Three scenarios for the revision of epistemic states *

Didier Dubois
IRIT-CNRS
Université Paul Sabatier
Toulouse, France
dubois@irit.fr

Abstract

This position paper discusses the difficulty of interpreting iterated belief revision in the scope of the existing literature. Axioms of iterated belief revision are often presented as extensions of the AGM axioms, upon receiving a sequence of inputs. More recent inputs are assumed to have priority over less recent ones. We argue that this view of iterated revision is at odds with the claim, made by Gärdenfors and Makinson, that belief revision and non-monotonic reasoning are two sides of the same coin. We lay bare three different paradigms of revision based on specific interpretations of the epistemic entrenchment defining an epistemic state and of the input information. If the epistemic entrenchment stems from default rules, then AGM revision is a matter of changing plausible conclusions when receiving specific information on the problem at hand. In such a paradigm, iterated belief revision makes no sense. If the epistemic entrenchment encodes prior uncertain evidence and the input information is at the same level as the prior information and possibly uncertain, then iterated revision reduces to prioritized merging. A third problem is one of the revision of an epistemic entrenchment by means of another one. In this case, iteration makes sense, and it corresponds to the revision of a conditional knowledge base describing background information by the addition of new default rules.

Introduction

The interest in belief revision as a topic of investigation in artificial intelligence was triggered by Gärdenfors (1988) book and the axiomatic approach introduced by C. Alchourrón, P. Gärdenfors and D. Makinson (1985) in the setting of propositional logic. This approach assumes that the set of accepted beliefs held by an agent is a deductively closed set of propositions. On this basis, axioms of belief change (revision, but also contraction) formulate constraints that govern the "flux" of information, i.e. that relate one belief set to the next one upon receiving a new piece of information. An important assumption is that belief revision takes place in a static world, so that the input information is supposed to bring insight to a case that the agent deals with, but is never meant to indicate that the world considered by the agent receiving it has evolved.

*This position paper was triggered by discussions with Jerome Lang and Jim Delgrande at a Belief Revision seminar in Dagstuhl, in August 2005

The crucial point of the AGM theory is that the axiomatic framework enforces the existence of a so-called epistemic entrenchment relation between propositions of the language. This relation acts like a priority assignment instrumental to determine the resulting belief set after revision. It is also similar (even if purely ordinal) to a probability measure. More specifically, an epistemic entrenchment is a complete preordering between propositions which looks like a comparative probability relation (Fishburn 1986), even if it has different properties. Properties of an epistemic entrenchments make it expressible in terms of a complete plausibility ordering of possible worlds, such that the resulting belief set after receiving input A is viewed as the set of propositions that are true in the most plausible worlds where A holds.

The AGM theory leaves the issue of iterated revision as an open problem. Since then, iterated revision has been the topic of quite a number of works (Nayak 1994), (Williams 1995), (Darwiche & Pearl 1997), (Lehmann 1995), (Jin & Thielscher 2005). However it also seems to have created quite a number of misunderstandings, due to the lack of insight into the nature of the problem to be solved.

A typical question that results from studying the AGM theory is: What becomes of the epistemic entrenchment after the belief set has been revised by some input information? Some researchers claimed it was simply lost, and that the AGM theory precludes the possibility of any iteration. Others claimed that it changes along with K, and tried to state axioms governing the change of the plausibility ordering of the worlds, viewing them as an extension of the AGM axioms. This trend led to envisage iterated belief revision as a form of prioritized merging where the priority assignment to pieces of input information reflected their recency.

However, this notion of iterated belief revision seems to be at odds with Gärdenfors and Makinson (1994) view of belief revision as the other side of non-monotonic reasoning, where the epistemic entrenchment relation is present from the start and describes the agent's expectations in the face of the available evidence. Such an epistemic entrenchment may also derive from the analysis of a set of conditionals, in the style of (Lehmann & Magidor 1992), yielding a ranking of worlds via the so-called rational closure. The revised belief set is then the result of a simple inference step of conditionals from conditionals, whereby propositional conclusions tentatively drawn are altered by the arrival of new pieces of evidence. In this framework, the conditional information, hence the plausibility ordering, is never revised and iteration comes down to inference of new conclusions and dismissal of former ones, in the spirit of non-monotonic reasoning.

Solving the clash of intuitions between iterated revision and non-monotonic reasoning leads us to considering that the AGM view of belief revision (related to non-monotonic reasoning) has more to do with inference under incomplete information than with iterated revision as studied by many subsequent researchers (see a critical discussion of Darwiche and Pearl(1997) axioms along this line in (Dubois, Moral, & Prade 1998)). Two settings for revision, namely revision as defeasible inference, and revision as prioritized merging emerge, that deal with distinct problems.

This note is also in the spirit of a former position paper by Friedman and Halpern (1996a). In that note, they complain that iterated belief revision research relies too much on the finding of new axioms justified by toy-examples, and representation results, while more stress should be put on laying bare an appropriate "ontology", that is, describing a concrete problem or scenario that iterated revision is supposed to address. Friedman and Halpern suggest two such ontologies, that basically differ by the meaning of the input information. According to the first one, the agent possesses knowledge and beliefs about the state of the world, knowledge being more entrenched than beliefs, and receives inputs considered as true observations. This view is similar to a form of conditioning in the sense of uncertainty theories. In the other scenario, the input information is no longer systematically held for true and competes with prior beliefs, thus corresponding to a kind of merging bearing much similarity to the combination of uncertainty in the theory of evidence (Shafer 1976).

