E g$ while $g \succ_{E^c} f$. If he chooses to go up in the tree of figure 3, he will choose f conditionally on E and g conditionally on E^c . Therefore, going up is equivalent to get the act f_{EG} . Yet, since $f_{EG} \sim_E f$ and $f_{EG} \sim_{E^c} g$ we have $f_{EG} \sim_E f \succ_E g$ and $f_{EG} \sim_{E^c} g \succ_{E^c} f$.⁵ As a consequence, the *Strong dynamic consistency axiom* implies that $f_{EG} \succ_{\Omega} f$ and $f_{EG} \succ_{\Omega} g$ and thus the agent should find valuable this information.

More systematically, it can be shown that the agent should always accept to be informed. If it were not the case, there would be instances where $f_{EG} \sim_E f \succ_E g$, $f_{EG} \sim_{E^c} g \succ_{E^c} f$ while $f \succ_{\Omega} f_{EG}$ which constitutes a violation of the *Strong dynamic consistency axiom*. Just as before, to go up in the tree of figure 3 means to get the act f_{EG} but there, this perspective is negative since by going down the agent could get f . The agent should refuse the information exhibiting a negative value of information. This phenomenon might happen in non expected utility models like that of Gilboa and Schmeidler (1989). We give such an example in an appendix.

⁵There is a hidden assumption behind the fact that $f_{EG} \sim_E f$, namely that the decision maker, when informed that E has occurred, considers that what might have happened had E^c occurred is irrelevant. This is a form of separability of preferences that we assume here and on which we will comment briefly in the conclusion.

We now turn to the restrictions that the *Strong dynamic consistency axiom* place on preferences, focussing on its implications concerning a corner-stone of Bayesian decision theory, the *Sure thing principle*, which we now recall.

Axiom (Sure thing principle).

$\forall E \subseteq \Omega, \forall f, g, f', g' \in \mathcal{A}$. If $f = f'$ and $g = g'$ on E , $f = g$ and $f' = g'$ on E^c then $f \succeq_{\Omega} g$ if and only if $f'_{\Omega} \succeq_{\Omega} g'$.

This axiom can be explained on the Ellsberg experiment presented in the introduction. Take the event E to be the event $\{Black, Red\}$. Then, on this event, act f and act f' are the same, and act g and act g' are the same. On E^c , that is on $\{Yellow\}$, f and g are equivalent, and f' and g' are equivalent. The sure thing principle then states that if f is preferred to g then f' should also be preferred to g' : indeed the only difference between the comparison of f versus f' and the comparison of g versus g' is to be found in the column *Yellow*, i.e., in the event the ball drawn is yellow. But conditionally on this event f and g are the same and f' and g' are the same. Hence, concludes the sure thing principle, the choice should not depend on what is the payoff in this event, as long as it is the same for f and g , and for f' and g' . Obviously, the choice pattern observed in Ellsberg experiment violates this principle.

Now, it can be shown that the *Sure thing principle* is a consequence of the *Strong dynamic consistency axiom*. Hence, adhering to the strong dynamic consistency axiom prevents Ellsberg type of behavior.

Proposition. *Strong dynamic consistency of preference implies that \succeq_{Ω} satisfies the Sure thing principle.*

As previously mentioned, the *Sure thing principle* is at the heart of the subjective expected utility model. Conversely, note that if we update subjective probabilities according to Bayes rule, then the subjective expected utility model satisfies the *Strong dynamic consistency axiom*. Therefore, if we accept the *Strong dynamic consistency axiom* then we do not have much latitude for the kind of preferences we can consider: they essentially have to conform to the subjective expected utility model. However, we will see in the following that this axiom is a bit too strong and that there is room for relaxing it in a sensible way.

4 Weak dynamic consistency of preference

We saw in the previous section that the rationale for the *Strong dynamic consistency axiom* was to ensure that an agent will not fear preferences conflicts in some dynamic choices and particularly, will exhibit a positive value of information. Let us examine now the following weaker axiom.

Axiom (Weak dynamic consistency of preference).

$\forall E \subseteq \Omega, \forall f, g \in \mathcal{A}$. If $f \succeq_E g$ and $f \succ_{E^c} g$ then $f \succeq_\Omega g$.

