

# Three Scenarios for the Revision of Epistemic States <sup>\*</sup>

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

January 18, 2008

## Abstract

This position paper discusses the difficulty of interpreting the iterated belief revision problem. Axioms of iterated belief revision are often presented as extensions of the AGM axioms, upon receiving a sequence of inputs, likely to alter not only the belief set, but also the epistemic entrenchment relation underlying the revision operator. Iterated belief revision presupposes that more recent inputs have priority over less recent ones. We argue that this view of iterated revision is at odds with the suggestion of Gärdenfors and Makinson, that belief revision and non-monotonic reasoning are two sides of the same coin. It is not clear that non-monotonic reasoning modifies the ranking of possible worlds implicit in default rules. We lay bare three different paradigms of revision based on specific interpretations of the epistemic entrenchment implicitly at work and of the input information. If the epistemic entrenchment stems from default rules and the input is contingent, then AGM revision is a matter of changing plausible conclusions, and iterated revision makes no sense. However, if the epistemic entrenchment encodes uncertain contingent evidence and the input information as well, then iterated revision reduces to prioritized merging. A third problem where iteration makes sense corresponds to the revision, by the addition of

---

<sup>\*</sup>This position paper was triggered by discussions with Jerome Lang and Jim Delgrande at a Belief Revision seminar in Dagstuhl, in August 2005, and presented at the 2006 Non-Monotonic Reasoning Workshop, Windermere, U.K. This version is to appear in the Journal of Logic and Computation

new default rules, of a conditional knowledge base describing background information. The three scenarios are compared with similar problems in the framework of probabilistic reasoning.

## 1 Introduction

The interest in belief revision as a topic of investigation in artificial intelligence was triggered by the book of P. Gärdenfors [31], following the approach introduced with his colleagues C. Alchourrón, and D. Makinson [2] 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. 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. On this basis, important examples of revision operators in the form of selection functions and partial meets are described. An important assumption is that belief revision takes place in a static context, namely, that the input information is supposed to bring insight to a case that the agent deals with, and the beliefs of the agent do not evolve due to the change of the nature of the case. They change due to the collection of new pieces of evidence about the case.

Gärdenfors [31] also shows that under the full set of AGM axioms, explicit revision operations can be replaced by a so-called epistemic entrenchment relation between propositions of the language, through which revision operators can be simply expressed. This relation indeed acts like a priority assignment instrumental to determine the resulting belief set after revision. Properties of an epistemic entrenchment make it expressible in terms of a complete plausibility ordering of possible worlds, such that the resulting belief set, after receiving an input information in the form of a proposition, is viewed as the set of propositions that are true in the most plausible worlds where the input proposition holds. Since an epistemic entrenchment is a complete preordering between propositions, it looks like a comparative probability relation [27], even if it has different properties. It suggests some analogy between probabilistic reasoning and belief revision.

The AGM approach is one-shot [53] and does not discuss iteration. Since then, iterated revision has been the topic of quite a number of works [47], [59], [16], [42], [37]. 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. This paper does not intend to dismiss existing approaches to iterated belief revision. It only tries to distinguish between several revision problems and see what

existing iterated belief revision theories have to say about these problems. Our ambition is here to propose an informal comparative discussion of three possible revision scenarios, with a view to help sorting out what concrete problems may lie behind the name “iterated belief revision”. We emphasise the existence of similar issues in probabilistic reasoning as in belief revision, the main difference between these areas being the languages used (symbolic vs. numerical), that do not have the same expressiveness.

## **2 Is iterated belief revision a well-defined problem?**

In the scope of iterated revision, 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 claim it is lost, since the AGM theory tells nothing about iteration. Others claim that it changes along with the belief set, and they tried to state axioms governing the change of the plausibility ordering of the worlds, viewing them as an extension of the AGM axioms [16]. This trend led to envisaging iterated belief revision as a form of prioritized merging where the priorities assigned to pieces of input information reflect their recency.

However, this notion of iterated belief revision does not seem to fit with Makinson and Gärdenfors [45] view of belief revision as the other side of non-monotonic reasoning. For these authors, non-monotonic reasoning relies on a set of “expectations that guide our beliefs without quite being part of them”. Such expectations can be modelled by a partial order over propositions, playing the role of the epistemic entrenchment relation. An expectation relation may also derive from the analysis of a set of conditionals, in the style of [43], yielding a ranking of worlds via the so-called rational closure. The revised belief set is then the result of a simple inference of conditionals from conditionals, whereby propositional conclusions tentatively drawn are altered by the arrival of new pieces of evidence. Noticeably, in this framework, there is no reason why the conditional information, hence the expectation ordering, should be revised. Iteration comes down to the inference of new conclusions and the 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 (along this line, see a critical discussion of Darwiche and

Pearl[16] axioms in [22]).

Friedman and Halpern [29] also complained that iterated belief revision research relies too much on the finding of new axioms justified by toy-examples with no practical significance, and representation results, while more stress should be put on laying bare an appropriate “ontology or scenario”, that is, “describing carefully what it means for something to be believed by an agent and what the status is of new information received by the agent”. It is not always clear which scenario 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. 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 conjunctive combination of uncertainty in the theory of evidence [54]. The difference between both scenarios lies in the strength of the input information. When the input is a true observation, revision corresponds to a form of conditioning found in the theory of evidence as a special case of the combination rule.