In this paper, we somewhat pursue this discussion by pointing out that the status of the epistemic entrenchment itself may also be understood differently: in some scenarios, it represents background information about the world, telling what is normal from what it is not, in a refined way. In that case, the plausibility ordering underlying the epistemic entrenchment is similar to a statistical probability distribution, except that the underlying population is ill-specified, and statistical data is not directly accessible. In other scenarios, the plausibility ordering expresses beliefs about unreliable observations about the solution to a problem at hand, the pieces of evidence gathered so far from witnesses on a whodunit case, for instance. In the latter situation, the resulting epistemic entrenchment is fully dependent on the case at hand and has no generic value.

It leads to propose three change problems that have little to do with each other even if they may share some technical tools. If we take it for granted that belief revision and non-monotonic reasoning are two sides of the same coin and if we rely on technical equivalence results between Lehmann and Magidor(1992) conditional logic under rational closure, and the AGM theory, then we come up with a qualitative counterpart of statistical reasoning, with inputs taken as incomplete but sure information about a case at hand. We call it Belief Revision as Defeasible Inference (BRDI). On the other hand, if we take it for granted that the epistemic entrenchment gathers uncertain evidence about a case, likely to evolve when new uncertain pieces of evidence are collected, we speak of Belief Revision as Prioritized Merging (BRPM). Finally, we consider the situation where our background knowledge is modified by new pieces of knowledge, whereby states of fact that we used to think as normal turn out not to be so, or conversely. We then speak of Revision of Background Knowledge by Generic Information (RBKGI). In the latter case, inputs take the form of conditionals.

It may be that other scenarios for belief change could be pointed out. However, we claim that iterated revision in each of the above scenarios corresponds to very different problems. A companion paper (Delgrande, Dubois, & Lang 2006) proposes a formal framework for the BRPM situation in full details. Here, we propose an informal comparative discussion of the three scenarios.

Belief Revision as Defeasible Inference (BRDI)

In the first setting, the AGM theory and non-monotonic reasoning are really regarded as two sides of the same coin. However, while in the AGM approach, only a flat belief set denoted K, composed of logical formulas, is explicitly available (since the epistemic entrenchment is implicit in the axioms of the theory), the nonmonotonic logic approach lays bare all the pieces of information that allows an agent to reason from incomplete reliable evidence and background knowledge. While in the AGM paradigm, the primitive object is the belief set, in the following, everything derives from conditional information, synthetized in the form of a partial ordering of propositions, and the available evidence. This view is fully developed by Dubois Fargier and Prade (2004) (2005) as a theory of accepted beliefs.

In the following, we consider a classical propositional language, and we do not distinguish between logically equivalent propositions. Hence, we consider propositions as subsets of possible worlds, in other words, events (to borrow from the probabilistic literature). The influence of syntax on revision is out of the scope of this paper. Under such a proviso, it is assumed that the agent's epistemic state is made of three components:

1. A confidence relation, in the form of a partial ordering $\succ$ on propositions $A, B, \ldots$ expressed in a given language. This relation, which should be in agreement with logical deduction, expresses that some propositions are more normally expected (or less suprizing) than others. It encodes the background information of the agent, which describes how (s)he believes the world behaves in general. It reflects the past experience of the agent. Such a confidence relation may directly stem from a set of conditionals Δ. Δ contains pieces of conditional knowledge of the form $A \to B$ where $\to$ is a nonclassical implication, stating that in the context where all that is known is A, B is generally true. Each such conditional is then encoded as the constraint $A \wedge B \succ A \wedge \neg B$, understood as the statement that $A \wedge B$ is generally more plausible (that is, less surprizing) than $A \wedge \neg B$ (Friedman & Halpern 1996b). A plausibility ordering of worlds $\geq_\pi$ can be de-

rived from such constraints via some information minimization principle (like rational closure of Lehmann and Magidor (1992), or equivalently, the most compact ranking compatible with the constraints(Pearl 1990), or yet the principle of minimal specificity of possibilistic logic (see (Benferhat, Dubois, & Prade 1997) for instance).

2. A set of contingent observations concerning a case of interest for the agent, under the form of a propositional formula A. The observations are sure evidence about this case, not general considerations about similar cases. Such pieces of evidence are sure facts (or at least accepted as such), hence consistent with each other. It means that a preliminary process is capable of handling conflicting observations and come up with a consistent report.

3. The belief set $K * A$ of the agent. It is made of propositions tentatively accepted as true by the agent about the case, in the face of the current observations. Propositions in $K * A$ are inferred from the observations and the background knowledge (so it is not an independent part of the epistemic state). K is the belief set of the agent before hearing about A. That input information is safe explains why the success postulate ($A \in K * A$) makes sense.

For instance consider a medical doctor about to diagnose a patient. It is assumed that the aim is to determine what the patient suffers from within a time-period where the disease does not evolve. The plausibility ordering reflects the medical knowledge of the medical doctor in general. Before seeing the patient, (s)he may have some idea of which diseases are more plausible than others. Observations consist of reports from medical tests and information provided by the patient on his state of health. The resulting belief set contains the diagnosis of the patient that will be formulated by the doctor on the basis of the available observations. This belief set concerns the patient, not people's health in general.