It is weaker than the *Strong dynamic consistency axiom* in the sense that this axiom permits situation where $f \sim_E g$, $f \succ_{E^c} g$ and $f \sim_\Omega g$. It seems that this axiom could do the job as well. Indeed, it also prevents the kind of reversal where $f \succ_E g$ and $f \succ_{E^c} g$ while $f \prec_\Omega g$, that we identified as being problematic.

It is worthy of interest to note that in the particular case of the information acquisition problem, this axiom is actually enough to guarantee that the agent values information positively. Indeed, in the situation where $f \succ_E g$ while $g \succ_{E^c} f$, the *Weak dynamic consistency axiom* implies that $f_{Eg} \succeq_\Omega f$ and $f_{Eg} \succeq_\Omega g$ and thus the agent will find this information valuable.

Still, one may find that it is a rather weak departure from the strong form of dynamic consistency: playing on weak and strict preferences might appear somewhat anecdotal. Note however that an important consequence of this apparently mild weakening is that the *Sure thing principle* cannot be deduced from the *Weak dynamic consistency axiom*. Hence, it opens the door to a weakening of the sure thing principle that might be of interest if one wants to be able to accept preferences that do not belong to the subjective expected utility class. More precisely, consider the following weakening of the sure thing principle.

Axiom (Weak sure thing principle).

$\forall E \subseteq \Omega, \forall f, f', g, g' \in \mathcal{A}$, if $f = f'$ and $g = g'$ on E , $f = g$ and $f' = g'$ on E^c , then $f \succ_\Omega g$ implies that $f' \succeq_\Omega g'$.

It can be shown that the *Weak sure thing principle* is a consequence of the *Weak dynamic consistency axiom*.

Proposition. *Weak dynamic consistency of preference implies that \succeq_Ω satisfies the Weak sure thing principle*

So far, we have weakened dynamic consistency in a rather mild way and showed that it lead to a weakened version of the sure thing principle. This would be of not much interest if the preferences that satisfy the weak sure thing principle but not the full fledged sure thing principle were of little interest. We now want to argue that this is not the case and that it includes an interesting set of preferences, in which the decision maker treats the information in a purely qualitative way. We present this class of models in the following section.

5 Possibilistic decision model

Zadeh (1978) introduced in the field of artificial intelligence a representation of uncertainty based on possibility measures (see also Dubois and Prade (1985)).

This work has been extended to a decision theoretic setting by Dubois, Prade, and Sabbadin (2000). Possibility measures were introduced to represent situations in which the information is qualitative as opposed to probabilities that are by nature quantitative. This notion of qualitative information can be seen as the formal representation of ideas introduced in economics by Shackle (1952).

Definition (Possibility measure).

A possibility measure is a set function $\Pi : 2^\Omega \rightarrow [0, 1]$, such that $\Pi(\emptyset) = 0$, $\Pi(\Omega) = 1$ and

$$\forall A, B \subseteq \Omega, \quad \Pi(A \cup B) = \text{Max}(\Pi(A), \Pi(B))$$

The interpretation of the possibility of event A is essentially to what degree the event is considered possible or not. It is qualitative in the sense that statements like “this event is two times as possible as that event” do not have a well defined meaning. Indeed, the definition makes it clear that the union of two events is “as possible as” the most possible of the two events. Observe that many events could be assigned the maximal possibility, i.e., 1 without any contradiction. In particular, for any event E , it must be the case that E or E^c has possibility 1, since $E \cup E^c = \Omega$. Possibility measures can be used in decision problems if one manages to come up with a decision criterion that is based on such objects. In much the same way as probabilities are naturally combined with utilities in the expected utility model, possibility measures can be combined with utility *via* the so-called Sugeno integral. Let $u : X \rightarrow [0, 1]$ be a utility function. It is possible to axiomatize a decision criterion that is based on the combination *via* min and max operations of the possibilistic representation of uncertainty and the utility of the consequence of an act. More precisely, an act f is evaluated by

$$\max_{\omega \in \Omega} [\min(\Pi(\omega), u(f(\omega)))]$$

The Sugeno integral, which is a generalization of the above formula, can be seen as a “regular” integral in which the min operator replaces the product operator while the max operator replaces the sum operator. If we consider simple acts like bets, it is easy to see that betting on the event A (i.e., win if A is observed loose if not), is simply evaluated by the possibility of A , $\Pi(A)$.