In this paper, we somewhat pursue this discussion by pointing out that there is no unique way of understanding the epistemic entrenchment itself : sometimes, it represents background information about the world, telling what is normal from what it is not, in a more or less 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 are not directly accessible. For instance it may encompass the claim that flying birds are more normal than non-flying ones, without making it precise exactly which class of birds is concerned. In other applications, the plausibility ordering expresses beliefs about unreliable observations concerning 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. Finally, one may argue that even the background information used for inferring plausible conclusions may evolve.

These considerations lead to lay bare 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 the AGM theory and Lehmann and Magidor’s conditional logic under rational closure [43], then we come up with a qualitative counterpart to statistical reasoning, where inputs stand for incomplete but safe information about a case at hand. We call it Belief Revision as Defeasible Inference (BRDI).

- If we take it for granted that the epistemic entrenchment gathers uncertain evidence about a case, integrating new uncertain pieces of evidence is a matter of Belief Revision as Prioritized Merging (BRPM). A recent paper [35] proposes a formal framework for the BRPM situation in full details.
- 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 call it Revision of Background Knowledge by Generic Information (RBKGI). In the latter case, inputs often take the form of conditionals.

### 3 Belief Revision as Defeasible Inference (BRDI)

Under this first view, the AGM theory and non-monotonic reasoning are really regarded as two sides of the same coin (see Rott [53], chap. 4 for a detailed comparison of belief change and nonmonotonic reasoning postulates). However, while in the AGM approach, only a flat belief set denoted  $K$ , composed of logical formulas  $p, q, \dots$ , is explicitly available (the epistemic entrenchment is implicit in the axioms of the theory), the nonmonotonic logic approach lays bare all the pieces of information that allow an agent to reason from incomplete reliable evidence and background knowledge. While in the AGM paradigm, the primitive objects are the belief set and the input information, we propose in this section a view where everything derives from the background knowledge, synthesized in the form of a partial ordering of propositions or a conditional knowledge base, and from the available evidence. This view, outlined in [32], is fully developed by Dubois Fargier and Prade [19] [20] as a theory of accepted beliefs.

#### 3.1 A semantic view of defeasible inference

In the following, we consider a classical propositional language, but we do not distinguish between logically equivalent propositions. Hence, we consider propositions  $p$  as subsets  $A, B, \dots$  of a set  $\Omega$  of possible worlds, in other words, events

(to borrow from the probabilistic literature). For clarity, the set of models of a proposition  $p$  is denoted  $[p]$ .  $A^c$  denotes the complement of set  $A$  ( $A^c = [\neg p]$ , where  $\neg p$  is the negation of proposition  $p$ ). The influence of syntax on revision is out of the scope of this paper. So, in the whole paper, the word “proposition” is deliberately taken as synonymous to “event”. As a consequence, the belief set  $K$  of an agent is understood as its set  $[K] = \bigcap_{p \in K} [p]$  of models. We see two advantages:

- In axiomatic studies about revision, it is simpler to stick to a semantic approach than to use syntax and write an axiom of independence from syntax.
- getting rid of syntax makes it easier to lay bare the analogies between revision notions and their counterparts in uncertainty theories

Under such a proviso, it is assumed that the agent’s epistemic state consists of three components:

1. A *confidence relation*, in the form of a partial ordering  $\succ$  on events  $A, B, \dots$  (induced by propositions expressed in a given language). This relation, which should be in agreement with logical deduction, expresses that some events are more normally expected [32] (or less surprizing) than others.  $A \succ B$  means that  $A$  is more entrenched, generally more plausible [30] (that is, less surprizing) than  $B$ . 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.
2. A set of *contingent observations* concerning a case of interest for the agent. It takes the form of a propositional formula with models forming the set  $A$ . The observations are evidence about this case, not general considerations about similar cases. Such pieces of evidence are reliable facts (or at least accepted as such), hence *consistent with each other*. So it is assumed that  $A \neq \emptyset$ . It means that a preliminary process is capable of handling conflicting observations and come up with a consistent report.
3. The *belief set* of the agent, what is denoted  $K * A$  in the AGM theory. It consists of propositions tentatively accepted as true by the agent about the case, in the face of the current observations  $A$ . Propositions in  $K * A$  are inferred from the observations and the background knowledge (so it is not an independent part of the epistemic state). Namely, in terms of a confidence relation  $\succ$  between events  $K * A = \{B, A \cap B \succ A \cap B^c\}$ .

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 confidence relation reflects the domain knowledge of the physician. 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.

Some remarks on the above setting are in order.

- As shown in Dubois et al.[20], if the set of inferred beliefs  $\{B, A \cap B \succ A \cap B^c\}$  is assumed to be deductively closed, and  $\succ$  is well-behaved with respect to classical inference, six AGM axioms of belief revision (restricted to consistent inputs) can be recovered<sup>1</sup>.
- If moreover  $\succ$  is the strict part of a complete preordering, one recovers all the eight AGM axioms and the setting of possibility theory [19]. In other words,  $\succ$  is a comparative possibility relation in the sense of Lewis[44], that derives from a complete plausibility preordering  $\geq_\pi$  of possible worlds, namely  $A \succ B$  if and only if  $\exists s_1 \in A, \forall s_2 \in B, s_1 >_\pi s_2$ . Under such a plausibility ordering  $\geq_\pi$ , it is well-known after Grove[33] that  $K$  (resp.  $K * A$ ) is the set of propositions true in the most plausible worlds (resp. where  $A$  is true).
- The closed set of contingent beliefs of the agent before hearing about  $A$  is  $K = \{B, B \succ B^c\}$ . Formally, under this view, the original belief set  $K$  is inferred from the background knowledge only. The lack of input information can be modelled by the tautology, viewed as a non-informative input ( $A = \Omega$ , assuming no observations are available<sup>2</sup>).  $K * A$  is derived likewise from input  $A$  and  $\succ$ .
- That input information is safe explains why the success postulate ( $A \in K * A$ ) makes sense. The assumption of reliable facts means that the agent does not question the validity of the inputs. It says nothing about the actual

---

<sup>1</sup>Axioms 4 (if  $\neg p \notin K$  then  $K * p$  contains all consequences of  $K \cup \{p\}$ ), and 8 (if  $\neg q \notin K * p$  then  $K * (p \wedge q)$  contains all consequences of  $K * p \cup \{q\}$ ) cannot be derived.