Formally, under this view, the original belief set K is inferred from Δ, or from $\succ$, or from $\geq_\pi$, (according to the choice of a representation) and from the tautology as input ($A = \top$, assuming no observations). $K * A$ is derived likewise from input A. In terms of conditionals, the change from K to $K*A$ stems from the fact that the conditionals $\top \to K$ and $A \to K * A$, respectively, can be inferred from Δ under some inferential system. In terms of a confidence relation $\succ$ between propositions $K * A = \{B, A \wedge B \succ A \wedge \neg B\}$. Dubois et al. (2005) show that requiring the deductive closure of $K * A$ is enough to recover system P of Kraus et al.(1990). Moreover if $\succ$ is the strict part of a complete preordering, one recovers the setting of possibility theory (Dubois, Fargier, & Prade 2004) and all the AGM axioms of belief revision (restricted to consistent inputs). In other words, $\succ$ is a comparative possibility relation in the sense of Lewis(1973), that derives from a plausibility ordering $\geq_\pi$ of possible worlds. Under a plausibility ordering $\geq_\pi$, it is well-known after Grove(1988) that K (resp. $K * A$) are the set of propositions true in the most plausible worlds (resp. where A is true).

This approach is very similar to probabilistic reasoning as emphasized by Pearl (1988), Dubois and Prade (1994). A set of conditionals Δ is the qualitative counterpart of a set of conditional probabilities of the form $P(B \mid A) = \alpha$ defining a family of probability measures. There is no need to resort to infinitesimals for bridging the gap between nonmonotonic reasoning and probabilistic reasoning. Recent works by Gilio and colleagues (2002) indicate that probabilistic reasoning with conditionals of the form $P(B \mid A) = 1$, precisely behaves like system P of Kraus et al. Benferhat *et al.* (1999), show that if we restrict to so-called big-stepped probabilities, conditionals can be interpreted by constraints $P(A \wedge B) > P(A \wedge \neg B)$.

Along the same lines, extracting a minimally informative plausibility ordering of worlds $\geq_\pi$ from a set of conditionals is very similar to the application of the maximal entropy principle from a set of conditional probabilities, an approach advocated by Paris (1994). This similarity has been studied by Maung(1995). So reasoning according to a plausibility ordering is also similar to probabilistic reasoning with Bayes nets (Pearl 1988). In this approach, the background knowledge is encoded by means of a (large) joint probability distribution on the state space defined by a set of (often Boolean) attributes. This probability distribution embodies statistical data pertaining to a population (e.g. of previously diagnosed patients, for instance) in the form of a directed acyclic graph and conditional probability tables. The advantage of the Bayes net format is to lay bare conditional independence assumptions and simplify the computation of inference steps accordingly. The network is triggered by the acquisition of observations on a case. Inferring a conclusion C based on observing A requires the computation of a conditional probability $P(C \mid A)$, and interpreting it as the degree of belief that C is true for the current situation for which all that is known is A. Apart from computing degrees of belief, one is interested in determining the most probable states upon learning A.

It is clear that the plausibility ordering in the above view of the AGM framework plays the same role as a Bayes net. Especially, $\geq_\pi$ might compile a population of cases, even if this population is ill-defined in the non-monotonic setting (the agent knows that "Birds fly" but it is not entirely clear which population of birds is referred to). It means that the input observations, since pertaining only to the case at hand, are not of the same nature as the plausibility ordering, and are not supposed to alter it, just like a Bayes net is not changed by querying it. In this framework, iterating belief change just means accumulating consistent observations and reasoning from them using the background knowledge. Interestingly, plausibility orderings, encoded as possibility distributions can be represented using the same graphical structures as joint probability distributions (see (Benferhat *et al.* 2002a)), and local methods for reasoning in such graphs can be devised (BenAmor, Benferhat, & Mellouli 2003). These graphical representations are equivalent to the use of possibilistic logic, but not necessarily more computationally efficient. In the purely ordinal case, CP-nets are also the counterparts of Bayes nets, and it is strange they are only proposed for preference modeling, while they could also implement a form of plausible reasoning compatible with the above "ontology" of qualitative reasoning under incomplete observations using background knowledge.

Belief Revision as Prioritized Merging

A radically different view is to consider that an epistemic state is made of uncertain evidence about a particular world of interest (a static world, again). It gathers the past uncertain observations obtained so far about a single case. So the belief set K is actually a completely ordered set (ordered by the epistemic entrenchment), and the underlying plausibility ordering on worlds describes what is the most plausible solution to the problem at hand. The epistemic entrenchment describes what should be more or less believed about the current case. In the BRPM view, the plausibility ordering is no longer like a statistical distribution.

The new observations A have the same status as the plausibility ordering, and are likely to modify it. They are testimonies or sensor measurements. They could be unreliable, uncertain.