One can check that the Sugeno integral with possibility satisfies the *Weak sure thing principle*. Furthermore, if the possibility measure is updated according to the following rule called the non-normalized Bayes rule

$$\Pi(A|E) = \Pi(A \cap E)$$

then it can be checked that this model satisfies the *Weak dynamic consistency axiom* as well and hence will never be subject to problems such as the negative value of information.

More generally, it is possible to extend this model in order to mix the expected utility model and the Sugeno integral in a general criterion which also satisfies the *Weak sure thing principle*. However, this kind of model, while generalizing the expected utility model, still can not explain the Ellsberg paradox as the latter constitutes a violation of the *Weak sure thing principle* as well. Since, as recalled in the introduction, the Ellsberg paradox is one of the main motivation for looking for non-Bayesian models, it might appear that the kind of possibilistic models we have just examined, while of theoretical interest, are somehow missing an important point of what imprecise information, as captured in Ellsberg experiment, is about.

6 The content of information

In this section, we explore properties of the *Weak dynamic consistency axiom* further, especially when coupled with information of a different kind than the one we dealt with so far. We do so by going back to the original motivation we gave for looking at non-Bayesian models, namely to capture imprecision or ambiguity in the data provided to the decision maker. For instance, in the Ellsberg paradox, the urn's content is incompletely described: there is a set of probability distributions that are compatible with the objective data. Now, the previous sections showed that we might run into trouble with respect to the dynamic consistency requirement in non-Bayesian decision models when the decision maker might learn whether the ball which was drawn in the urn is yellow or not. Arguably, non-Bayesian models are not meant to deal with this type of learning. Rather, it is easy to imagine other forms of information and thus other kinds of learning. For instance, the decision maker might learn that there is less than 20 yellow balls in the urn. This operation would amount to a decrease of the imprecision in the data the agent has.

Another, related, kind of learning, which fits well with non-Bayesian models is sampling. Suppose the agent is allowed to sample in the urn before he chooses and the final drawing is made. For instance if he observes 2 red balls, 7 black balls and 1 yellow balls, he will learn something on the urn. Intuitively, this kind of learning which also reduces the imprecision about the probabilistic data will probably be valuable to a decision maker who is adverse to such imprecision.

In the remainder of this section, we suggest a rather abstract argument to establish that in these situations, *Weak dynamic consistency axiom* can be satisfied. Since we want to consider information that bears on the imprecision of probabilistic data, it is convenient to enlarge the description of the states of nature in order to explicitly capture this dimension. For instance, assume that the set Ω is the cross product of two sets: $\Omega = \Theta \times S$. Assume furthermore that the objective data about this problem is that, conditionally on a realization

$\theta \in \Theta$, there is a probability distribution p_θ over S but, on the other hand, there is no available probability distribution over Θ . Loosely speaking, this represents a world with two layers of uncertainty: a first layer, on which one has very little data or information, the set Θ , which captures general features about the world and a second layer, S , on which, conditionally on the realization of a general feature θ , there is a probability distribution. For instance, in the Ellsberg's urn example, an element $\theta \in \Theta$ is the description of the content of the urn while $s \in S$ is the color of the ball drawn.

Beyond the anecdotal problem of the Ellsberg experiment, this formalization also captures features of problems of scientific uncertainty. For instance, in the climate change problem, Θ can be understood as a list of hypothesis and predictions provided by the scientific community. $\theta \in \Theta$ could be the level of temperature increase in 50 years. S is the description of natural catastrophic events such as storm, flood... Conditionally on θ we may know p_θ , which gives the frequency of catastrophic events if temperature has increased by a given amount: e.g., if the temperature raises by 2°C, then a type of storm that occurred every 100 years in the past will now occur every 10 years.