<sup>2</sup>A non-informative input should not to be confused with a certain, reliable, input.

truth of these facts. Only that the agent has no reason to doubt them. This point implies that the input should be consistent. If a new observation is made and it is inconsistent with the previous ones, then no inference can be drawn, and some preprocessing of the observed facts is needed (this is the topic of the next section, in fact).

- In the case where prior beliefs in  $K$  are consistent with the new input information  $A$ , the revision step should reduce to an expansion. Indeed there is no need to change our previous beliefs about the current situation when new evidence does not contradict such beliefs. This requirement is sanctioned by the AGM setting, and is satisfied when a possibility relation stands for the confidence relation.

### 3.2 Expressing background knowledge by conditional knowledge bases

A confidence relation representing generic knowledge may directly stem from a set of conditionals  $\Delta$ .  $\Delta$  contains pieces of conditional knowledge of the form  $A \rightarrow B$  where  $\rightarrow$  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  $A \cap B \succ A \cap B^c$ , understood as in terms of the above confidence relation [30]. A plausibility ordering of worlds  $\geq_\pi$  can then be derived from such constraints via some information minimization principle (like rational closure of Lehmann and Magidor [43], or equivalently, the most compact ranking compatible with the constraints[52], or yet the principle of minimal specificity of possibilistic logic [7]).

In terms of conditionals, the change from  $K$  to  $K * A$  stems from the fact that the conditionals  $\Omega \rightarrow K$  and  $A \rightarrow K * A$ , respectively, can be inferred from  $\Delta$  under some inferential system. Dubois et al. [20] show that requiring the deductive closure of  $K * A$  is enough to recover system P of Kraus et al.[41].

### 3.3 Belief revision and probabilistic reasoning

The BRDI framework is very similar to probabilistic reasoning as emphasized by Pearl [51], Dubois and Prade [24]. The set of conditionals  $\Delta$  is the qualitative counterpart to a set of conditional probabilities of the form  $P(B | A) = \alpha$  defining a family of probability measures. In fact, Pearl [51] indicates that each conditional can be viewed as an infinitesimal conditional probability  $P(B | A) = 1 - \epsilon$ ,

and Lehmann and Magidor [43] provide a full-fledged account of the infinitesimal probability version of non-monotonic reasoning, that comes close to Adams conditional logic [1]. So in some sense, non-monotonic reasoning is a qualitative form of probabilistic reasoning.

In fact, there is no need to resort to infinitesimals for bridging the gap between nonmonotonic reasoning and probabilistic reasoning. Benferhat *et al.* [8], show that if we restrict to so-called big-stepped probabilities, conditionals interpreted by constraints  $P(A \cap B) > P(A \cap B^c)$  obey system P of Kraus Lehmann and Magidor. A big-stepped probability on  $\Omega$  defines a total order of states, whereby each state is more probable than the disjunction of less probable states.

Recent works by Gilio and colleagues [12] also indicate that probabilistic reasoning with conditionals of the form  $P(B | A) = 1$ , understood in the setting of de Finetti, behaves like a nonmonotonic logic. Namely, in de Finetti's view, it is possible to have that  $P(B | A) = 1$  while  $P(A) = 0$ . It implies that the solution to a set of conditional statements interpreted as  $P(B | A) = 1$  is a sequence of probability measures with disjoint supports, defining an ordered partition of the possible worlds precisely corresponding to the plausibility rankings obtained in the rational closure approach. For instance the set of conditional probabilities  $\{P(B | A) = 1, P(C | B) = 1, P(C^c | A) = 1\}$  is interpreted as the set of constraints

$$\begin{aligned} P(B \cap A) &= P(A); \\ P(C \cap B) &= P(B); \\ P(C^c \cap A) &= P(A). \end{aligned}$$

It implies  $P(A) = 0$ , i.e any probability measure  $P_1$  satisfying such equations trivially satisfies the first and the last ones. Following De Finetti, when  $P(A) = 0$  (hence  $P(A^c) = 1$ ), this is interpreted as the idea that  $A$  is one order of magnitude less likely than  $A^c$ . It is then possible to consider a second step where non-trivial solutions  $P_2$  of  $P(B \cap A) = P(A); P(C^c \cap A) = P(A)$  are searched for, with  $P_2(A) > 0$ . This example (which corresponds to the famous Tweety penguin case) suggests that this probabilistic approach singles out the rule  $B \rightarrow C$  from the other two, just like rational closure does.

Note that extracting a minimally informative plausibility ordering of worlds  $\geq_\pi$  from a set of conditionals via rational closure is very similar to the selection, by application of the maximal entropy principle, of a unique probability distribution satisfying a set of conditional probability constraints, an approach advocated by Paris [50]. This similarity has been studied by Maung[46].