So this kind of belief change is particularly adapted to the robotics environment for the fusion of unreliable measurements. It also accounts for the problem of collecting evidence, where the main issue is to validate facts relevant to a case on the basis of unreliable testimonies and incomplete observations. As an example, consider a criminal case where the guilty person is to be found on the basis of (more or less unreliable) testimonies and clues. The investigator's beliefs reflect all evidence gathered so far about the case. The input information consists of an additional clue or testimony.

Under this view, belief revision means changing the pair $(K, \geq_\pi)$ into another pair $(K * A, \geq_{\pi_A})$. Again the belief set K is induced by the plausibility ordering, but here there is no background knowledge at work. A new input should be merged with the existing information, with its own reliability level. If this level is too weak, it may be contradicted by the original belief set. Note that K cannot be viewed as knowledge (as opposed to belief). It is just what the agent thinks is more likely. Here, iterating the revision process makes sense, and comes down to a merging process because the a priori information and the input information are of the same nature. The success postulate just expresses the fact that the newest information is the most reliable. Not questioning this postulate has led to a view of iterated belief revision where the newest piece of information is always more reliable than the previous ones. One may argue that iterated belief revision can be more convincingly considered as a form of prioritized merging. Indeed, it seems that assigning priorities on the sole basis of the recency of observations in a static problem about which information accumulates is not always a reasonable assumption. Sherlock Holmes would not dismiss previously established facts on the basis of new evidence just because such evidence is new.

At the computational level, an epistemic state $(K, \geq_\pi)$ is best encoded as an ordered belief base using possibilistic logic (Dubois, Lang, & Prade 1994) or kappa rankings (Williams 1995). However the meaning of a prioritized belief base differs according to whether it is viewed as a partial epistemic entrenchment (what Williams calls an "ensconcement") or as a set of constraints on a family of possible epistemic entrenchments (possibilistic logic). Practical methods for merging ordered belief bases were devised in (Benferhat *et al.* 1999), (Benferhat *et al.* 2000) and in the special case when the success postulate is acknowledged see (Benferhat *et al.* 2002c).

The numerical counterpart to this view of iterated revision here is to be found in Shafer(1976)'s mathematical theory of evidence. In this theory, an unreliable testimony takes the form of a proposition E and a weight $m(E)$ reflecting the probability that the source providing E is reliable. It means that with probability $1 - m(E)$, the input information is equivalent to receiving no information at all. More generally, a body of evidence is made of a set of propositions E_i along with positive masses $m(E_i)$ summing to 1. $m(E_i)$ is the probability that proposition E_i correctly reflects the agent's evidence about the case at hand. The degree of belief $Bel(C)$ of a proposition C is the probability that C can be logically inferred from the agent's body of evidence (summing the masses of propositions E_i that imply C). Revising the agent belief upon arrival of a sure piece of information A ($m'(A) = 1$) comes down to a conditioning process ruling out all states or worlds that falsify A. If the input information is not fully reliable, Dempster's rule of combination, an associative and commutative operation, carries out the merging process. Note that the symmetry of the operation is due to the fact that the new pair $(A, m'(A))$ is merged with the body of evidence. The smaller $m'(A)$, the less effective is the input information A in the revision process.

When the input information is legitimately considered as more reliable than what has been acquired so far, merging the plausibility ordering and the new observation in a non-commutative way is a possible option. A similar view was advocated by (Dubois & Prade 1992) where the plausibility ordering was encoded by means of a possibility distribution. The AGM axioms were extended to plausibility orderings $\geq_\pi$ and are thus discussed in terms of their relevance for characterizing the revision of possibility distributions by input information. The success postulate led us to consider belief revision as a form of conditioning, in the tradition of probability kinematics (Domotor 1980).

Darwiche and Pearl (1997) axioms of iterated belief change embody the principle of minimal change of the ordering that is expected when the priority is always given to the new information. Among revision operations satisfying these postulates (applied to plausibility orderings) Boutilier's natural revision (Boutilier 1993) can be viewed as iterated revision of a plausibility ordering $\geq_\pi$, with priority to the new input A. In this scheme, the resulting most plausible worlds are the $\geq_\pi$-best A-worlds, all other things remaining equal, while possibilistic conditioning flatly eliminates worlds not in agreement with the input information (thus not obeying the Darwiche-Pearl postulates). Papini and colleagues (Benferhat *et al.* 2002b) adopt the view that in the resulting plausibility ordering all A-worlds are more plausible than any $\neg A$-world all things being equal. This method also satisfies the Darwiche-Pearl postulates.

The case of uncertain inputs is discussed in (Dubois & Prade 1992). It is pointed out that two situations may occur: one whereby the degree of certainty of the new piece of information is considered as a constraint. Then, this piece of information is to be entered into the a priori ordered belief set with precisely this degree of certainty. If this degree of certainty is low it may result in a form of contraction (if the source reliably claims that a piece of information cannot be known, for instance). In probability theory this is at work when using Jeffrey's revision rule (Jeffrey 1965). Darwiche and Pearl (1997) propose one such revision operation in terms of kappa-functions. The other view is that the degree of uncertainty attached to the input is an estimation of the reliability of the source, and then the piece of information is absorbed or not into the belief set. The latter view is more in line with the prioritized merging setting.

The companion paper (Delgrande, Dubois, & Lang 2006) reconsiders postulates for iterated revision without making any recency assumption: there is a certain number of more less reliable pieces of information to be merged, one of them being the new one. If we postulate that all uncertain observations play the same role and have the same reliability, a symmetric (and possibly associative) merging process can take place.