Consider now the following (non-Bayesian) decision criterion, known as the multiple-prior model, which was axiomatized by Gilboa and Schmeidler (1989).⁶ An act f , which is a mapping from $\Theta \times S$ to the space of consequences X , is evaluated by

$$\min_{\theta \in \Theta} [E_{p_\theta} u(f(\theta, s))]$$

The distinction between the two kinds of learning we suggested can easily be embodied in this formalization. Indeed, one can oppose learning about S (learning about the color of the ball) and learning about Θ (learning about the content of the urn). To learn about S is to consider an event $E \subseteq \Omega$ of the form $E = \Theta \times F$ where $F \subseteq S$. The multiple-prior criterion clearly does not satisfy the *Weak sure thing principle* with respect to such an event. However, if we consider an event $E \subseteq \Omega$ of the form $E = \Phi \times S$ where $\Phi \subseteq \Theta$, then the multiple-prior criterion does satisfy the *Weak sure thing principle*. Moreover, if we suppose that $f \succeq_{\Phi \subseteq \Theta} g$ if and only if

$$\min_{\theta \in \Phi} [E_{p_\theta} u(f(\theta, s))] \geq \min_{\theta \in \Phi} [E_{p_\theta} u(g(\theta, s))]$$

one can easily check that the *Weak dynamic consistency axiom* can be satisfied.

Therefore, for the type of learning we have introduced, and that we could call *reduction of imprecision*, there is no problem of dynamic consistency. Thus, one is lead to conclude from this that when modelling the arrival of information, one should be careful as for which type of information it is and whether the underlying

⁶This model is used in the appendix to illustrate how information, as modelled in the previous sections, might have a negative value for non-Bayesian decision makers.

decision model is adapted to treat that particular type of information. One can therefore look back at Hammond's result as saying that the type of information that is taken into account goes well with the expected utility criterion. The result of this section, on the other hand, leads to the conclusion that non-Bayesian decision models such as the multiple prior model is tailored to treat arrival of information that refines the non-probabilistic initial data and leads to a reduction in imprecision.

Remark that in the perspective of formalizing scientific uncertainty, reduction of imprecision can be interpreted as the usual evolution of science. It does not make sense to attribute a probability of validity to a scientific hypothesis and, when it happens to be refuted, the theory resting on it is dropped from the scientific knowledge.

7 Conclusion

In this chapter, we reviewed Hammond's argument founding expected utility on a dynamic consistency principle. We weakened Hammond's *Strong dynamic consistency axiom* and showed that this weakened version is compatible with a qualitative treatment of uncertainty through the possibilistic decision model. However, this model does not account for the type of imprecision adverse behavior that we discussed in the introduction. We then enlarged the discussion and showed that non-Bayesian representation of imprecision adverse behavior admit a form of dynamic consistency property with respect to the arrival of a particular form of information. Indeed, reduction in imprecision (such as rejection of scientific hypothesis) will not lead to dynamically inconsistent behavior in these models, even though the sure thing principle is not satisfied.

Another stream of research that we will simply mention attacked another aspect of Hammond (1988) and (1989) argument. Indeed, implicit⁷ in our discussion was another principle put forward by Hammond (1988), namely *consequentialism*. Consequentialism, loosely speaking, is a separability principle that states that at a given decision node in a tree, the decision maker's choice should depend only on the continuation of the tree from that node on. The part of the tree that was not reached (and cannot be reached anymore) could simply be dropped without changing the choice of the decision maker. This principle, which allows one to perform backward induction and dynamic programming, has been maintained in the approach we followed here. Dropping it would be a way to break the implication from the dynamic consistency requirement to the sure thing principle and hence consider non-Bayesian models. However, if one were to drop consequentialism, as some authors like Machina (1989) or Mac Clennen

⁷Although see the remark in footnote 5.

(1990) did, one has to specify how the decision maker analyzes a decision tree. A way that has been explored in the literature is to assume that the decision maker is resolute enough to follow his initial intentions. Hence, the equivalence between the optimal choice in the reduced strategic form of a tree and the choice operated in the dynamic decision tree is postulated. If such a postulate is taken for granted, dropping the sure thing principle is of no real consequence as far as dynamic properties of choices are concerned. Machina (1989) and Mac Clennen (1990) propose different means of doing so.⁸ Another approach to this problem is to consider that the decision maker is composed of different selves (one at each point in time and for each realization of uncertainty) and the final outcome reached is some sort of Nash equilibrium of the game played by the different selves. This so-called “sophisticated behavior” has been introduced by Strotz (1955) and is compared with resolute choice in Machina (1989). In this paper, we have taken a different route to the problem posed by dynamic inconsistency, arguing that they might be the consequence of an inappropriate way of modelling information in various contexts.