Reasoning according to a plausibility ordering is also similar to probabilistic reasoning with Bayes nets [51]. 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 confidence relation the above view of the AGM framework plays the same role as a Bayes net. Especially, the plausibility ordering  $\geq_\pi$  on states is similar to a joint probability distribution and 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). Interestingly, plausibility orderings, encoded as possibility distributions can be represented using the same graphical structures as joint probability distributions (see [4]), and local methods for reasoning in such graphs can be devised [3]. These graphical representations are equivalent to the use of possibilistic logic, but are not necessarily more computationally efficient. In the purely ordinal case, CP-nets [14] are also counterparts to Bayes nets representing preference relations on complex spaces. It is strange they are only proposed for preference modeling, while they could also implement plausibility orderings as background knowledge, for qualitative reasoning under incomplete observations.

Adopting this analogy between non-monotonic and probabilistic reasoning 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, so that only plausible contingent conclusions are modified by the arrival of new observations.

Gärdenfors [31] also discusses conditional probability in the light of belief revision. However, he argues that conditioning by a positive probability event is a form of expansion rather than revision. This is because the belief set  $K$  associated

to a probability measure  $P$  on  $\Omega$  is understood as  $K = \{B, P(B) = 1\}$ , so that  $K \subseteq K * A = \{B, P(B | A) = 1\}$ , if  $P(B) > 0$ . However this very conservative view can be challenged if we consider beliefs as propositions with sufficiently high probabilities ( $P(B) = 1 - \epsilon$  using infinitesimals, or  $P(B) > P(B^c)$ , using big-stepped probabilities). In the latter situation, since a high value of  $P(B)$  can be consistent with a small value of  $P(B | A)$ , it is more intuitively satisfying to consider probabilistic conditioning as revision, as proposed in [22]. Indeed, the main advantage of the probabilistic conditioning over classical inference is the possibility to make likely conclusions unlikely by acquiring new information (see also [19] for a comparison of nonmonotonic and probabilistic notions of defeasible acceptance). Moreover, if the knowledge of an agent is represented by a single probability measure (as in the subjectivist Bayesian tradition), it can be argued that such a representation extends the concept of *complete* belief set (the autoduality property  $P(B) = 1 - P(B^c)$  standing for the condition  $B \in K$  or  $B^c \in K$ ), and it makes no sense to expand a complete belief set. From this standpoint, changing a probability measure cannot mean anything but revising it.

## 4 Belief Revision as Prioritized Merging

A radically different view of belief revision considers an epistemic state as 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 epistemic state is modelled as a completely ordered set of propositions  $(K, \succ)$  (ordered by the epistemic entrenchment  $\succ$ ), and the underlying plausibility ordering  $\geq_\pi$  on worlds indicates what are the most plausible solutions to the problem at hand. The input information  $A$  is of the same nature as the epistemic state.  $A$  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. Absorbing the input information is then likely to modify the original plausibility ordering.

### 4.1 The framework of prioritized merging

Suppose the epistemic entrenchment  $\succ$  describes what should be more or less believed *about the current case*. Then, the plausibility ordering  $\geq_\pi$  is no longer like a statistical distribution, and the new observations  $A$  are additional testimonies. They could be unreliable, uncertain.

This kind of belief change is particularly adapted to the robotics environment for the fusion of unreliable sensor 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 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  $(K, \succ)$  into another pair  $(K * A, \succ_A)$ . The plausibility ordering of worlds  $\geq_\pi$  is changed into another one  $\geq_{\pi_A}$  under the constraint expressed by the input information. Here, however, there is no background knowledge at work. The pair  $(K, \succ)$  cannot be viewed as background knowledge (as opposed to contingent belief). It is just what the agent thinks is more likely. The new input, with its own reliability level, should be merged with the existing information. If this level is too weak, it may be contradicted by the original belief set.

Now, iterating the revision process makes sense, and comes down to a merging process because the prior 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 (BRPM). 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.

Clearly this discussion emphasises a crucial difference between BRDI and BRPM: the former is a fundamentally asymmetric problem because the input information and the background knowledge do not play the same role (nor can we exchange the prior beliefs  $K$  and the input  $A$ ), while BRPM is a symmetric process (this symmetry being possibly broken by priority assignments based on reliability). This point is also discussed by Maynard-Reid II and Shoham [34] who try to reduce belief revision to the fusion of an expert opinion with a novice opinion.

At the computational level, an epistemic state  $(K, \succ)$  is best encoded as an ordered belief base using possibilistic logic [21] or kappa rankings [59]. However the meaning of a prioritized belief base differs according to whether it is viewed as a partial epistemic entrenchment (what Williams [60] calls an “ensconcement”)

or as a set of constraints on a family of possible epistemic entrenchments (possibilistic logic). In the former case, the priority attached to a piece of information is fixed, and cannot change via inference processes. It means that if  $A$  and  $B$  are present in a prioritized belief base,  $A$  logically implies  $B$ , and  $A$  is specified as more reliable than  $B$ , this is a conflict only due to an improper allocation of priorities. On the contrary, in possibilistic logic, only lower bounds on priority levels are assumed and no conflicts are generated by priority assessments. But, the certainty level of a belief can be upgraded via inference. Practical methods for merging ordered belief bases were devised in [9], [5] and, for the special case when the success postulate is acknowledged, see [10].

## 4.2 An extended framework for iterated revision

Darwiche and Pearl [16] axioms are stated in terms of iterated revision of an ordered belief set. Under our notations, they take the following form (where  $K * A$  is the belief set induced from a plausible ordering of states  $\geq_{\pi_A}$ ):

- (C1) if  $A \subseteq B$  then  $\geq_{\pi_A} = \geq_{(\pi_B)_A}$ ;
- (C2) if  $A \subseteq B^c$  then  $\geq_{\pi_A} = \geq_{(\pi_B)_A}$ ;
- (C3) if  $B \in K * A$  then  $B \in (K * B) * A$  induced from  $\geq_{(\pi_B)_A}$ ;
- (C4) if  $B^c \notin K * A$  then  $B^c \notin (K * B) * A$  induced from  $\geq_{(\pi_B)_A}$ .

