

Iterated revision as prioritized merging

James Delgrande
Simon Fraser University
Burnaby, BC, Canada
jim@cs.sfu.ca

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

Jérôme Lang
IRIT-CNRS
Toulouse, France
lang@irit.fr

Abstract

Standard accounts of iterated belief revision assume a static world, about which an agent receives a sequence of observations. More recent items are assumed to have priority over less recent items. We argue that there is no reason, given a static world, for giving priority to more recent items. Instead we suggest that a sequence of observations should be *merged* with the agent's beliefs. Since observations may have differing reliability, arguably the appropriate belief change operator is prioritized merging. We develop this view here, suggesting postulates for prioritized merging, and examining existing merging operators with respect to these postulates. As well, we examine other suggested postulates for iterated revision, to determine how well they fit with the prioritized merging interpretation. All postulates for iterated revision that we examine, except for Darwiche and Pearl's controversial C2, are consequences of our suggested postulates for prioritized merging.

Introduction

In knowledge representation, the area of *belief change* addresses the specification and construction of systems for reasoning about a possibly uncertain and possibly evolving world. A fundamental belief change operation is *belief revision* (along with its dual operation of *belief contraction*). Belief revision concerns the situation in which new information may be inconsistent with the reasoner's beliefs, and needs to be incorporated in a consistent manner where possible. The common assumption is that in revision an agent receives information about a *purely inertial* (or *static*) world¹. That is, the agent performs no actions that can cause the world to evolve, nor do any exogenous actions occur.

A belief revision operator is not arbitrary, but rather is usually guided by various *rationality criteria*. One of the

most widely accepted of the rationality criteria is the *success postulate*: that a new item of information (which we'll refer to as an *observation*) is always accepted. Thus if we use K to denote the agent's initial belief state, and the agent receives the observation α , then in the revised state $K * \alpha$, α is believed. Much attention has been paid to the problem of *iterated belief revision*, in which an agent receives a stream or sequence of (possibly conflicting) observations. An assumption common to all approaches to iterated revision is that revision takes place whenever an observation is received. Hence for a sequence of observations $\alpha_1, \dots, \alpha_n$, the result of revising by this information is $(\dots (K * \alpha_1) * \dots * \alpha_{n-1}) * \alpha_n$. This assumption, together with the success postulate, implies that more recent observations are assigned a higher priority than less recently received observations. For example, p will be believed in the state resulting from $(K * \neg p) * p$.

This ordering of observations for revision is reasonable in a dynamic framework, where events can occur and induce unpredicted changes in the world. However, this does not carry over to a purely inertial framework, where the state of the world does not change. In this case, the order in which the observations are made is not really significant in itself, and we might have received the very same observations in a different order. Thus, for example, imagine coming home to three messages on your telephone answering machine, each from a friend independently reporting on a party that you failed to attend; clearly the order of messages is irrelevant. In fact there are examples wherein priority is given to the *older* items of information. Thus in history, all other things being equal, older reports concerning some ancient event may be given more weight than more recent reports, since they are closer to the event itself and so presumably more accurate. Even if in *some* contexts it makes sense that reliability coincides with recency, these contexts are specific and should not be considered as a general case.

In this paper we address iterated belief revision based upon these intuitions. Thus, given an inertial framework, the order of observations is irrelevant. However, it is quite possible that some observations may be more reliable than others (since sources of information may have varying degrees of reliability). Hence, for us the problem of iterated revision is, in fact, a problem of prioritized merging of information: The agent has (perhaps as part of an epistemic state) a set of be-

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

¹(Friedman & Halpern 1999) argue that more generally, what counts in belief revision is not that the world itself is static but that the *propositions used to describe the world* are static, therefore, time-stamping variables also allows for dealing with an evolving world in a pure belief revision setting. Thus, belief revision is also typical of the situation where an agent investigates a past event and tries to reason about what was the real state of the world when this event took place.

iefs, and must modify its epistemic state so as to account for the new observation. Each observation is attached an evaluation of its reliability (or priority); these evaluations induce a total preorder over the observations (and, in fact, the agent's beliefs), and the problem becomes one of merging this information into a single set of beliefs while taking the priorities into account. Thus, "standard" iterated belief revision corresponds to the situation in which observations are linearly ordered, and an observation's priority corresponds with its temporal position.

We give a number of postulates governing prioritized merging, and examine existing approaches with respect to these postulates. It turns out that these postulates also imply the AGM revision postulates. As well, we examine postulates that have been proposed for iterated revision. It proves to be the case that all postulates for iterated revision that we examine, except for Darwiche and Pearl's controversial C2, are consequences of our suggested postulates for prioritized merging.

First though, we need to address some crucial questions about the meaning of belief revision, especially when it comes to revising epistemic states and iteration. We fully agree with Friedman and Halpern(1996) claim that before stating postulates for iterated revision, the concerned problem addressed by the theory must be laid bare, namely what they call the "underlying *ontology* or scenario". This is the topic of the next section, where we will argue that there are (at least) three different scenarios for the revision of epistemic states. In particular, we note that the view of iterated belief revision as a merging problem adopted in this paper is at odds with the view of belief revision as resulting from non-monotonic inference based on background knowledge, originally suggested by Gärdenfors and Makinson as a reinterpretation of the belief revision axioms. In the following section, we develop a general approach for prioritized merging and investigate properties of prioritized merging as well as its relation to unprioritized merging. After this, in the next section we come back to iterated revision, and show that it is essentially a specific case of prioritized merging. In the next to last section, we briefly question the assumption that the result of merging is a plain formula, and consider the case of merging complex observations into complex observations. Last, we conclude with a summary and some remarks concerning future research.

Revision of epistemic states: three views

Background Interest in belief revision as a foundational topic in artificial intelligence arguably goes back to the *AGM approach* (Alchourrón, Gärdenfors, & Makinson 1985), which provides a well-known set of *rationality postulates* for belief revision. This approach assumes that belief states are modelled by sets of sentences, called *belief sets*, closed under the logical consequence operator of a logic that includes classical propositional logic. An important assumption is that belief revision takes place in an inertial (or static) world, so that the input information is always with respect to the same, static world. The axiomatic framework given by the rationality postulates has a corresponding semantic

model, given by a so-called *epistemic entrenchment* relation between propositions of the language. Properties of an epistemic entrenchment make it representable by means of a complete plausibility ordering over possible worlds, and the resulting belief set, after receiving input α , is viewed as the set of propositions that are true in the most plausible worlds where α holds. An epistemic entrenchment relation (or plausibility ordering on worlds) can be viewed as an *epistemic state*, containing not just an agent's set of beliefs, but also sufficient information for carrying out revision, given any input. The AGM approach is silent on what happens to an epistemic state following a revision.

Subsequent work on iterated revision has addressed this issue. However, we suggest that the study of iterated revision has led to a number of misunderstandings. Some researchers have claimed that, once the belief set has been revised by some input information, the epistemic entrenchment relation is simply lost, thus precluding the possibility of any further iteration. Others have claimed that the epistemic entrenchment relation changes along with the belief state. This has led to axioms being proposed which are intended to govern the change of the plausibility ordering of worlds; these additional axioms are viewed as extending the AGM axioms. This approach has led to a view of belief revision as a form of prioritized merging, where the priority assignment to pieces of input is reflected by their recency.

I Belief revision as defeasible inference However, this view of iterated revision seems to be at odds with (Gärdenfors & Makinson 1994), which argues that belief revision is the other side of non-monotonic reasoning. This view can be characterised as *belief revision as defeasible inference (BRDI)*. The BRDI problem can be stated as follows: given a plausibility ordering on worlds describing background knowledge and an input information α representing sure observations about the case at hand, find the most plausible worlds where α is true. It is thus assumed that the agent possesses three kinds of information: an epistemic entrenchment, induced by a plausibility ordering of worlds, a set of contingent observations about the current world, under the form of sentences, and a set of beliefs about the current world induced by observations and the epistemic entrenchment (Dubois, Fargier, & Prade 2004). The role of the latter is to guide the agent in tasks of inquiry and deliberation, sorting what is credible from what is less credible in view of the contingent observations, considered as sure facts.