⁸For yet another approach, see Gul and Pesendorfer (2001).

Appendix

Negative value of information with the multiple prior expected utility criterion

Non bayesian decision criterion might lead to a negative value of information, when coupled with the idea that agents analyze trees by backward induction and use a revision rule to determine their beliefs at each point of the tree.⁹ This is for instance the case for the so-called multiple prior model in Gilboa and Schmeidler (1989). In this model the agent's beliefs are given by a set of probability measures. The decision criterion is then to choose the act that gives the highest minimal expected utility (where the minimum is taken over the set of probability representing the agent's beliefs).

We present here a slightly more complicated version of the type of behavior exhibited in Ellsberg's experiments. There are 100 balls in an urn and the decision maker is given the following information: 60 balls are marked *A* or *D* in an unknown proportion, while 20 are marked *B* and 20 are marked *C*. Thus, the set of probability measures (p_A, p_B, p_C, p_D) representing the decision maker's beliefs is given by $\{(p, .2, .2, .6 - p) | p \in [0, .6]\}$.

Assume that the decision maker's set of available actions is $\{f_{AUD}, f_B, f_C\}$, i.e., bet on *A* or *D*, bet on *B*, bet on *C* with the following payoffs:

	A	B	C	D
f_{AUD}	1	0	0	1
f_B	0	1	0	0
f_C	0	0	1	0

Now assume that the decision maker has access to some information, that will tell him whether the ball is marked *A* or *B*, or whether it is marked *C*, or whether it is marked *D*. Formally, the decision maker is told whether $\{A, B\}$, or $\{C\}$ or $\{D\}$ occurred. Using this information, the decision maker has now access to a richer set of strategies if he conditions the available bets on the information received. For instance, the strategy, bet on *B* if told $\{A, B\}$, bet on *C* if told $\{C\}$, bet on *D* if told $\{D\}$ will amount to f_{BUCUD} . One can check that the set of available choices with information, in the strategic form decision tree is

$$\{f_{AUD}, f_B, f_C, f_{AUCUD}, f_{BUCUD}, f_{CUD}, f_{BUD}, f_{BUC}\}$$

When asked for his optimal choice in this set, the decision maker will say f_{AUCUD} . Indeed, this bet gives him a 80% chance of winning. The other bets have smaller chances to win, when evaluated by the minimal probability. For instance, f_{BUCUD} has only a minimal probability of winning of 40%. Hence, the decision maker

⁹In non bayesian models, several revision rules can be introduced. We focus on one such rule here but essentially, a similar argument could be made for any rule.

will want to use the information in the sense that he chooses a bet which is contingent on the information received.

Let's now look at the dynamic choices associated to this story. Assume that the decision maker revises his beliefs by updating in a bayesian fashion all the probability measures in his sets of priors. For instance, conditionally on the fact that the ball is marked A or B , the beliefs are given by the set $\left\{ \left(\frac{p}{.2+p}, \frac{.2}{.2+p}, 0, 0 \right) \mid p \in [0, .6] \right\}$ of probability distribution over $\{A, B, C, D\}$.

Hence, at the node where he receives the information $\{A, B\}$, the decision maker, having to choose among $\{f_{AUD}, f_B, f_C\}$, will choose f_B , since its minimal expected value (with respect to the set of updated priors) is .8, while f_{AUD} yields a minimal expected value of 0. Trivially, when the information is $\{C\}$, the decision maker will have (revised) beliefs given by $(0, 0, 1, 0)$ and will choose f_C , while if given the information $\{D\}$, revision of beliefs leads to the beliefs $(0, 0, 0, 1)$, and the act f_{AUD} is thus optimal. Hence, the optimal choice is:

$$\left\{ \begin{array}{ll} \text{If } \{A, B\}, & \text{choose } f_B \\ \text{If } \{C\}, & \text{choose } f_C \\ \text{If } \{D\}, & \text{choose } f_{AUD} \end{array} \right.$$