Reliability degrees are no longer a matter of recency, but can be decided on other grounds. In (Delgrande, Dubois, & Lang 2006), four axioms, for the prioritized merging of unreliable propositions into a supposedly accepted one are proposed. They embody the BRPM scenario of evidence collection and sorting producing a clearly established fact (a propositional formula representing a belief set). Informally they express the following requirements:

- A piece of information at a given reliability level should never make us disbelieve something we accepted after merging pieces of information at strictly higher reliability levels.
- The result of merging should be consistent.
- Vacuous evidence does not affect merging.
- Optimism: The result of merging consistent propositions is the conjunction thereof.

The important postulate is optimism, which suggests that if supposedly reliable pieces of information do not conflict, we can take them for granted. In case of conflicts, one may then assume as many reliable pieces of information as possible so as to maintain local consistency. It leads to optimistic assumptions on the number of truthful sources, and justify procedures for extracting maximal consistent subsets of items of information, see (Dubois & Prade 2001). This may be viewed as an extended view of the minimal change postulate, via the concern of keeping as many information items as possible. A restricted form of associativity stating that merging can be performed incrementally, from the most reliable to the least reliable pieces of information is proposed as optional. These axioms for prioritized merging recover Darwiche and Pearl postulates (except the controversial C2 dealing with two successive contradictory inputs) as well as two other more recent postulates from (Nayak *et al.* 1996; Nayak, Pagnucco, & Peppas 2003), and from (Jin & Thielscher 2005), when the reliability ordering corresponds to recency. It also recovers the setting of Konieczny and Pino-Perez (2002) for flat merging under integrity constraints for the fusion of equally reliable items in the face of more reliable ones. The prioritized merging setting of (Delgrande, Dubois, & Lang 2006) can also be viewed as a framework for extracting a set of preferred models from a potentially inconsistent prioritized belief base. Extending the postulates to outputs in the form of an ordered belief set is a matter of further research.

Interestingly, the BRPM scenario can be articulated with the previous BRDI scenario. One may see the former as a prerequisite for the latter: first evidence must be sorted out using a BRPM step, and then once a fact has been sufficiently validated, the agent can revise plausible conclusions about the world, based on this fact using BRDI (in order to suggest the plausible guilty person in a case, thus guiding further evidence collection).

AGM = BRDI or BRPM ?

Considering the relative state of confusion in the iterated revision literature, it is not completely clear what the AGM theory is talking about: BRDI or BRMP. Due to the stress given subsequently by Gärdenfors and Makinson (1994) to the similarity between non-monotonic reasoning and belief revision, it is natural to consider that BRDI is the natural framework for understanding their results. But then it follows that iterated revision deals with a different problem, and the above discussion suggests it can be BRMP.

1. In the AGM theory you never need K to derive $K * A$, you only need the revision operation * (in other words the plausibility ordering) and A. So the notation $K * A$ is in some sense misleading, since it suggests an operation combining K and A. This point was also made by Friedman and Halpern (1996a) In the BRPM view, the resulting epistemic state is also a function of the prior epistemic state and the input information only.
2. The AGM postulates of belief revision are in some sense written from a purely external point of view, as if an observer had access to the agent's belief set from outside, would notice its evolution under input information viewed as stimuli, and describe its evolution laws (the AGM theory says: if from the outside, an agent's beliefs seem to evolve according to the postulates, then it is as if there were a plausibility ordering that drives the belief flux). In this view, the background knowledge remains hidden to the observer, and its existence is only revealed through the postulates (like small particles are revealed by theories of microphysics, even if not observed yet). In the BRPM problem, the prior plausibility ordering is explicitly stated. Under the BRDI view, for practical purposes, it also looks more natural to use the plausibility ordering as an explicit primitive ingredient (as done by (Gärdenfors & Makinson 1994) and to take an insider point of view on the agent's knowledge, rather than observing beliefs change from the outside.
3. The belief revision step in the AGM theory leaves the ordering of states unchanged under the BRDI view. This is because inputs and the plausible ordering deal with different matters, resp. the particular world of interest, and the class of worlds the plausible ordering refers to. The AGM approach, in the BRDI view is a matter of "querying" the

epistemic entrenchment relation, basically, by focusing it on the available observation. Under this point of view, axioms for revising the plausibility ordering, as proposed by (Darwiche & Pearl 1997), for instance, cannot be seen as additional axioms completing the AGM axioms. On the contrary, the prioritized merging view understands the AGM axioms as relevant for the revision of epistemic states and apply them to the plausibility ordering. As such they prove to be insufficient for its characterization, hence the necessity for additional axioms.