Their representation theorem shows that these axioms embody the principle of minimal change of the ordering when priority is always given to the new information. Among revision operations satisfying these postulates (applied to plausibility orderings) Boutilier's natural revision [13] can be viewed as iterated revision of a plausibility ordering  $\geq_{\pi}$ , with priority to the new input  $A$  (interpreted as  $A \succ A^c$ ). In this scheme, the resulting most plausible worlds are the  $\geq_{\pi}$ -best  $A$ -worlds, all other things remaining equal. In contrast, possibilistic conditioning eliminates worlds not in agreement with the input information, making them equally unlikely [23] (thus violating Darwiche-Pearl postulate C2). Papini and colleagues [6] adopt the view that in the resulting plausibility ordering all  $A$ -worlds are more plausible than any  $A^c$ -world all things being equal. This method also satisfies the Darwiche-Pearl postulates.

Delgrande et al. [35] reconsider these postulates for iterated revision without making any recency assumption: there is a certain number of more less reliable or important pieces of information to be merged, one of them being the new one. If we postulate that all observations play the same role and have the same priority, a

symmetric (and possibly associative) merging process can take place.

Priorities are no longer a matter of recency, but can be assigned on other grounds. In [35], 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, that eventually produces a clearly established fact (a propositional formula representing a belief set). Informally these axioms express the following requirements:

- A piece of information at a given priority level should never make us disbelieve something we accepted after merging pieces of information at strictly higher priority 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 justifies procedures for extracting maximal consistent subsets of items of information, see [26]. 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 [48, 49], and from [37], when the priority ordering corresponds to recency. It also recovers flat merging under integrity constraints after Konieczny and Pino-Perez [40], for the fusion of equally reliable items in the face of more reliable ones. The prioritized merging setting of [35] is tailored to the extraction of 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 needs further research. A related question is studied by Benferhat and Kaci [11].

When the input information is legitimately considered as more reliable than what has been acquired so far, the success postulate suggests considering belief revision as a form of conditioning, in the tradition of probability kinematics [17]. Namely, conditional probability  $P(\cdot | A) = P_{new}(A)$  is provably equal to the closest probability measure (in the sense of relative entropy) to the original measure  $P$ , under the constraint  $P_{new}(A) = 1$ . A similar view was advocated in [23] for plausibility orderings encoded by means of a possibility distribution. The AGM axioms were extended to possibility distributions for characterizing their revision in terms of conditioning by the input information, and respecting the minimal change principle. Conditioning in this setting is a particular case of the above prioritized merging paradigm.

In the case of uncertain inputs, two situations may occur [23]. One view is that the degree of priority 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. This is in line with the prioritized merging setting. According to the other view, 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 prior 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 [36]. Darwiche and Pearl [16] propose one such revision operation in terms of kappa-functions.

### 4.3 Prioritized merging in Dempster-Shafer theory

The numerical counterpart to the prioritized merging view of iterated revision here is to be found in Shafer's mathematical theory of evidence [54]. In this theory, a body of evidence is made of propositions  $E_i$  along with a *mass assignment*, i.e. positive masses  $m(E_i)$  summing to 1.  $m(E_i)$  is the probability that  $E_i$  is the proposition that correctly reflects the evidence about the case at hand. It is the probability of only knowing  $E_i$ , or equivalently, the probability that  $E_i$  properly reflects the agent's belief set. In particular, if the agent's only information is based on some unreliable testimony, it takes the form of a proposition  $E$  and a weight  $m(E)$  reflecting the probability that the source providing  $E$  is reliable. With probability  $m(E)$  the agent's belief set is  $K$  with  $[K] = E$ . It means that with probability  $1 - m(E)$ ,  $K = \emptyset$  (i.e.  $[K] = \Omega$ ), i.e. the input information is equivalent to receiving no information at all. A non-dogmatic mass assignment should always have  $m(\Omega) > 0$ , so as to allow for revision via any kind of input

information. For epistemic entrenchments, it corresponds to assuming that any contingent proposition is less entrenched than tautologies.

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's belief upon arrival of a fully reliable 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 a new mass assignment  $\{(A, m'(A)), A \subseteq \Omega\}$  (and not only an input  $A$ ) is merged with the original body of evidence. The smaller  $m'(A)$ , the less effective is the input information  $A$  in the revision process. This is the numerical translation of the first postulate of prioritized merging. Dempster rule of combination obeys the other postulates: normalization ensures the consistency of the result of merging non-dogmatic bodies of evidence; combining a body of evidence with a vacuous input ( $m'(\Omega) = 1$ ) does not create any change; finally the combination rule generalizes set-intersection, i.e. obeys optimism.

## 5 AGM = BRDI or BRPM ?

The AGM theory does not take sides on its potential extensions to iterated revisions since it is one-shot. So how should the AGM theory be interpreted : a special case of BRDI or of BRPM? Due to the stress put by Makinson and Gärdenfors [45] on 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 BRPM. The existence of these two paradigms, corresponding to two kinds of revision problems blurs the meaning of the AGM theory, iterated revision and postulates governing them.