Under this view, the AGM approach to belief revision (as related to non-monotonic reasoning) has little to do with iterated revision as studied by subsequent researchers. According to Gärdenfors and Makinson, a revised belief set is the result of an inference step involving (nonmonotonic or defeasible) conditionals, from which propositional conclusions are tentatively drawn. These conclusions are altered by the arrival of new pieces of evidence, supposed to be sure, hence consistent ((Friedman & Halpern 1996)). In this framework, there is no clear reason why the conditional information, hence the plausibility ordering, should be revised

upon making new contingent observations, and “iteration” in fact corresponds to the nonmonotonic inference of new conclusions – that is, different inputs simply yield different non-monotonic conclusions. The Darwiche and Pearl (1997) axioms that were intended to extend the AGM axioms so as to allow for iterated revision by modifying the plausibility ordering seem to have little relevance here. Moreover, in the AGM theory you never need the original belief set when deriving the revised belief set (a point also made by Friedman and Halpern (1996)). You only need the epistemic entrenchment and the input information to construct the latter, while the original belief set is based on the most plausible worlds induced by the epistemic entrenchment.

II Belief revision as incorporation of evidence A quite different view is to assume that an epistemic state represents uncertain evidence about a particular (static) world of interest. An agent now gathers and “compiles” possibly uncertain or inaccurate observations about a particular world. So the underlying plausibility ordering on worlds represents the agent’s subjective ranking as to which world is most likely the actual world, and the degree of entrenchment of a proposition (evaluated on the basis of the most plausible world that violates it (Gärdenfors 1988; Dubois & Prade 1991)) is an agent’s *degree of belief* that a proposition is true or not. The instigating philosophical work here arguably is (Spohn 1988), on ordinal conditional functions, and most work on iterated belief revision is placed within this framework. However the mathematical theory of evidence by Shafer (1976) discusses a very similar problem (the one of merging uncertain evidence), as well as possibility theory (Dubois & Prade 1988). This overall approach can be characterised as *belief revision as incorporation of evidence (BRIE)*.

Under this view, belief revision means changing the plausibility ordering in response to new information, and it makes sense to talk about iterating the process of revision. The success postulate then expresses the fact that the most recent information is the most reliable. However, given that observations concern a static world, it is by no means clear why the most recent should be taken as the most reliable. If all observations are equally reliable, then it seems most natural to somehow merge these observations with the agent’s beliefs. If observations come with varying degrees of reliability, then it seems most natural to exploit *this* reliability ordering while merging the observations with the agent’s beliefs.

Comparison of the first two views To sum up the main differences between these views: under the BRDI view, the belief revision step leaves the epistemic entrenchment relation (i.e., the plausibility ordering on states) unchanged. This is because inputs and the plausibility ordering deal with different matters, resp. the particular world of interest, and the class of worlds the plausibility ordering refers to. Under this view, AGM revision is a matter of “querying” the epistemic entrenchment relation; so axioms for revising the epistemic state (e.g. (Darwiche & Pearl 1997)), cannot be seen as

additional axioms completing the AGM axioms. In contrast, the BRIE view understands the AGM axioms as relevant for the revision of epistemic states including the plausibility ordering. Since the AGM axioms neither explicitly dealt with ranked belief sets, they could not address iterated revision, and additional axioms are needed to this end. Under this view, the prior epistemic state and the inputs can be handled in a homogeneous way, since they both consist of uncertain evidence about the world of interest; thus, it makes sense to have a new epistemic state be a function of the prior epistemic state and the input information; and it is then natural to iterate the process. But then, the plausibility ordering does not contain any information about how it should be revised, while, in the BRDI view, its role is precisely to provide a revision procedure.

III Belief revision of background knowledge There is a third form of belief revision which we will just briefly mention: *revision of background knowledge by generic information*. This problem is the one, not addressed in the AGM theory, of revising the epistemic entrenchment relation within the BRDI view, not upon receiving a contingent input, but upon receiving a new piece of generic knowledge. Here, in some fashion, the generic knowledge of an agent’s epistemic state is revised. This could also be characterised as *theory change* (in the same sense as changing a scientific theory).

Since in the BRDI view the epistemic entrenchment is equivalently represented as a set of conditionals, this problem is also the one of revising a set of conditionals by a new conditional (Boutilier & Goldszmidt 1993). Since the epistemic entrenchment is induced by a plausibility ordering on worlds, and since an input conditional can be modeled as another plausibility ordering, this third revision problem is also akin to the revision of a preference relation by another one as studied in (Freund 2004). Rules for revising a plausibility ordering can be found in (Williams 1995) and (Darwiche & Pearl 1997) in terms of Spohn’s ordinal conditional functions or (Dubois & Prade 1997) in terms of possibility theory. Several authors, as for instance Friedman and Halpern (1996), Dubois, Moral and Prade (1998) note that the results of (Darwiche & Pearl 1997) do not make it clear whether the considered issue is that of revising an epistemic entrenchment ordering, or that of allowing for iteration in the revision of belief sets.

Alternatively, one might adopt a principle of minimal change of the background knowledge under the constraint of accepting the new conditional. Then this is no longer a symmetric merging process. This seems to be the rationale behind Boutilier’s natural revision (Boutilier 1993). This asymmetric approach has been mainly studied in the probabilistic literature (Domotor 1980). For instance, Jeffrey’s rule (Jeffrey 1965) revises probability distributions by enforcing a constraint $P(A) = x$. But there have been extensions of such probability kinematics to conditional inputs of the form $P(A | B) = x$ (van Fraassen 1981). Devising basic principles and rational change operations for qualita-

tive non-probabilistic representations of generic knowledge is left for further research. In the rest of this paper the focus is on prioritized merging of qualitative uncertain evidence.

A principled approach to prioritized merging

Prioritized merging

Prioritized merging consists of aggregating pieces of information, all partially describing the same world and attached with priority (reliability) levels, such that any information at a higher level has precedence over information at a lower level. There is nothing that says that these pieces of information are in fact correct, or that the full set is consistent. As suggested earlier, iterated belief revision is a specific case of prioritized belief merging, in which priority depends directly on the time when the observations were performed, following the *priority to recency* principle, which asserts that an older piece of information is weaker than a more recent item. While this would be perfectly acceptable in a dynamic framework, as previously argued, it is inapplicable in a purely inertial framework.

Moreover, there are other assumptions in the revision literature that we might not want to take for granted. First, we might wish to allow for more than *a single piece of information* at each reliability level or, with the usual interpretation of iterated revision, that a single, indecomposable observation is gathered at each time point. Second, we might wish to allow that a piece of information is not automatically incorporated into the agent's beliefs; it may, for instance, conflict with these beliefs, and be *less reliable* than the current beliefs.

In the following, L_{PS} denotes a propositional language generated from a finite set PS of propositional symbols. Let 2^{PS} be the induced set of possible states of nature.

Definition 1 *An observation α is defined as a consistent formula of L_{PS} . A prioritized observation base (POB) is defined as a set of observations with an attached reliability degree:*

$$\sigma = \langle \sigma(1), \dots, \sigma(n) \rangle, \text{ for some } n \geq 1, \quad (1)$$

where each $\sigma(i)$ is a (possibly empty) multiset² of propositional formulas, namely, the observations of reliability level i .

The assumption that each observation is individually consistent is in accordance with common sense; moreover it makes the presentation simpler without inducing any loss of generality, as our notions and results can be easily generalized so as to allow for inconsistent observations.

Expressing a POB as a collection of multisets of formulas, i.e., under the form $\sigma = \langle \sigma(1), \dots, \sigma(n) \rangle$, is equivalent to

²A multiset is a set in which different occurrences of the same object are distinguished, or more formally a mapping from L_{PS} to \mathbb{N} . $\sigma(i)$ is a multiset (this assumption is usual) rather than a set because we want, in some cases, to allow for counting occurrences of identical observations.

expressing it as a multiset of formulas, each being attached with a reliability degree, i.e.,

$$\sigma = \{ \langle \alpha_1, r_1 \rangle, \dots, \langle \alpha_m, r_m \rangle \}. \quad (2)$$

We have that (2) and (1) are equivalent provided that for all α_i, α_j , (a) if α_i and α_j are in the same $\sigma(k)$ then $r_i = r_j$ and (b) if $\alpha_i \in \sigma(k)$ and $\alpha_j \in \sigma(k')$ with $k < k'$ then $r_i < r_j$. Given these assumptions, all notions and results in this paper deal with ordinal scales and are thus unaffected by applying any monotone transformation to reliability degrees.