4. In BRDI, while belief sets seem to evolve (from K to $K * A$ to $(K * A) * B \ldots$) as if iterated belief revision would take place, $(K * A) * B$ is really obtained by gathering the available observations A and B and inferring plausible beliefs from them. Again we do not compute $(K * A) * B$ from $K * A$. But $(K * A) * B$ means $K * (A \wedge B)$ (itself not obtained from K), with the proviso that A and B should be consistent. And indeed, within the BRDI view,

$$(K * A) * B = K * (A \wedge B) \text{ if } A \wedge B \neq \bot$$

is a consequence of AGM revisions (especially Axioms 7 and 8), if we consider that after revision by A the plausibility ordering does not change (we just restrict it to the A-worlds). Strictly speaking, these axioms say that the identity holds if B is consistent with $K * A$ (not with A). However, if the relative plausibility of worlds is not altered after observing A, the subsequent revision step by observation B will further restrict $\geq_\pi$ to the $A \wedge B$-worlds since $A \wedge B \neq \bot$, and the corresponding belief set is thus exactly $K * (A \wedge B)$ corresponding the most plausible among $A \wedge B$-worlds. It underlies an optimistic assumption about input information, namely that both A and B are reliable if consistent (a postulate of prioritized merging). This situation is similar to probabilistic conditioning whereby iterated conditioning $(P(C \mid A \mid B))$ comes down to simple conditioning on the conjunction of antecedents $(P(C \mid A \wedge B))$. Of course this is also a restricted view of the AGM theory, forbidding not only the revision by $\bot$, but also by a sequence of consistent inputs that are globally inconsistent. But we claim that this restriction is sensible in the BRDI scenario.

5. If in the AGM setting, observations A, B are inconsistent then the BRDI scenario collapses, because it means that some of the input facts are wrong. In this case, even if the AGM theory proposes something, the prospect it offers is not so convincing, as this is clearly a pathological situation. Similarly, in probabilistic reasoning, conditioning on a sequence of contradicting pieces of evidence makes no sense. Within the BRDI view, the natural approach is to do a merging of observations so as to restore a consistent context prior to inferring plausible beliefs (and as suggested above, the BRPM could be applied to the merging of such inconsistent input observations). In the medical example, it is clear that the physician receiving contradictory reports about the patient will first try to sort out the correct information prior to formulating a diagnosis. In the BRPM view, there is nothing anomalous with the situation of several conflicting inputs, because this conflict is expected as being of the same nature as the possible conflict between the agent's epistemic state and one piece of input information.

In summary, under the BRDI view, the belief revision problem (moving from K to $K * A$) is totally different from the problem of revising the plausible ordering of states of nature, while in the BRPM view both are essentially the same problem and must be carried out conjointly. In particular, it makes no sense to "revise an ordering by a formula", in the AGM framework. In the BRPM view, the input proposition A is viewed as an ordering of worlds such that at least one world where A is true is more likely than any world where A is false. In other words, belief revision can be cast within a more general setting of merging uncertain pieces of evidence (encoded by plausibility orderings).

Revision of Background Knowledge by Generic Information (RBKGI)

In the BRDI view, apart from the (contingent) belief revision problem addressed by the non-pathological part of the AGM theory and non-monotonic inference, there remains the problem of revising the generic knowledge itself (encoded or not as a plausibility ordering) by means of input information of the same kind. The AGM theory tells nothing about it. This problem is also the one of revising a set of conditionals by a new conditional (Boutilier & Goldszmidt 1993). Comparing again to probabilistic reasoning, contingent belief revision is like computing a conditional probability using observed facts instantiating some variables, while revising a plausibility ordering is like revising a Bayes net (changing the probability tables and/or the topology of the graph). In the medical example, the background knowledge of the physician is altered when reading a book on medicine or attending a specialized conference on latest developments of medical practice.

One interesting issue is the following: since background knowledge can be either encoded as a plausibility ordering $\geq_\pi$ or as a conditional knowledge base Δ, should we pose the RBKGI problem in terms of revising Δ or revising $\geq_\pi$?

Suppose Δ is a conditional knowledge base, which, using rational closure, delivers a plausibility ordering $\geq_\pi$ of possible worlds. Let $A \to B$ be an additional generic rule that is learned by the agent. If $\Delta \cup \{A \to B\}$ is consistent (in the sense that a plausibility ordering $\geq_{\pi'}$ can be derived from it), it is natural to consider that the revision of $\geq_\pi$ yields the plausibility ordering $\geq_{\pi'}$, obtained from $\Delta \cup \{A \to B\}$ via rational closure. Viewed from the conditional knowledge base this form of revision is just an expansion process. The full-fledged revision would take place when the conditional $A \to B$ contradicts Δ, so that no plausibility ordering is compatible with $\Delta \cup \{A \to B\}$ (Freund 2004). This kind of knowledge change needs specific rationality postulates for the revision of conditional knowledge bases, in a logic that is not classical logic, but the logic of conditional assertions of Kraus et al.(1990).

Alternatively, one may attempt to revise the plausibility ordering $\geq_\pi$ (obtained from Δ via a default information minimisation principle), using a constraint of the form

$A \wedge B \succ A \wedge \neg B$. To do so, Darwiche-Pearl postulates can be a starting point, but they need to be extended in the context of this particular type of change. Results of (Freund 2004) and (Kern-Isberner 2001) seem to be particularly relevant in this context. For instance it is not clear that the change process should be symmetric. One might adopt a principle of minimal change of the prior beliefs under the constraint of accepting the new conditional or ordering as a constraint (Domotor 1980). A set of postulates for revising a plausibility ordering (encoded by a kappa-function) by a conditional input information of the form $A \wedge B \succ A \wedge \neg B$ is proposed by Kern-Isberner (2001). They extend the Darwiche-Pearl postulates and preserve the minimal change requirement in the sense that they preserve the plausibility ordering $\geq_\pi$ among the examples $A \wedge B$ of the input conditionals, its counterexamples $A \wedge \neg B$, and its irrelevant cases $\neg A$.