1. In the AGM theory, you do not need  $K$  to derive  $K * A$ , you only need the revision operation  $*$  (in other words the plausibility ordering on  $\Omega$ ) and  $A$ . So the notation  $K * A$  is in a sense misleading, since it suggests an operation combining  $K$  and  $A$ . This point was also made by Friedman and Halpern [29], and later by Maynard-Reid II and Shoham [34]. Of course, if  $K$  and  $A$  are consistent, it turns out that the most plausible models of  $A$  are precisely  $[K] \cap A$ . This is the expansion case, where  $K * A$  can be computed from

$K$  and  $A$ . In case of a genuine revision, this is no longer true. In the BRPM view,  $K$  is often viewed as an ordered belief set, i.e encoding the whole prior epistemic state, and the resulting epistemic state is then a function of  $K$  and the input information.

2. The AGM postulates of belief revision are in a 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 (as 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 in [32]) 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. In particular, it makes no sense to "revise an ordering by a formula". Under this point of view, axioms for revising the plausibility ordering, as proposed by [16], 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 full epistemic states and applies 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  $(K * A) * B$  ...) 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  $K * (A \cap B)$  (itself not obtained from  $K$ ),

with the proviso that  $A$  and  $B$  should be consistent. And indeed, within the BRDI view, the following *reduction axiom*

$$(K * A) * B = K * (A \cap B) \text{ if } A \cap B \neq \perp$$

looks natural. It is a consequence of AGM framework, *if it is assumed that after revision by  $A$  the plausibility ordering does not change* (it is just restricted to the  $A$ -worlds)<sup>3</sup>. Strictly speaking, the AGM axioms say that the identity  $(K * A) * B = K * (A \cap B)$  holds if  $B$  is consistent with  $K * A$  (not only with  $A$ ). However, after observing  $A$ , the plausibility ordering is restricted to  $A$  and the relative plausibility of  $A$ -worlds is not altered. The subsequent revision step due to observation  $B$  will further restrict  $\geq_\pi$  to the  $A \cap B$ -worlds since  $A \cap B \neq \perp$ , and the corresponding belief set is thus exactly  $K * (A \cap B)$  corresponding the most plausible among  $A \cap B$ -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 ( $P(C \mid A \cap B)$ ). Of course this is also a restricted view of the AGM theory, forbidding not only the revision by  $\perp$ , 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 (it is impossible to condition on the empty set in standard probability theory). Within the BRDI view, the natural approach is to do a merging (using BRPM) of observations so as to restore a consistent context prior to inferring plausible beliefs. One may indeed see BRPM as a prerequisite for BRDI: first, evidence must be sorted out using a BRPM step (so as to establish the facts, as in a crime case), and then once a

---

<sup>3</sup>The assumption that the epistemic entrenchment remains the same after the revision step does not belong to the AGM theory. Hence the reduction axiom is not *stricto sensu*, derivable from the AGM axioms. However since AGM Axioms 7 and 8 are needed to derive the epistemic entrenchment and they involve a second input, one may argue that these axioms involve two revision steps, and that the epistemic entrenchment thus laid bare is valid for these two steps.

fact is validated, the agent can revise plausible conclusions about the world, based on this fact, using BRDI (in order to suggest the plausible guilty person, thus guiding further evidence collection). 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 the input information.

In summary, under the BRDI view, the problem of revising contingent beliefs (moving from  $K$  to  $K * A$ ) is totally different from the problem of revising the epistemic entrenchment relation, while in the BRPM view both are essentially the same problem and must be carried out conjointly.

## 6 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 [15]. 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  $\Delta$ , should we pose the RBKGI problem in terms of revising  $\Delta$  or revising  $\geq_\pi$ ?

Suppose  $\Delta$  is a conditional knowledge base, which, using rational closure, delivers a plausibility ordering  $\geq_\pi$  of possible worlds. Let  $A \rightarrow B$  be an additional generic rule that is learned by the agent. If  $\Delta \cup \{A \rightarrow B\}$  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  $\Delta \cup \{A \rightarrow B\}$  via rational closure. In terms of the conditional knowledge base, this form of revision is just an expansion process. The full-fledged revision would take place when the conditional  $A \rightarrow B$  contradicts  $\Delta$ , so that no plausibility ordering is compatible with  $\Delta \cup \{A \rightarrow B\}$  [28]. 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.[41].

Alternatively, one may attempt to revise the plausibility ordering  $\geq_{\pi}$  (obtained from  $\Delta$  via a default information minimisation principle), using a constraint of the form  $A \cap B \succ A \cap B^c$ . 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 [28] and Kern-Isberner [39] 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, accepting the new conditional or ordering as a constraint like in probability kinematics [18]. A set of postulates for revising a plausibility ordering (encoded by a kappa-function) by a conditional input information of the form  $A \cap B \succ A \cap B^c$  is proposed by Kern-Isberner [39]. 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  $A \cap B$  of the input conditionals, its counterexamples  $A \cap B^c$ , and its irrelevant cases  $A^c$ .

Some insights can also be obtained from the probabilistic literature [56] [17]. For instance Jeffrey's rule [36] 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  $P^*(\cdot) = \alpha.P(\cdot | A) + (1 - \alpha)P(\cdot | A^c)$ . 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 in [57]. Rules for revising a plausibility ordering can be found in [59], [58], [39] (using the kappa functions of Spohn [55]) and [25] using possibility distributions. These formal proposals have some relevance to the BRPM problem as well, since in the BRPM and the RBKGI problems both prior information and inputs are at the same level of generality: contingent evidence in BRPM, generic knowledge in RBKGI.