Higher values of i indicate more reliable formulas. There is no specific object for prior beliefs: since, according to the discussion in the second section, under the BRIE view prior beliefs and observations have the same status (they all bring evidence about the actual state of the world), a prior belief state, if any, will be represented as a prioritized formula (or set of formulas)³.

Note that *reliability* should not be confused with *likelihood*. The reliability of an observation reflects one's confidence in the source of the observation; hence $\alpha \wedge \beta$ may be more reliable than β if the source of the former is more reliable than that of the latter, although it certainly wouldn't be more likely. The numerical counterpart of a set of sentences $\sigma(i)$ with a reliability level used here is to be found in Shafer (1976)'s mathematical theory of evidence. In this theory, unreliable testimony takes the form of a proposition α and a weight $m(\alpha)$ reflecting the probability that the source providing α is reliable. In the theory of belief functions the reliability level r_i would be denoted $m(\alpha_i)$ and would reflect the probability, for the agent receiving α_i , of knowing only α_i (so the information is considered vacuous with probability $1 - \alpha_i$). There is no constraint relating $m(\alpha \wedge \beta)$ and $m(\alpha)$. In this approach, if the reliability of α_i is r_i it means that the (subjective) probability of α_i is *at least* r_i . This is in full agreement with possibilistic logic (Dubois & Prade 1991) where the pair (α_i, r_i) is interpreted as the inequality $N(\alpha_i) \geq r_i$, and $N(\alpha_i)$ is a degree of necessity (certainty) or, equivalently, entrenchment.

POBs and prioritized merging are nothing new: the merging of POBs (elsewhere called prioritized or stratified belief bases, prioritized defaults, or infobases) is actually (implicitly) employed in a series of works including base revision, prioritized default reasoning, possibilistic logic and inconsistency handling. Arguably, up to minor differences, what these streams of work are doing is, in the end, prioritized merging.

We use the following notation and shorthands:

- If $S \subseteq L_{PS}$ then $\bigwedge(S)$ is the conjunction of all formulas in S , with the usual convention $\bigwedge(\emptyset) = \top$. In particular, $\bigwedge \sigma(i)$ is the conjunction of all formulas in $\sigma(i)$ and $\bigwedge(\sigma) = \bigwedge_{i=1, \dots, n} \bigwedge \sigma(i)$.
- $\sigma_{i \rightarrow j} = \langle \sigma(i), \dots, \sigma(j) \rangle$ for $1 \leq i \leq j \leq n$.

³We find this view of a prior belief state as a list of input formulas in other works, e.g. (Konieczny & Pino-Pérez 2000; Lehmann 1995b). We shall return to this issue later.

- If $\sigma = \langle \sigma(1), \dots, \sigma(n) \rangle$ and $\sigma' = \langle \sigma'(1), \dots, \sigma'(p) \rangle$ then (σ, σ') (also sometimes denoted by $\sigma.\sigma'$) is the concatenation of σ and σ' , defined by $(\sigma, \sigma')(i) = \sigma(i)$ for $i \leq n$ and $(\sigma, \sigma')(i) = \sigma'(i - n)$ for $n + 1 \leq i \leq n + p$.

- If $\sigma = \langle \sigma(1), \dots, \sigma(n) \rangle$ and $\sigma' = \langle \sigma'(1), \dots, \sigma'(n) \rangle$ are two POBs, we write $\sigma \approx \sigma'$ iff for each i , we can write $\sigma(i) = \{\beta_1, \dots, \beta_p\}$ and $\sigma'(i) = \{\beta'_1, \dots, \beta'_p\}$ such that $\beta_j \equiv \beta'_j$ holds for every $j = 1, \dots, p$.

Similarly, we write $\sigma' \subseteq \sigma$ iff for every i , $\sigma'(i) \subseteq \sigma(i)$, where \subseteq is here multiset inclusion, and we simply say that σ' is a subset of σ .

- $\hat{\sigma}$ is the multiset $\bigcup_{i=1, \dots, n} \sigma(i)$.
- When $\sigma(i)$ is a singleton $\{\alpha_i\}$, we abuse notation and write $(\sigma_{1 \rightarrow i-1}, \alpha_i, \sigma_{i+1 \rightarrow n})$ instead of $(\sigma_{1 \rightarrow i-1}, \{\alpha_i\}, \sigma_{i+1 \rightarrow n})$.
- $Cons(\sigma)$ is the set of consistent subsets of σ , that is, the set of all POBs $S = (S_1, \dots, S_n)$ such that $S \subseteq \sigma$ and $\bigwedge(S)$ is consistent.

A *prioritized merging* operator \star maps any POB σ to a propositional formula $\star(\sigma)$ ⁴.

Two specific cases of prioritized merging can be distinguished:

- $n = 1$: all pieces of information have the same reliability degree. This is the case of *non-prioritized, commutative merging* (e.g., (Konieczny & Pino Pérez 2002)).
- $\sigma(i)$ is a singleton $\{\alpha_i\}$ for each i : there is a *single piece of information* for each reliability level. This is precisely the case for iterated revision, where reliability levels coincide with the position of observations in the sequence, and where a single, indecomposable observation is present at each time point. We call such a σ a *strictly ordered prioritized observation base* (SOPOB) and we write σ as an ordered list $\sigma = \langle \alpha_1, \alpha_2, \dots, \alpha_n \rangle$.

Oddly enough, while each of these two specific cases has been addressed from the axiomatic point of view in many works, we are not aware of such a study for the general case. Yet, some concrete prioritized merging operators have been proposed in the literature.

In the following we use the definition: If \succ is a strict order (i.e. a transitive and asymmetric binary relation) on a set X , then for any $Y \subseteq X$ we denote by $Max(\succ, Y)$ the set of undominated elements of Y with respect to \succ , i.e.,

$$Max(\succ, Y) = \{y \in Y \mid \text{there is no } z \in Y \text{ such that } z \succ y\}.$$

Then we have:

best-out (Benferhat *et al.* 1993)

⁴Alternatively, the outcome of a merging operator can be defined as a closed theory. Both are equivalent, since the set of propositional symbols is finite; therefore any closed theory can be described as the set of all consequences of some formula.

Let $k(\sigma) = \min\{i, \bigwedge \sigma_{i \rightarrow n} \text{ consistent}\}$. Then $\star_{bo}(\sigma) = \bigwedge \sigma_{k(\sigma) \rightarrow n}$.⁵

discrimin (Brewka 1989; Nebel 1991; Benferhat *et al.* 1993)

For $S, S' \in Cons(\sigma)$, define $S' \succ_{discrimin} S$ iff $\exists k$ such that

- (a) $\sigma_{k \rightarrow n} \cap S' \supset \sigma_{k \rightarrow n} \cap S$, and
- (b) for all $i > k$, $\sigma_{i \rightarrow n} \cap S' = \sigma_{i \rightarrow n} \cap S$.

Then $\star_{discrimin}(\sigma) =$

$$\bigvee \left\{ \bigwedge S, S \in Max(\succ_{discrimin}, Cons(\sigma)) \right\}$$

leximin (Benferhat *et al.* 1993; Lehmann 1995a)

For $S, S' \in Cons(\sigma)$, define $S' \succ_{leximin} S$ iff $\exists k$ such that

- (a) $|\sigma_{k \rightarrow n} \cap S'| > |\sigma_{k \rightarrow n} \cap S|$ and
- (b) $\forall i > k, |\sigma_{i \rightarrow n} \cap S'| = |\sigma_{i \rightarrow n} \cap S|$.