This choice amounts to choose f_{BUCUD} in the equivalent strategic form. Hence, the choice the decision maker would make in the dynamic decision tree when he decides to acquire the information would be f_{BUCUD} , which is dominated in the equivalent strategic form by f_{AUCUD} . Furthermore, if the decision maker were to refuse the information (i.e., go down in the analogue of the tree of figure 3 for the case at hand), his optimal choice would be f_{AUD} . To conclude the argument, observe that, f_{AUD} gives a 60% chance of winning, while the minimal probability of winning with f_{BUCUD} is only of 40%. Hence, according to the criterion of the maximal minimal expected utility, f_{AUD} is preferred to f_{BUCUD} . Thus, the decision maker will prefer to remain uninformed, whereas the information is potentially useful (as defined in section 2, i.e., the decision maker would condition his choice on such a piece of information in the equivalent strategic forms). The information can then be said to have a "negative" value for him.

References

- ALLAIS, M. (1953): “Le comportement de l’homme rationnel devant le risque: critique des postulats de l’école américaine,” *Econometrica*, 21, 503–546.
- COHEN, M., AND J.-M. TALLON (2000): “Décision dans le risque et l’incertain : l’apport des modèles non-additifs,” *Revue d’Economie Politique*, 110, 631–681.
- DUBOIS, D., AND H. PRADE (1985): *Théorie des possibilités*. Masson, Paris.
- DUBOIS, D., H. PRADE, AND R. SABBADIN (2000): “Qualitative decision theory with Sugeno Integrals,” in *Fuzzy measures and Integrals. Theory and Applications.*, ed. by M. Grabisch, T. Murofushi, and M. Sugeno, pp. 314–322. Physica-Verlag.
- ELLSBERG, D. (1961): “Risk, ambiguity, and the Savage axioms,” *Quarterly Journal of Economics*, 75, 643–669.
- EPSTEIN, L., AND M. LEBRETON (1993): “Dynamically consistent beliefs must be Bayesian,” *Journal of Economic Theory*, 61(1), 1–22.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin expected utility with a non-unique prior,” *Journal of Mathematical Economics*, 18, 141–153.
- GUL, F., AND W. PESENDORFER (2001): “Temptation and self control,” *Econometrica*, 69, 1403–1435.
- HAMMOND, P. (1988): “Consequentialist foundations for expected utility,” *Theory and Decision*, 25, 25–78.
- (1989): “Consistent plans, consequentialism, and expected utility,” *Econometrica*, 57(6), 1445–1449.
- (1998a): “Objective Expected Utility: A Consequentialist Perspective,” in *Handbook of Utility Theory*, ed. by S. Barbera, and C. Seidl. Kluwer, chapter 5.
- (1998b): “Subjective Expected Utility,” in *Handbook of Utility Theory*, ed. by S. Barbera, and C. Seidl. Kluwer, chapter 6.
- KARNI, E., AND D. SCHMEIDLER (1991): “Utility theory with uncertainty,” in *Handbook of Mathematical Economics*, ed. by W. Hildenbrand, and H. Sonnenschein, vol. IV, chap. 33, pp. 1763–1831. North-Holland.
- KEYNES, J.-M. (1921): *A treatise on probability*. London: McMillan.
- KNIGHT, F. (1921): *Risk, uncertainty and profit*. Houghton Mifflin.

- MAC CLENNEN, E. (1990): *Rationality and Dynamic Choice Foundational Explorations*. Cambridge University Press.
- MACHINA, M. (1989): “Dynamic consistency and non-expected utility models of choice under uncertainty,” *Journal of Economic Literature*, 28, 1622–1668.
- SAVAGE, L. (1954): *The foundations of statistics*. New-York, John Wiley.
- SCHMEIDLER, D. (1989): “Subjective probability and expected utility without additivity,” *Econometrica*, 57(3), 571–587.
- SHACKLE, G. (1952): *Expectation in economics*. Cambridge University Press.
- STROTZ, R. (1955): “Myopia and inconsistency in dynamic utility maximization,” *Review of Economic Studies*, 23, 165–180.
- WANG, T. (2002): “Conditional preferences and updating,” Discussion paper, University British Columbia, forthcoming *Journal of Economic Theory*.
- ZADEH, L. (1978): “Fuzzy sets as a basis for a theory of possibility,” *Fuzzy sets and systems*, 1, 3–28.