However it is not clear that revising the plausibility ordering  $\geq_{\pi}$  obtained from  $\Delta$  by a constraint of the form  $A \cap B \succ A \cap B^c$  has any chance to always produce the same result as deriving the plausibility ordering  $\geq_{\pi'}$  from the revised conditional

knowledge base  $\Delta$  after enforcing a new rule  $A \rightarrow B$ .

Solving this question is beyond the scope of this paper. At least, this section 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. Generic knowledge being the result of experience, one may conjecture it is more stable and safer than contingent evidence. So it may be less liable to revision than to update. Indeed, RBKGI might reflect a (slowly) 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 should perhaps be revisited when generic knowledge is the subject of change.

## 7 Conclusion

This position paper tried to lay bare three problems of belief change corresponding to different scenarios with specific features summarised in Table 1 where the similarity between symbolic and numerical frameworks is highlighted. Results in the literature of iterated belief change should be scrutinized further in the context of these scenarios.

It may be that other scenarios for belief change could be pointed out. For instance, one may, by symmetry imagine a problem of revising contingent information by generic knowledge acting as input, exchanging the roles of the two basic items in BRDI. However it is hard to make sense of this possibility and find a natural example for it. This situation may happen indirectly: generic input knowledge alters prior generic knowledge that in turn leads to revising the agent's contingent beliefs about the current situation. Viewed as such, this fourth case could be reduced to a sequence of RBKGI and BRDI steps.

It is also clear that addressing these problems separately is a simplification. For instance in the BRDI approach, observations are always considered as reliable, 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

<b>Problem</b>	<b>Generic Knowledge</b>	<b>Prior Beliefs</b>	<b>Input Information</b>	<b>After Change</b>
<b>BRDI</b>	Conditional Knowledge Base	Propositions in Belief Set	Contingent Observation	Revised Belief Set
<b>Probabilistic Reasoning</b>	Bayes Net	Prior Degrees of Belief	Variable Instanciation	Posterior Degrees of Belief
<b>BRPM</b>	None	Prioritized Propositions	Prioritized Propositions	Accepted Fact (+ Revised Priorities)
<b>Evidence Combination</b>	None	Mass Assignment	Mass Assignment	Merged Mass Assignment
<b>RBKGI</b>	Conditional Knowledge Base	(Irrelevant)	New Conditional	Revised Conditional Knowledge Base
<b>Probability Kinematics</b>	Probability Measure	(Irrelevant)	Conditional Probability	Revised Probability Measure

Table 1: Comparison of different change problems

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. More generally, experience, viewed as collecting evidence over a long period of time, alters generic knowledge of agents. The latter is the very purpose of learning theory, and the question of the relationship between learning and belief revision is a natural one even if beyond the scope of this paper (see discussions and results by Kevin Kelly [38] on this topic).

### **Acknowledgements**

Many thanks to three referees that provided numerous useful comments that led to a far better presentation of the thesis of this paper.

## **References**

- [1] E. Adams. *The Logic of Conditionals*. Reidel, U.K, Dordrecht, The Netherlands, 1975.

- [2] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change : partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [3] N. Ben-Amor, S. Benferhat, and K. Mellouli. Anytime propagation algorithm for min-based possibilistic graphs. *Soft Computing*, 8:150–161, 2003.
- [4] S. Benferhat, D. Dubois, L. Garcia, and H. Prade. On the transformation between possibilistic logic bases and possibilistic causal networks. *Int. J. Approximate Reasoning*, 29:135–173, 2002.
- [5] S. Benferhat, D. Dubois, S. Kaci, and H. Prade. Encoding information fusion in possibilistic logic:a general framework for rational syntactic merging. In *Proc.14th Europ.Conf.on Artificial Intelligence (ECAI2000)*, pages 3–7. IOS Press, 2000.
- [6] S. Benferhat, D. Dubois, S. Lagrue, and O. Papini. Making revision reversible: an approach based on polynomials. *Fundamenta Informaticae*, 53:251–280, 2002.
- [7] S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artificial Intelligence*, 92:259– 276, 1997.
- [8] S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge. *Journal of Logic and Computation*, 9:873–895, 1999.
- [9] S. Benferhat, D. Dubois, H. Prade, and M.A. Williams. A practical approach to fusing prioritized knowledge bases. In *Proc. 9th Portuguese Conference on Artificial Intelligence*, Lecture Notes in Artificial Intelligence, pages 222–236. Springer, 1999.
- [10] S. Benferhat, D. Dubois, H. Prade, and M.A. Williams. A practical approach to revising prioritized knowledge bases. *Studia Logica*, 70:105–130, 2002.
- [11] S. Benferhat and S. Kaci. Fusion of possibilistic knowledge bases from a postulate point of view. *Int. J. of Approximate Reasoning*, 33:255–285, 2003.