Then $\star_{leximin}(\sigma) =$

$$\bigvee \left\{ \bigwedge S, S \in Max(\succ_{leximin}, Cons(\sigma)) \right\}$$

When σ is a SOPOB, $\star_{discrimin}$ and $\star_{leximin}$ coincide with *linear merging*⁶:

linear merging (Dubois & Prade 1991; Nebel 1994)

$\sigma = \langle \alpha_1, \dots, \alpha_n \rangle$ is a SOPOB, and $\star_{linear}(\alpha_1, \dots, \alpha_n)$ is defined inductively by: $\star_{linear}() = \top$ and for $k \leq n$, $\star_{linear}(\alpha_k, \dots, \alpha_n) =$

$$\begin{cases} \alpha_k \wedge \star_{linear}(\alpha_{k+1}, \dots, \alpha_n) & \text{if consistent} \\ \star_{linear}(\alpha_{k+1}, \dots, \alpha_n) & \text{otherwise} \end{cases}$$

Remark: There are two different ways of defining merging operators on POBs, whether the output is an epistemic state (another ordered set of formulas) or a plain formula. If $\otimes(\alpha_1, \dots, \alpha_n)$ is the epistemic state resulting from the prioritized merging of $\sigma = \langle \alpha_1, \dots, \alpha_n \rangle$ then $\star(\alpha_1, \dots, \alpha_n) = Bel(\otimes(\alpha_1, \dots, \alpha_n))$ is its projection (defined in the usual sense by focusing on the most normal states). One may argue against defining a prioritized merging operator as producing a plain formula: one would expect to get an epistemic entrenchment as the result, because it is much more informative than the plain formula obtained by focusing on the most plausible states. We adopt this simplified setting here because it is sufficient to unify a large part of the literature on iterated revision and commutative merging⁷.

⁵this operation is basically at work in possibilistic logic (Dubois & Prade 1991; Dubois, Lang, & Prade 1994) where the inference from a prioritized base σ is explicitly defined as $\sigma \vdash_{Poss} \alpha$ if and only if $\star_{bo}(\sigma) \vdash \alpha$.

⁶actually called *linear base revision* in (Nebel 1994). Here we view it as a merging operator.

⁷Obviously, the problem dealt with in this article is also relevant for how to turn an ordered belief base into a flat one (by extracting its best models). This is confirmed by the examples of merging operations above, whose original motivation is the quest for variants of possibilistic logic inference (Benferhat *et al.* 1993).

Properties of prioritized merging

We now consider some desirable properties for prioritized merging. The first property, *prioritized monotonicity*, says that a piece of information at a given level i should never make us disbelieve something we accepted after merging pieces of information at strictly higher levels.

(PMon) for every $i < n$, $\star(\sigma_{i \rightarrow n}) \vdash \star(\sigma_{i+1 \rightarrow n})$

This postulate reflects the principle of prioritization: higher priority observations can be dealt with independently of lower-ranked observations. It is remarkable that, as far as we know, this postulate has never been considered within the framework of iterated belief revision. The reason for this might be that **PMon** does not compare the *leftmost* (what papers on iterated revision consider as oldest) k observations to the *leftmost* $k+1$ (as do most postulates for iterated revision, except those in (Lehmann 1995b)), but rather the *rightmost* k observations to the *rightmost* $k+1$. Note that (PMon) is equivalent to the following:

for every i and $j < i$, $\star(\sigma_{j \rightarrow n}) \vdash \star(\sigma_{i \rightarrow n})$

(the proof from left to right consist in applying (PMon) iteratively to $i-1, \dots, j$; and from right to left, in taking $j = i-1$).

Taking $i = n$ and $j = 1$ in the latter formulation, we find that (PMon) implies the following *success property*:

(Succ) $\star(\sigma) \vdash \star(\sigma(n))$

The four postulates below are classical.

(Cons) $\star(\sigma)$ is consistent

(Taut) $\star(\sigma, \top) \equiv \star(\sigma)$

(Opt) if $\bigwedge(\sigma)$ is consistent then $\star(\sigma) \equiv \bigwedge(\sigma)$

(IS) if $\sigma \approx \sigma'$ then $\star(\sigma) \equiv \star(\sigma')$

Even if uncontroversial, some of these postulates deserve some comments. First, consistency may not be a relevant concern in an enlarged setting where prioritized merging produces a prioritized set of formulas. (Thus consider, by analogy, the debate around the renormalization step in the Dempster rule of combination (Shafer 1976) with belief functions, where renormalizing the result comes down to a consistency preservation assumption). However, since we adopt an approach that produces a single formula, it is natural to require the consistency of this formula. (Taut) is reasonable given our assumption that every formula in a POB is (individually) consistent. (Opt) is *closure by optimism*: observations are all accepted if they are jointly consistent. The last postulate actually means that we could do away with the propositional language and directly work with sets of possible worlds (for instance simple support belief functions (Shafer 1976), or possibility distributions (Dubois & Prade 1988), or systems of spheres (Grove 1988)). We keep the logical setting because it adheres more closely with common practice, and readers may be more familiar with it. Lastly, the above postulates seem to be general enough to be appropriate for other kinds of merging problems, for instance merging a set of more or less important goals in the scope of decision making.

Proposition 1 (PMon) and (Opt) imply generalized success:

(GS) if, for a given value of i , $\bigwedge \sigma_{i \rightarrow n}$ is consistent then $\star(\sigma) \vdash \bigwedge \sigma_{i \rightarrow n}$.

Proof: Assume $\bigwedge \sigma_{i \rightarrow n}$ is consistent. Then by (Opt), $\star(\sigma_{i \rightarrow n}) = \bigwedge \sigma_{i \rightarrow n}$. Now, by (PMon), $\star(\sigma) \vdash \star(\sigma_{i \rightarrow n})$, therefore $\star(\sigma) \vdash \bigwedge \sigma_{i \rightarrow n}$. ■

It seems that (GS) has also never been considered in the literature on iterated revision. Again, the reason might be that (GS) focuses on the *rightmost* k observations. We believe that generalized success is very natural, and fully in the spirit of prioritization and iterated revision. In the latter context it reads: if the most recent k observations are jointly consistent then accept them all.

Corollary 1 If \star satisfies (PMon) and (Opt) then

$\star(\sigma) \vdash \star_{bo}(\sigma)$

Proof: Just take $i = k^*(\sigma) = \min\{j, \bigwedge(\sigma_{j \rightarrow n}) \text{ is consistent}\}$ in (GS). ■

Thus, any \star satisfying (PMon) and (Opt) is at least as strong as best-out merging.

(PMon), (Cons), (Taut), (IS) and (Opt) seem uncontroversial and will be called *basic* properties for prioritized merging operators. We now consider additional properties. The first we consider is *right associativity*; after an extended exploration of the ramifications of right associativity, we consider a second property that gives sufficient expressivity to imply the AGM belief revision postulates.

(RA) $\star(\sigma_{i \rightarrow n}) = \star(\sigma(i), \star(\sigma_{i+1 \rightarrow n}))$

The idea here is that prioritized merging can be carried out on a given reliability level independently of lower levels. Unlike the basic properties, right associativity, although natural and intuitive, is not totally unquestionable, as we may well imagine prioritized merging operators which are not right associative. Note that (RA) and (Succ) imply (PMon).

From flat merging to prioritized merging

(RA), together with the above properties, leads us to the following construction: since $\star(\sigma)$ can be computed incrementally by merging $\sigma(n)$ and then by incorporating $\sigma(n-1)$ and so on to the result, we can *induce a prioritized merging operator from a non-prioritized – or basic – merging operator, and vice versa*. Let us define a (*constrained*) *flat merging operator* \times as a mapping from a multiset of formulas Φ and a consistent formula θ to a closed theory $\times(\Phi|\theta)$ (the merging of Φ under integrity constraint θ , similarly as in (Konieczny & Pino Pérez 2002)) such that:

(Cons \times) $\times(\Phi|\theta)$ is consistent

(Succ \times) $\times(\Phi|\theta) \vdash \theta$

(Opt \times) if $\bigwedge(\Phi) \wedge \theta$ is consistent then $\times(\Phi|\theta) \equiv \bigwedge(\Phi) \wedge \theta$

(IS \times) if $\Phi \approx \Phi'$ and $\theta \equiv \theta'$ then $\times(\Phi|\theta) \equiv \times(\Phi'|\theta')$

These properties are called respectively (IC1), (IC0), (IC2) and (IC3) in (Konieczny & Pino Pérez 2002). Additional properties could be considered; they would, of course,

induce additional properties of *prioritized merging*; here, for the sake of brevity we stick to these four basic, undebatable properties of merging (besides, the aim of our paper is not to find an exhaustive list of properties for prioritized merging, but rather to show how it relates to, and generalizes, flat merging and iterated revision.) Moreover, this way of inducing prioritized merging operators from flat merging operators gives us a way of lifting representation theorems of Konieczny and Pino-Pérez (2002) from flat to prioritized merging (again, this issue is not central to our paper).