Some insights can also be obtained from the probabilistic literature (van Fraassen 1980) (Domotor 1985). For instance Jeffrey's rule consists in revising a probability distribution P, enforcing a piece of knowledge, of the form $P(A) = \alpha$, as a constraint which the resulting probability measure P^* must satisfy. The probability measure "closest" to P in the sense of relative entropy, and obeying $P^*(A) = \alpha$ is of the form $P^*(.) = \alpha.P(. \mid A) + (1 - \alpha)P(. \mid \neg A)$. The problem of revising a probability distribution by means of a conditional input of the form $P(A|B) = \alpha$ has been considered in the probabilistic literature by (van Fraassen 1981). Rules for revising a plausibility ordering can be found in (Williams 1995), (Weydert 2000), (Kern-Isberner 2001) (using the kappa functions of (Spohn 1988)) and (Dubois & Prade 1997) using possibility distributions.

However it is not clear that revising the plausibility ordering $\geq_\pi$ obtained from Δ by a constraint of the form $A \wedge B \succ A \wedge \neg B$ has any chance to always produce the same result as deriving the plausibility ordering $\geq'_\pi$ from the revised conditional knowledge base Δ after enforcing a new rule $A \to B$.

While our aim is not to solve this question, at least our paper claims that revising generic knowledge whether in the form of a conditional knowledge base, or in the form of a plausibility ordering, is a problem distinct from the one of contingent belief revision (BRDI, which is only a problem of inferring plausible conclusions), and from the prioritized merging of uncertain information. The RBKGI problem can be subject to iterated revision, as well. One may argue that RBKGI underlies an evolving world in the sense of accounting for a global evolution of the context in which we live. In some respects, the normal course of things to-day is not the same as it used to be fifty years ago, and we must adapt our generic knowledge accordingly. The distinction between updates and revision is not so clear when generic knowledge is the subject of change.

Conclusion

This position paper tried to lay bare three problems of belief change corresponding to different scenarios. Results in the literature of iterated belief change should be scrutinized further in the context of these problems. It is clear that addressing these problems separately is a simplification. For instance in the BRDI approach, observations are always considered as sure facts, but one may consider the more complex situation of inferring plausible conclusions from uncertain contingent information using background knowledge. Also the assumption that in the BRDI approach, contingent inputs never alter the background knowledge is also an idealization: some pieces of information may destroy part of the agent's generic knowledge, if sufficiently unexpected (think of the destruction of the Twin Towers); moreover, an intelligent agent is capable of inducing generic knowledge from a sufficient amount of contingent observations. The latter is a matter of learning, and the question of the relationship between learning and belief revision is a natural one even if beyond the scope of this paper.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change : partial meet contraction and revision functions. *J. Symbolic Logic* 50:510–530.

BenAmor, N.; Benferhat, S.; and Mellouli, K. 2003. Anytime propagation algorithm for min-based possibilistic graphs. *Soft Computing* 8:150–161.

Benferhat, S.; Dubois, D.; Prade, H.; and Williams, M. 1999. A practical approach to fusing prioritized knowledge bases. In *Proc. 9th Portuguese Conference on Artificial Intelligence*, Lecture Notes in Artificial Intelligence, 222–236. Springer.

Benferhat, S.; Dubois, D.; Kaci, S.; and Prade, H. 2000. Encoding information fusion in possibilistic logic:a general framework for rational syntactic merging. In *Proc.14th Europ.Conf.on Artificial Intelligence (ECAI2000)*, 3–7. IOS Press.

Benferhat, S.; Dubois, D.; Garcia, L.; and Prade., H. 2002a. On the transformation between possibilistic logic bases and possibilistic causal networks. *Int. J. Approximate Reasoning* 29:135–173.

Benferhat, S.; Dubois, D.; Lagrue, S.; and Papini, O. 2002b. Making revision reversible: an approach based on polynomials. *Fundamenta Informaticae* 53:251–280.

Benferhat, S.; Dubois, D.; Prade, H.; and Williams, M. 2002c. A practical approach to revising prioritized knowledge bases. *Studia Logica* 70:105–130.

Benferhat, S.; Dubois, D.; and Prade, H. 1997. Nonmonotonic reasoning, conditional objects and possibility theory. *Artificial Intelligence* 92:259– 276.

Benferhat, S.; Dubois, D.; and Prade, H. 1999. Possibilistic and standard probabilistic semantics of conditional knowledge. *J. Logic and Computation* 9:873–895.

Biazzo, V.; Gilio, A.; Lukasiewicz, T.; and Sanfilippo, G. 2002. Probabilistic Logic under Coherence, Model-Theoretic Probabilistic Logic, and Default Reasoning in System P. *J. Applied Non-Classical Logics* 12(2):189–213.

Boutilier, C., and Goldszmidt, M. 1993. Revision by conditionals beliefs. In *Proc. of the 11th National Conf. on Artificial Intelligence (AAAI'93)*.