- [12] V. Biazzo, A. Gilio, T. Lukasiewicz, and G. Sanfilippo. Probabilistic Logic under Coherence, Model-Theoretic Probabilistic Logic, and Default Reasoning in System P. *Journal of Applied Non-Classical Logics*, 12(2):189–213, 2002.
- [13] C. Boutilier. Revision sequences and nested conditionals. In *Proceedings of IJCAI'93*, 1993.
- [14] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res.*, 21:135–191, 2004.
- [15] C. Boutilier and M. Goldszmidt. Revision by conditionals beliefs. In *Proc. of the 11th National Conf. on Artificial Intelligence (AAAI'93)*, 1993.
- [16] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
- [17] Z. Domotor. Probability kinematics and representation of belief change. *Philosophy of Science*, 47:284–403, 1980.
- [18] Z. Domotor. Probability kinematics - conditional and entropy principles. *Synthese*, 63:74–115, 1985.
- [19] D. Dubois, H. Fargier, and H. Prade. Ordinal and probabilistic representations of acceptance. *J. Artificial Intelligence Research*, 22:23–56, 2004.
- [20] D. Dubois, H. Fargier, and H. Prade. Acceptance, conditionals, and belief revision. In *Conditionals, Information, and Inference*, volume 3301 of *Lecture Notes in Artificial Intelligence*, pages 38–58. Springer, 2005.
- [21] D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *Handbook of logic in Artificial Intelligence and logic programming*, volume 3, pages 439–513. Clarendon Press - Oxford, 1994.
- [22] D. Dubois, S. Moral, and H. Prade. Belief change rules in ordinal and numerical uncertainty theories. In *Belief Change*, pages 311–392. Kluwer, 1998.
- [23] D. Dubois and H. Prade. Belief change and possibility theory. In P. Gärdenfors, editor, *Belief Revision*, pages 142–182. Cambridge University Press, 1992.

- [24] D. Dubois and H. Prade. Non-standard theories of uncertainty in knowledge representation and reasoning. *The Knowledge Engineering Review*, 9:399–416, 1994.
- [25] D. Dubois and H. Prade. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *Int. J. of Approximate Reasoning*, 17:295–324, 1997.
- [26] D. Dubois and H. Prade. Possibility theory in information fusion. In *Data Fusion and Perception*, volume 431 of *CISM Courses and Lectures*, pages 53–76. Springer, 2001.
- [27] P. C. Fishburn. The axioms of subjective probabilities. *Statistical Science*, 1:335–358, 1986.
- [28] M. Freund. On the revision of preferences and rational inference processes. *Artificial Intelligence*, 152:105–137, 2004.
- [29] N. Friedman and J. Halpern. Belief revision: A critique. In *Proceedings of KR'96*, pages 421–631, 1996.
- [30] N. Friedman and J. Y. Halpern. Plausibility measures and default reasoning. In *Proceedings of AAAI'96*, pages 1297–1304, 1996.
- [31] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.
- [32] P. Gärdenfors and D. Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65:197–245, 1994.
- [33] A. Grove. Two modellings for theory change. *J. Philos. Logic*, 17:157–170, 1988.
- [34] Pedrito Maynard-Reid II and Yoav Shoham. Belief fusion: Aggregating pedigreed belief states. *Journal of Logic, Language and Information*, 10(2):183–209, 2001.
- [35] J. Lang J. Delgrande, D. Dubois. Iterated belief revision as prioritized merging. In *Proceedings of KR'06, Windermere, U.K*, 2006.
- [36] R. Jeffrey. *The logic of decision*. McGraw-Hill, 1965.

- [37] Y. Jin and M. Thielscher. Iterated revision, revised. In *Proc. IJCAI'05*, pages 478–483, 2005.
- [38] Kevin T. Kelly. Iterated belief revision, reliability, and inductive amnesia. *Erkenntnis*, 50:7–53, 1999.
- [39] G. Kern-Isberner. *Conditionals in Nonmonotonic Reasoning and Belief Revision*, volume 2087 of *Lecture Notes in Artificial Intelligence*. Springer, 2001.
- [40] S. Konieczny and R. Pino Pérez. Merging information under constraints: a qualitative framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.
- [41] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Art. Int.*, 44(1-2):167–207, 1990.
- [42] D. Lehmann. Belief revision, revised. In *Proceedings of IJCAI'95*, pages 1534–1540, 1995.
- [43] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.
- [44] D. Lewis. *Counterfactuals*. Basil Blackwell, U.K, 1973.
- [45] D. Makinson and P. Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence, pages 185–205. Springer, 1991.
- [46] I. Maung. Two characterizations of a minimum information principle for possibilistic reasoning. *Int. J. Approximate Reasoning*, 12:133–156, 1995.
- [47] A. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 1994.
- [48] A. Nayak, N. Foo, M. Pagnucco, and A. Sattar. Changing conditional beliefs unconditionally. In *Proceedings of TARK96*, pages 119–135, 1996.
- [49] A. Nayak, M. Pagnucco, and P. Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146:193–228, 2003.

- [50] J. Paris. *The Uncertain Reasoner's Companion*. Cambridge University Press, Cambridge, 1994.
- [51] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [52] J. Pearl. 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)*, pages 121–135. Morgan & Kaufmann, San Mateo, CA, 1990.
- [53] H. Rott. *Change, Choice and Inference*. Oxford University Press, 2001.
- [54] G. Shafer. A mathematical theory of evidence. *Princeton University Press*, 1976.
- [55] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision, Belief Change and Statistics*, volume 2, pages 105–134. Kluwer Academic Pub., 1988.
- [56] B. van Fraassen. Rational belief and probability kinematics. *Philosophy of Science*, 47:165–187, 1980.
- [57] B. van Fraassen. A problem for relative information minimizers. *British J. of the Philosophy of Science*, 33:375–379, 1981.
- [58] E. Weydert. How to revise ranked probabilities. In *Proc. 14th Europ. Conf. on Artificial Intelligence (ECAI2000)*, pages 38–44. IOS Press, 2000.
- [59] M.A. Williams. Iterated theory-based change. In *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95)*, pages 1541–1550, 1995.
- [60] M.A. Williams. On the logic of theory base change. In *Proc. of the Europ. Workshop on Logics in Artificial Intelligence (JELIA'95)*, volume 838 of *Lecture Notes in Computer Sciences*, pages 1541–1550, Berlin, 1995. Springer Verlag.