Proposition 2 \star satisfies (Cons), (IS), (PMon), (Opt), (Taut) and (RA) if and only if there exists a constrained flat merging operator \times such that

1. $\star(\sigma(n)) = \times(\sigma(n)|\top)$;
2. for each $i < n$, $\star(\sigma_{i \rightarrow n}) = \times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n}))$.

We denote by H the mapping $\times \mapsto \star$ defined above, that is, $\star = H(\times)$.

Proof:

(\Leftarrow) Let \star be defined from a flat merging operator \times by the above construction, that is, $\star = H(\times)$.

1. \star satisfies (Cons) because \times satisfies (Cons \times).
2. \star satisfies (PMon) because $\star(\sigma_{i \rightarrow n}) = \times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n}))$ and, due to (Succ \times), $\times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n})) \vdash \star(\sigma_{i+1 \rightarrow n})$.
3. We show by induction that \star satisfies (Opt). Assume $\bigwedge(\sigma)$ consistent. Then for all i , $\bigwedge(\sigma_{i \rightarrow n})$ is consistent. In particular, $\bigwedge(\sigma(n))$ is consistent and then $\star(\sigma(n)) = \times(\sigma(n)|\top) = \bigwedge(\sigma(n))$ by (Opt \times). Assume that for some i , $\star(\sigma_{i+1 \rightarrow n}) = \bigwedge(\sigma_{i+1 \rightarrow n})$. Then $\star(\sigma_{i \rightarrow n}) = \times(\sigma(i) | \bigwedge(\sigma_{i+1 \rightarrow n}))$, and since $\bigwedge(\sigma(i)) \wedge \bigwedge(\sigma_{i+1 \rightarrow n}) = \bigwedge(\sigma_{i \rightarrow n})$ is consistent, we have $\star(\sigma_{i \rightarrow n}) = \bigwedge(\sigma_{i \rightarrow n})$. Hence by backward induction we get $\star(\sigma) = \bigwedge(\sigma)$, therefore \star satisfies (Opt).
4. $\star(\sigma(i), \star(\sigma_{i+1 \rightarrow n})) = \times(\sigma(i) | \star(\star(\sigma_{i+1 \rightarrow n})))$; but $\star(\sigma_{i+1 \rightarrow n})$ is a consistent formula (since it has been proven at point 1 above that (Cons) holds), therefore, by (Opt) (which has been shown to hold as well, at point 3 above), $\star(\star(\sigma_{i+1 \rightarrow n})) \equiv \star(\sigma_{i+1 \rightarrow n})$. Thus, $\star(\sigma(i), \star(\sigma_{i+1 \rightarrow n})) = \times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n})) = \star(\sigma_{i \rightarrow n})$, which shows that (RA) holds.
5. We show (Taut) by induction. $\star(\sigma(n), \top) = \times(\sigma(n) | \star(\top)) = \times(\sigma(n) | \top)$ using (Opt); thus, $\star(\sigma(n), \top) = \star(\sigma(n))$. Now, assume $\star(\sigma_{i+1 \rightarrow n}, \top) = \star(\sigma_{i+1 \rightarrow n})$. Then $\star(\sigma_{i \rightarrow n}, \top) = \times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n}, \top)) = \times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n})) = \star(\sigma_{i \rightarrow n})$. Therefore, $\star(\sigma, \top) = \star(\sigma)$.
6. We show (IS) by induction. Let $\sigma \approx \sigma'$. $\star(\sigma(n)) = \times(\sigma(n) | \top) \equiv \times(\sigma'(n) | \top)$ using (IS \times), therefore $\star(\sigma(n)) \equiv \star(\sigma'(n))$. Assume that $\star(\sigma_{i+1 \rightarrow n}) \equiv \star(\sigma'_{i+1 \rightarrow n})$. Then $\star(\sigma_{i \rightarrow n}) = \times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n})) \equiv \times(\sigma'(i) | \star(\sigma'_{i+1 \rightarrow n}))$ using (IS \times) together with $\sigma(i) \approx \sigma'(i)$ and $\star(\sigma_{i+1 \rightarrow n}) \equiv \star(\sigma'_{i+1 \rightarrow n})$. Therefore, $\star(\sigma_{i \rightarrow n}) \equiv \star(\sigma'_{i \rightarrow n})$.

(\Rightarrow) Let \star be a prioritized merging operator satisfying (Cons), (IS), (PMon), (Opt), (Taut) and (RA). Define $\times = G(\star)$ by $\times(\Phi|\theta) = \star(\Phi, \{\theta\})$.

It is easily checked that \times satisfies (Cons \times), (Succ \times) and (Opt \times).

It remains to be checked that $H(G(\star)) = \star$. Let $\times = G(\star)$ and $\star' = H(G(\star))$. Then $\star'(\sigma(n)) = \times(\sigma(n) | \top) = \star(\sigma(n), \top) = \star(\sigma(n))$ by (Taut). Suppose that $\star'(\sigma_{i+1 \rightarrow n}) = \star(\sigma_{i+1 \rightarrow n})$. Then $\star'(\sigma_{i \rightarrow n}) = \times(\sigma(i) | \star'(\sigma_{i+1 \rightarrow n})) = \times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n}))$ (using the induction hypothesis $\star'(\sigma_{i+1 \rightarrow n}) = \star(\sigma_{i+1 \rightarrow n})$); now, $\times(\sigma(i) | \star(\sigma_{i+1 \rightarrow n})) = \star(\sigma(i), \star(\sigma_{i+1 \rightarrow n})) = \star(\sigma_{i \rightarrow n})$; hence, $\star'(\sigma_{i \rightarrow n}) = \star(\sigma_{i \rightarrow n})$ for all i , which proves that $\star' = \star$.

Therefore, there exists a constrained flat merging operator $\times (= G(\star))$ such that $\star = H(\times)$. ■

Proposition 2 gives us a practical way of defining prioritized merging operators, namely by inducing them from any flat constrained merging operator. A first question is whether the well-known prioritized merging operators reviewed above can be built up this way, or, equivalently, whether they satisfy the above properties or not. This turns out to be the case for $\star_{discrimin}$ and $\star_{leximin}$, and *a fortiori* for \star_{linear} . We first define the constrained flat merging operators \times_i and \times_c and then we show that they respectively induce $\star_{discrimin}$ and $\star_{leximin}$.

Definition 2 For any multiset Φ of propositional formulas, and any propositional formula θ ,

- a θ -consistent subset of Φ is a subset X of Φ such that $\bigwedge(X) \wedge \theta$ is consistent; $C(\Phi, \theta)$ denotes the set of all θ -consistent subsets of Φ ;
- $X \subseteq \Phi$ is a maximally θ -consistent subset of Φ if and only if $X \in C(\Phi, \theta)$ and there is no $X' \in C(\Phi, \theta)$ such that $X \subset X'$; $MaxCons(\Phi, \theta)$ denotes the set of all maximally θ -consistent subsets of Φ ;
- $X \subseteq \Phi$ is a maxcard θ -consistent subset of Φ if and only if $X \in C(\Phi, \theta)$ and there is no $X' \in C(\Phi, \theta)$ such that $|X| < |X'|$; $MaxcardCons(\Phi, \theta)$ denotes the set of all maxcard θ -consistent subsets of Φ .

Definition 3

- $\times_i(\Phi|\theta) = \bigvee\{(\bigwedge X) \wedge \theta \mid X \in MaxCons(\Phi, \theta)\}$
- $\times_c(\Phi|\theta) = \bigvee\{(\bigwedge X) \wedge \theta \mid X \in MaxcardCons(\Phi, \theta)\}$

We easily verify that \times_i and \times_c are constrained flat merging operators, that is, they satisfy (Cons \times), (Succ \times), (Opt \times) and (IS \times).