Boutilier, C. 1993. Revision sequences and nested conditionals. In *Proceedings of IJCAI'93*.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89:1–29.

Delgrande, J.; Dubois, D.; and Lang, J. 2006. Iterated belief revision as prioritized merging. In *Proceedings of KR'06, Windermere, U.K.*

Domotor, Z. 1980. Probability kinematics and representation of belief change. *Philosophy of Science* 47:284–403.

Domotor, Z. 1985. Probability kinematics - conditional and entropy principles. *Synthese* 63:74–115.

Dubois, D., and Prade, H. 1992. Belief change and possibility theory. In Gärdenfors, P., ed., *Belief Revision*. Cambridge University Press. 142–182.

Dubois, D., and Prade, H. 1994. Non-standard theories of uncertainty in knowledge representation and reasoning. *The Knowledge Engineering Review* 9:399–416.

Dubois, D., and Prade, H. 1997. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *Int. J. Approximate Reasoning* 17:295–324.

Dubois, D., and Prade, H. 2001. Possibility theory in information fusion. In *Data Fusion and Perception*, volume 431 of *CISM Courses and Lectures*. Springer. 53–76.

Dubois, D.; Fargier, H.; and Prade, H. 2004. Ordinal and probabilistic representations of acceptance. *J. Artificial Intelligence Research* 22:23–56.

Dubois, D.; Fargier, H.; and Prade, H. 2005. Acceptance, conditionals, and belief revision. In *Conditionals, Information, and Inference*, volume 3301 of *Lecture Notes in Artificial Intelligence*. Springer. 38–58.

Dubois, D.; Lang, J.; and Prade, H. 1994. Possibilistic logic. In Gabbay, D.; Hogger, C.; and Robinson, J., eds., *Handbook of logic in Artificial Intelligence and logic programming*, volume 3. Clarendon Press - Oxford. 439–513.

Dubois, D.; Moral, S.; and Prade, H. 1998. Belief change rules in ordinal and numerical uncertainty theories. In *Belief Change*. Kluwer. 311–392.

Fishburn, P. C. 1986. The axioms of subjective probabilities. *Statistical Science* 1:335–358.

Freund, M. 2004. On the revision of preferences and rational inference processes. *Artificial Intelligence* 152:105–137.

Friedman, N., and Halpern, J. 1996a. Belief revision: A critique. In *Proceedings of KR'96*, 421–631.

Friedman, N., and Halpern, J. Y. 1996b. Plausibility measures and default reasoning. In *Proceedings of AAAI'96*, 1297–1304.

Gärdenfors, P., and Makinson, D. 1994. Nonmonotonic inference based on expectations. *Artificial Intelligence* 65:197–245.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.

Grove, A. 1988. Two modellings for theory change. *J. Philos. Logic* 17:157–170.

Jeffrey, R. 1965. *The logic of decision*. McGraw-Hill.

Jin, Y., and Thielscher, M. 2005. Iterated revision, revised. In *Proc. IJCAI'05*, 478–483.

Kern-Isberner, G. 2001. *Conditionals in Nonmonotonic Reasoning and Belief Revision*, volume 2087 of *Lecture Notes in Artificial Intelligence*. Springer.

Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: a qualitative framework. *J. Logic and Computation* 12(5):773–808.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44(1-2):167–207.

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.

Lehmann, D. 1995. Belief revision, revised. In *Proceedings of IJCAI'95*, 1534–1540.

Lewis, D. 1973. *Counterfactuals*. Basil Blackwell, U.K.

Maung, I. 1995. Two characterizations of a minimum information principle for possibilistic reasoning. *Int. J. Approximate Reasoning* 12:133–156.

Nayak, A.; Foo, N.; Pagnucco, M.; and Sattar, A. 1996. Changing conditional beliefs unconditionally. In *Proceedings of TARK96*, 119–135.

Nayak, A.; Pagnucco, M.; and Peppas, P. 2003. Dynamic belief revision operators. *Artificial Intelligence* 146:193–228.

Nayak, A. 1994. Iterated belief change based on epistemic entrenchment. *Erkenntnis*.

Paris, J. 1994. *The Uncertain Reasoner's Companion.* Cambridge University Press, Cambridge.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

Pearl, J. 1990. System z: A natural ordering of defaults with tractable applications to default reasoning. In *Proc. of the 3rd Conf. on Theoretical Aspects of Reasoning about Knowledge (TARK'90)*, 121–135. Morgan & Kaufmann, San Mateo, CA.

Shafer, G. 1976. A mathematical theory of evidence. *Princeton University Press*.

Spohn, W. 1988. Ordinal conditional functions: a dynamic theory of epistemic states. In Harper, W. L., and Skyrms, B., eds., *Causation in Decision, Belief Change and Statistics*, volume 2. Kluwer Academic Pub. 105–134.

van Fraassen, B. 1980. Rational belief and probability kinematics. *Philosophy of Science* 47:165–187.

van Fraassen, B. 1981. A problem for relative information minimizers. *British J. Philosophy of Science* 33:375–379.

Weydert, E. 2000. How to revise ranked probabilities. In *Proc.14th Europ.Conf.on Artificial Intelligence (ECAI2000)*, 38–44. IOS Press.

Williams, M. 1995. Iterated theory-based change. In *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95)*, 1541–1550.