Now, we have the following result:

Proposition 3

1. $H(\times_i) = \star_{discrimin}$
2. $H(\times_c) = \star_{leximin}$

Before proving Proposition 3 we first state the following Lemma (the proof of which is omitted).

Lemma 1 If $S \subseteq \hat{\sigma}$ then we write $S = (S_1, \dots, S_n)$, where $S_j = S \cap \sigma(j)$, and we also write $S_{j \rightarrow n} = (S_j, \dots, S_n)$.

1. for all $j \leq n$,

$$S_{j \rightarrow n} \in \text{Max}(\succ_{\text{discrimin}}, \sigma_{j \rightarrow n}) \Leftrightarrow \begin{cases} S_j \in \text{MaxCons}(\sigma(j), \star_{\text{discrimin}}(\sigma_{j+1 \rightarrow n})) \\ S_{j+1 \rightarrow n} \in \text{Max}(\succ_{\text{discrimin}}, \sigma_{j+1 \rightarrow n}) \end{cases}$$
2. for all $j \leq n$,

$$S_{j \rightarrow n} \in \text{Max}(\succ_{\text{leximin}}, \sigma_{j \rightarrow n}) \Leftrightarrow \begin{cases} S_j \in \text{MaxcardCons}(\sigma(j), \star_{\text{leximin}}(\sigma_{j+1 \rightarrow n})) \\ S_{j+1 \rightarrow n} \in \text{Max}(\succ_{\text{leximin}}, \sigma_{j+1 \rightarrow n}) \end{cases}$$

Proof of Proposition 3: We give the proof for (1) in full detail; the proof for 2 is very similar. The result is shown by backward induction. Let $\star_i = H(\times_i)$ and $\text{Hyp}(j)$ the induction hypothesis

$$\text{Hyp}(j) : \star_{\text{discrimin}}(\sigma_{j \rightarrow n}) = \star_i \sigma_{j \rightarrow n}$$

$\text{Hyp}(n)$ is easily verified: $\star_{\text{discrimin}}(\sigma_{n \rightarrow n}) = \star_{\text{discrimin}}(\sigma(n)) = \bigvee(\bigwedge S | S \in \text{MaxCons}(\sigma(n))) = \times_i(\sigma(n) | \top) = \star_i(\sigma_{n \rightarrow n})$.

We now show that $\text{Hyp}(j+1)$ implies $\text{Hyp}(j)$. Assume $\text{Hyp}(j+1)$ holds. We have $\star_{\text{discrimin}}(\sigma_{j \rightarrow n}) = \bigvee\{\bigwedge(S) | S \in \text{Max}(\succ_{\text{discrimin}}, \text{Cons}(\sigma_{j \rightarrow n}))\}$. Now, using Lemma 1, $S \in \text{Max}(\succ_{\text{discrimin}}, \text{Cons}(\sigma_{j \rightarrow n}))$ holds iff $S = (S_j, S_{j+1 \rightarrow n})$ where $S_j \in \text{MaxCons}(\sigma(j), \star_{\text{discrimin}}(\sigma_{j+1 \rightarrow n}))$ and $S_{j+1 \rightarrow n} \in \text{Max}(\succ_{\text{discrimin}}, \sigma_{j+1 \rightarrow n})$. Therefore, $\star_{\text{discrimin}}(\sigma_{j \rightarrow n}) = \bigvee(\bigwedge S | S \in \text{MaxCons}(\sigma(i), \star_{\text{discrimin}}(\sigma_{j+1 \rightarrow n}))) = \times_i(\sigma(j) | \star_{\text{discrimin}}(\sigma_{j+1 \rightarrow n}))$. Thus, $\text{Hyp}(j)$ is satisfied. ■

Therefore, we recover $\star_{\text{discrimin}}$ and \star_{leximin} as the prioritized generalizations of maxcons and maxcard merging (which is what we expected). As a corollary, $\star_{\text{discrimin}}$ and \star_{leximin} satisfy (Cons), (IS), (Taut), (PMon), (Opt) and (RA).

There is no reason to stop here: we can take other flat merging operators and see what happens when we ‘‘prioritize’’ them. For instance, if we take the degenerate merging operator $\times_{\text{deg}}(\Phi | \theta) = \bigwedge \Phi \wedge \theta$ if consistent, $\times_{\text{deg}}(\Phi | \theta) = \theta$ otherwise, then the induced prioritized merging corresponds to linear merging, where each $\sigma(i)$ is considered as a single indivisible formula $\bigwedge \sigma(i)$. We can also take for \times a distance-based merging operator as in (Konieczny, Lang, & Marquis 2004; Delgrande & Schaub 2004), a quota-based one as in (Everaere, Konieczny, & Marquis 2005), etc. It should be noticed that we cannot recover \star_{bo} this way, because it fails to satisfy (RA) (even though it satisfies all other properties): take for example the SOPOB $\sigma = \langle b, \neg a, a \rangle$. Then $\star_{bo}(\sigma) = a$. Now, $\star_{bo}(\sigma_{2 \rightarrow 3}) = a$ and $\star_{bo}(b, \star_{bo}(\sigma_{2 \rightarrow 3})) = \star_{bo}(\langle b, a \rangle) = (a \wedge b)$.

Lastly, we may also consider the following property:

(Add)

$$\begin{array}{l} \text{If } \star(\sigma) \wedge \phi \not\vdash \perp \\ \text{then } \star(\sigma_{1 \rightarrow n-1}, \sigma(n) \cup \{\phi\}) \equiv \star(\sigma) \wedge \phi \end{array}$$

Thus, informally, if a formula is consistent with a merging operation, then the conjunction of that formula with the

merging is the same as doing the merging with the formula added to the highest priority level. Interestingly, if we restrict this postulate to 2-level SOPOBs (that is, SOPOBs of the form $\langle \alpha_1, \alpha_2 \rangle$) we recover the AGM postulates (using the reformulation of (Katsuno & Mendelzon 1991)):

Proposition 4 *Let \star be a prioritized merging operator satisfying the basic postulates for prioritized merging operators, and let \star be defined by*

$$\alpha \star \beta = \star(\alpha, \beta)$$

Then

1. \star satisfies the AGM postulates (R1), (R2), (R3) and (R4);
2. if \star satisfies (Add) then \star satisfies (R5) and (R6).

Therefore, the basic postulates together with (Add) imply the full set of AGM revision postulates. Note that (Add) is satisfied by most intuitive merging operators, including $\star_{\text{discrimin}}$, \star_{leximin} (a fortiori \star_{linear}) and \star_{bo} . Further, (Add) implies (Opt).

Back to iterated revision

Recall that the specific case of iterated revision is obtained by specializing prioritized merging to SOPOBs; we may then ask what the axiomatic studies for belief revision become in our framework. In the following we therefore consider prioritized merging operators \star restricted to SOPOBs. For the sake of simplicity, we assume furthermore that each individual observation is consistent.

Various researchers have studied a number of postulates that, apart from (C2), are generally considered desirable:

- (C1) if $\beta \vdash \mu$ then $\star(\sigma.\mu.\beta) = \star(\sigma.\beta)$;
- (C2) if $\beta \vdash \neg\mu$ then $\star(\sigma.\mu.\beta) = \star(\sigma.\beta)$;
- (C3) if $\star(\sigma.\beta) \vdash \mu$ then $\star(\sigma.\mu.\beta) \vdash \mu$;
- (C4) if $\star(\sigma.\beta) \wedge \mu$ is consistent then $\star(\sigma.\mu.\beta) \wedge \mu$ is consistent;
- (Rec) if $\beta \wedge \mu$ is consistent then $\star(\sigma.\mu.\beta) \vdash \mu$;
- (Ind) if $\star(\sigma.\beta) \wedge \mu$ is consistent then $\star(\sigma.\mu.\beta) \vdash \mu$.

(C1) - (C4) come from (Darwiche & Pearl 1997), (Rec) from (Nayak *et al.* 1996; Nayak, Pagnucco, & Peppas 2003), and (Ind) from (Jin & Thielscher 2005). (C2) has been criticized in several places, including (Freund & Lehmann 1994; Nayak *et al.* 1996; Konieczny & Pino Pérez 2002); see (Jin & Thielscher 2005) for a recent discussion.) It is shown in (Jin & Thielscher 2005) that (C3) and (C4) follow from (Ind).

Interestingly enough, all of these postulates, except the controversial (C2), are consequences of our properties (and even of our *basic* properties, except for (C1)).

Proposition 5

- (Opt) and (RA) imply (C1);
- (PMon), (Opt) and (Cons) imply (Rec), (Ind) and a fortiori (C3) and (C4).

Proof:

(Opt) and (RA) imply (C1) Let β and μ be consistent (recall that we consider only consistent inputs) such that $\beta \vdash \mu$. Then $\beta \wedge \mu$ is consistent (since it is equivalent to β). By (Opt) we have $\star(\mu, \beta) = \beta$. By (RA) we get $\star(\sigma, \mu, \beta) = \star(\sigma, \star(\mu, \beta)) = \star(\sigma, \beta)$.

(PMon) and (Opt) imply (Rec) Assume $\beta \wedge \mu$ consistent. By (GS), $\star(\sigma, \mu, \beta) \vdash \mu \wedge \beta$.

(PMon) and (Opt) imply (Ind) Assume $\star(\sigma, \beta) \wedge \mu$ consistent. By (Succ), $\star(\sigma, \beta) \vdash \beta$, therefore $\beta \wedge \mu$ is consistent. Then by (GS) we have $\star(\sigma, \mu, \beta) \vdash \beta \wedge \mu$, therefore $\star(\sigma, \mu, \beta) \vdash \mu$. ■

This means (C1), (C3), (C4), (Rec), and (Ind) are here given a new justification (and a new interpretation) in the more general context of prioritized merging. Actually, (C3), (C4), (Rec) and (Ind) are all consequences of the following:

If $\mu \wedge \beta$ is consistent then $\star(\sigma, \mu, \beta) \equiv \star(\sigma, \mu \wedge \beta)$ and $\star(\sigma, \mu, \beta) \vdash \mu \wedge \beta$.

This property reminds us that in the AGM framework, we have that if μ and β are inputs consistent with each other then Postulates 7 and 8 collapse into $(\sigma * \mu) * \beta = \sigma * (\mu \wedge \beta)$. There is also a property of probabilistic conditioning, whereby conditioning $P(\sigma \mid \mu)$ on β is equivalent to computing $P(\sigma \mid \mu \wedge \beta)$, when $\mu \wedge \beta$ is consistent.

We might then wonder about the addition of the postulate (C2). Reinterpreted in the context of prioritized merging, (C2) means that when we have two contradicting observations, the less reliable one is totally overridden by the more reliable one. Hence, observations are considered as atomic, and are undecomposable: if one wants to reject a part of the observation then we have to reject it all, so that, for example, when merging $a \wedge b$ together with the more reliable observation $\neg b$, we get $\neg b$ and we ignore the support on a given by the weaker observation. Note that if we merge $\sigma(1) = \{a, b\}$ with $\sigma(2) = \{\neg b\}$ using discrim, we get $\{a, \neg b\}$, while we get $\{\neg b\}$ if $\sigma(1) = \{a \wedge b\}$. The point is that in postulates of iterated revision (C1-4), revision is made by a formula, not by a multiset of formulas. Rewriting (C2) for multi-sets in the form

if $\sigma(3) \vdash \neg \wedge \sigma(2)$ then $\star(\sigma(1), \sigma(2), \sigma(3)) = \star(\sigma(1), \sigma(3))$

clearly highlights the debatable nature of this axiom.

In the case of SOPOBs, we get an interesting result: adding (C2) to the previous postulates restricts the set of possible prioritized merging operators to a single one, namely, linear merging.

Proposition 6 Assume σ is a SOPOB and \star satisfies (Opt), (PMon), (RA), (Cons) and (C2). Then $\star(\sigma) = \star_{linear}(\sigma)$.

Proof: By induction we show that $\star(\sigma_{i \rightarrow n}) = \star_{linear}(\sigma_{i \rightarrow n})$, denoted $A(i)$. $A(n)$ is obviously true. Suppose $A(i+1)$ holds. Then $\star(\sigma_{i \rightarrow n}) = \star(\sigma(i), \star(\sigma_{i+1 \rightarrow n}))$ by (RA), thus $\star(\sigma_{i \rightarrow n}) = \star(\sigma(i), \star_{linear}(\sigma_{i+1 \rightarrow n}))$ by $A(i+1)$. Now, two cases:

- if $\sigma(i) \wedge \star_{linear}(\sigma_{i+1 \rightarrow n})$ is consistent then by (Opt), $\star(\sigma_{i \rightarrow n}) = \sigma(i) \wedge \star_{linear}(\sigma_{i+1 \rightarrow n})$, therefore $\star(\sigma_{i \rightarrow n}) = \star_{linear}(\sigma_{i \rightarrow n})$;
- if $\sigma(i) \wedge \star_{linear}(\sigma_{i+1 \rightarrow n})$ is inconsistent then by (C2), $\star(\sigma_{i \rightarrow n}) = \star(\star_{linear}(\sigma_{i+1 \rightarrow n}))$; by (Cons), $\star_{linear}(\sigma_{i+1 \rightarrow n})$ is consistent, therefore, by (Opt), $\star(\star_{linear}(\sigma_{i+1 \rightarrow n})) = \star_{linear}(\sigma_{i+1 \rightarrow n})$; therefore $\star(\sigma_{i \rightarrow n}) = \star_{linear}(\sigma_{i+1 \rightarrow n}) = \star_{linear}(\sigma_{i \rightarrow n})$. ■

If (C2) is not required, other operators for SOPOBs are possible. From Proposition 2, we know that a prioritized merging operator satisfying (Opt), (PMon), (RA), (Taut) and (Cons) is defined from a flat merging operator. Here, when restricting POBs to SOPOBs, a constrained flat merging operator \times for SOPOBs maps any pair of consistent formulas φ and θ to a formula $\times(\varphi|\theta)$, satisfying

1. $\times(\varphi|\theta)$ is consistent;
2. $\times(\varphi|\theta) \vdash \theta$;
3. if $\varphi \wedge \theta$ is consistent then $\times(\varphi|\theta) \equiv \varphi \wedge \theta$.

\times can be seen as a “pre-revision” operator. Take for instance $\times_{Hammm}(\varphi, \theta)$, where $Mod(\times_{Hammm}(\varphi, \theta))$ is defined as the models of θ minimizing the Hamming distance to $Mod(\varphi)$. Then the associated prioritized merging operator \star_{Hammm} for SOPOBs satisfies (RA) and the basic properties, but not (C2): for instance, we have

$$\star_{Hammm}(a \wedge b, \neg a) = \neg a \wedge b \not\equiv \star_{Hammm}(\neg a)$$

Complex observations

Prioritized observation bases consist of very specific pieces of information: a propositional formula, together with a reliability level, that partitions the set of states into two subsets (states that are accepted and states that are rejected by the piece of information). More generally, we may want to allow pieces of information expressing that some states are more plausible than others, and thus allow for more than a simple dichotomy. Thus, for example, (Nayak 1994) addresses belief revision where the formula for revision comes with its own system of spheres; similarly, (Meyer 2001) looks at specific ways in which plausibility orderings may be combined; see also (Chopra, Ghose, & Meyer 2003). Thus, complex observations can be identified with epistemic states, and merging them comes down to merging epistemic states. Note that merging complex observations is more general than prioritized merging, which itself is more general than iterated revision.

We just give an example of a model handling complex observations, using *kappa functions* (see also (Williams 1995; Jin & Thielscher 2005) for iterated revision operators in this framework). A *kappa function* κ (Spohn 1988) is a mapping from S to $\bar{\mathbb{N}}$ such that $\min_{s \in S} \kappa(s) = 0$ (the higher $\kappa(s)$, the less plausible s .) A complex observation is here any kappa function κ_{obs} . The *combination* of two kappa functions κ_1 and κ_2 is defined (only when $\min_{s \in S} \kappa_1(s) + \kappa_2(s) < \infty$) by $(\kappa_1 \oplus \kappa_2)(s) = \kappa_1(s) + \kappa_2(s) - \min_{s \in S} \kappa_1(s) + \kappa_2(s)$.

\oplus is associative, therefore we can define the merging of n observations $\oplus(\kappa_1, \dots, \kappa_n)$. Define a *simple observation* $\kappa_{\varphi, k}$, where φ is a consistent formula and $k \in \overline{N}$, by $\kappa_{\varphi, k}(s) = k$ if $s \models \varphi$, $\kappa_{\varphi, k}(s) = \infty$ if $s \models \neg\varphi$. (Note that the restriction to simple observations implies a loss of generality (Laverny & Lang 2005).) Prioritized merging, in this framework, is obtained by considering only simple observations and taking sufficiently high values of k in $\kappa_{\varphi, k}$ so that (PMon) holds. Then, in this specific case, the prioritized merging operator obtained is $\star_{leximin}$, and specialization to SOPOBs gives \star_{linear} .

The above example could be reformulated likewise in the setting of (numerical) possibility theory. A possibility distribution π can be obtained by rescaling a kappa function inside the unit interval, namely (Dubois & Prade 1991) $\pi_{\kappa}(s) = 2^{-\kappa(s)}$. Then the most plausible states are when $\pi_{\kappa}(s) = 1$, and the impossible states when $\pi_{\kappa}(s) = 0$. An epistemic entrenchment is the ordering of formulas induced by necessity degrees $N(\alpha) = 1 - \max_{s \models \neg\alpha} \pi_{\kappa}(s)$. A simple observation is just a weighted formula in possibilistic logic. The view of the AGM setting proposed in (Dubois & Prade 1992) is precisely that of prioritized merging of a complex observation (an epistemic entrenchment) with a high priority simple observation in order to ensure the success postulate. In possibility theory, the above associative rule on kappa functions reads $(\pi_1 \oplus \pi_2)(s) = \frac{\pi_1(s) \cdot \pi_2(s)}{\max_{s \in S} \pi_1(s) \cdot \pi_2(s)}$. This merging operation was proposed by Dubois and Prade (1988) as generalizing the associative merging operation in the old MYCIN expert system. It is also very close to Dempster rule of combination of belief functions after Shafer (1976) (it only uses a different normalization factor).

These remarks indicate that enlarging the present framework of prioritized merging to complex observations has potential for bridging the gap between logical and numerical approaches to merging uncertain information.

Concluding remarks

The traditional account of iterated revision assumes that an agent receives information about a static world, in the form of a sequence of formulas that are incorporated into her belief set according to the principle of recency. We have argued that iterated revision should instead be treated more generally, as prioritized belief merging; thus an agent receives information about a static world, where the information (or perhaps the sources of information) come with varying degrees of reliability. This information is then incorporated with the agent's belief set, taking reliability into account. What is usually called iterated belief revision turns out to be a restriction of this more general setting.

We began by surveying approaches to prioritized merging, and provided a set of postulates, with respect to which existing approaches to merging were examined; as well it was shown how the full postulate set implies the AGM revision postulate set; further (given right associativity) it was shown that prioritized merging is expressible as sequential "flat" merging with integrity constraints. Given this framework, we re-examined postulates for iterated revision that have been previously proposed. Lastly, we briefly consid-

ered complex observations.

Our focus then has been on examining the relation between prioritized merging and iterated revision; to this end we presented a minimal set of merging postulates that were adequate for this purpose. What remains then for future work is to carry out a systematic examination of prioritized merging, along the lines of what (Konieczny & Pino Pérez 2002) do for flat merging. This presumably would entail presenting a general and, in some sense complete, set of postulates for prioritized merging. It is not difficult to come up with reasonable extended variants of the basic postulates. For example, the following extended version of (Opt) seems uncontentious:

if $\bigwedge(\sigma)$ is consistent then $\star(\sigma' \cdot \sigma) \equiv \star(\sigma', \wedge(\sigma))$.

As well, (Taut) would seem to generalize readily to:

$$\star(\sigma' \cdot \top \cdot \sigma) \equiv \star(\sigma' \cdot \sigma).$$

On the other hand, not all flat merging operators readily generalize to prioritized counterparts; for example, a prioritized counterpart to a majority merging operator, characterized by the postulate:

$$\exists k \text{ such that } \star(\sigma_{1 \rightarrow i-1}, \sigma_i \cup \{\alpha\}^k, \sigma_{i+1 \rightarrow n}) \vdash \alpha$$

is incompatible with (PMon). Hence, a general approach to prioritized merging may not be a straightforward extrapolation from unprioritized counterparts.

References

- Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change : partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.
- Benferhat, S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency management and prioritized syntax- based entailment. In *Proceedings of IJCAI'93*, 640–645.
- 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*.
- Brewka, G. 1989. Preferred subtheories: an extended logical framework for default reasoning. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI'89)*, 1043–1048.
- Chopra, S.; Ghose, A.; and Meyer, T. 2003. Non-prioritized ranked belief change. *Journal of Philosophical Logic* 32(4):417–443.
- Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89:1–29.
- Delgrande, J., and Schaub, T. 2004. Two approaches to merging knowledge bases. In *9th European Conference on Logics in Artificial Intelligence*, 426–438.
- Domotor, Z. 1980. Probability kinematics and representation of belief change. *Philosophy of Science* 47:284–403.

- Dubois, D., and Prade, H. 1988. Representation and combination of uncertainty with belief functions and possibility measures. *Computational Intelligence* 4(4):244–264.
- Dubois, D., and Prade, H. 1991. Epistemic entrenchment and possibilistic logic. *Artificial Intelligence* 50:223–239.
- 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. 1997. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *Int. J. of Approximate Reasoning* 17:295–324.
- Dubois, D.; Fargier, H.; and Prade, H. 2004. Ordinal and probabilistic representations of acceptance. *J. Artificial Intelligence Research* 22:23–56.
- 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.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2005. Quota and gmin merging operators. In *Proceedings of IJCAI'05*.
- Freund, M., and Lehmann, D. 1994. Belief revision and rational inference. Technical Report Tech. Report TR-94-16, Institute of Computer Science, The Hebrew University of Jerusalem.
- Freund, M. 2004. On the revision of preferences and rational inference processes. *Artificial Intelligence* 152:105–137.
- Friedman, N., and Halpern, J. 1996. Belief revision: A critique. In *Proceedings of KR'96*, 421–631.
- Friedman, N., and Halpern, J. 1999. Modeling beliefs in dynamic systems. part 2: revision and update. *JAIR* 10:117–167.
- 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.
- Katsuno, H., and Mendelzon, A. 1991. On the difference between updating a knowledge base and revising it. In *Proceedings of KR'91*, 387–394.
- Konieczny, S., and Pino-Pérez, R. 2000. A framework for iterated revision. *Journal of Applied Non-classical logics* 10(3-4).
- Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: a qualitative framework. *Journal of Logic and Computation* 12(5):773–808.
- Konieczny, S.; Lang, J.; and Marquis, P. 2004. Da2 merging operators. *Artificial Intelligence*.
- Laverny, N., and Lang, J. 2005. From knowledge-based programs to graded belief based programs, part I: on-line reasoning. *Synthese* 147(2):277–321.
- Lehmann, D. 1995a. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15(1):61–82.
- Lehmann, D. 1995b. Belief revision, revised. In *Proceedings of IJCAI'95*, 1534–1540.
- Meyer, T. 2001. On the semantics of combination operations. *Journal of Applied NonClassical Logics* 11(1-2):59–84.
- 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*.
- Nebel, B. 1991. Belief revision and default reasoning: Syntax-based approaches. In *Proceedings of KR'91*, 417–428.
- Nebel, B. 1994. Base revision operations and schemes: Semantics, representation, and complexity. In *Proceedings of ECAI'94*, 341–345.
- 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. 1981. A problem for relative information minimizers. *British J. of the Philosophy of Science* 33:375–379.
- Williams, M. 1995. Iterated theory-based change. In *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95)*, 1541–1550.