A sequential reversible belief revision method based on polynomials

Salem Benferhat and **Didier Dubois**
IRIT, université Paul Sabatier.
118 route de Narbonne. 31062 Toulouse
dubois@irit.fr and benferha@irit.fr

Odile Papini
LIM, université de la Méditerranée.
163 av. de Luminy, 13288 Marseille, cedex 09
papini@lim.univ-mrs.fr

Abstract

This paper deals with iterated belief change and proposes a drastic revision rule that modifies a plausibility ordering of interpretations in such a way that any world where the input observartion holds is more plausible that any world where it does not. This change rule makes sense in a dynamic context where observations are received, and the newer observations are considered more plausible than older ones. It is shown how to encode an epistemic state using polynomials equipped with the lexicographical ordering. This encoding makes it very easy to implement and iterate the revision rule using simple operations on these polynomials. Moreover, polynomials allow to keep track of the sequence of observations. Lastly, it is shown how to efficiently compute the revision rule at the syntactical level, when the epistemic state is concisely represented by a prioritized belief base. Our revision rule is the most drastic one can think of, in accordance with Darwiche and Pearl's principles, and thus contrasts with the minimal change rule called natural belief revision.

Introduction

One of the most fascinating problems in reasoning about knowledge is the one of belief change, and more specifically the one of iterated belief change. The most noticeable result obtained in the eighties by the AGM school is that rational revision steps require an ordering on interpretations. This ordering represents an epistemic state which distinguishes between interpretations which are more or less plausible. In the last ten years, after noticing that AGM revision could not be iterated, because it did not affect the underlying plausibility ordering, the focus point has been the construction of a rational approach to iterated belief revision. Modifying a plausibility ordering upon the arrival of a proposition μ that should be true can actually be done in several ways, called transmutations by (Williams 1994). In order to accept μ in the new epistemic state, the minimal requirement is that there are some models of μ which become more plausible than its countermodels. In natural revision, proposed by Boutilier (Boutilier 1993) but already pointed out by Spohn (Spohn 1988), only the best models of μ are made the most plausible interpretations. Alternative plausibility ordering revision rules have been proposed by (Williams 1994), based on Spohn's ordinal conditional functions, where the plausibility of all models of μ is affected. At the opposite of Boutilier's natural revision, another possible change rule, also evoked by Spohn (Spohn 1988), where upon receiving μ, *each* model of μ becomes more plausible than *all* countermodels of μ. This rule respects Darwiche and Pearl axioms (Darwiche & Pearl 1997) since it does not affect the relative ranks of models of μ, nor the relative countermodels of μ, and if a model of μ is more plausible than a countermodel, it remains so. Such a belief change rule makes sense in a dynamic context where observations are received, and the newer observations are considered more plausible than older ones. In this context, the meaning of the input observation suggests that at the time t when the observation is received, the real situation is necessarily one of the models of μ, so that μ should remain accepted whatever additional assumption is made at time t. This paper investigates this type of revision process from a semantic and syntactic point of view. It is shown that at the semantic level, the iterated belief change is made very easy if the plausibility orderings are encoded by means of lexicographical ordering of polynomials. The dual change rule, where older observations take precedence over newly acquired information, is studied and encoded likewise. Lastly, assuming that an epistemic state is represented by a prioritized set of formulas, it is shown how to perform the belief revision rule at the syntactic level, in full agreement with plausibility ordering change studied here.

Representation of epistemic states

Let $\mathcal{W}$ be the set of interpretations of propositional calculus, denoted $\mathcal{L}_{\mathcal{PC}}$. An epistemic state, denoted Ψ, encoding a set of beliefs about the real world (based on the available information), is represented by a total pre-order on $\mathcal{W}$, denoted $\leq_\Psi$. $\omega_1 \leq_\Psi \omega_2$ (resp. $\omega_1 <_\Psi \omega_2$) means that ω_1 is preferred to (strictly preferred to) ω_2. $Bel(\Psi)$ denotes a belief set associated to $\mathcal{W}$, representing agent's current beliefs, obtained from $\leq_\Psi$. It is a propositional formula [1] whose models are the most preferred ones w.r.t. $\leq_\Psi$, namely: $\text{Mod}(Bel(\Psi)) = \{\omega : \nexists \omega' \textit{ such that } \omega' <_\Psi \omega\}$.

Total pre-orders have been represented according to different points of view: binary relations, kappa-rankings which associate to each world an ordinal as in (Spohn 1988), (Williams 1994), possibility distributions (Dubois & Prade 1997) which associate to each world a degree between [0,1], vectors, etc. Changing a binary relation can't be easily expressed, it requires the use of another binary relation. Possibility distributions and kappa-rankings need to verify a so-called normalization condition (the existence of a world having a rank 0 in kappa-distributions, or degree 1 in possibility distributions), after each revision operation, which makes the process more complex. Moreover, these two representations in general are not reversible, namely there is no operations to reinstall the old ranking. Vectors can't be used because after shift operations their length grows. They involve difficulties for computation and comparison.

In this paper, we propose a suitable representation of $\leq_\Psi$ based on polynomials (Papini 1999), which allow to easy formalize the change of $\leq_\Psi$ according to the incoming observation. Each interpretation is assigned a weight, defined as a polynomial. Polynomials allow to keep track of the sequence of observations, they represent the history of the observations. They easily provide the models or the countermodels of successive observations. They allow to come back to previous orders, which is not possible with the other representations and they are tailored to representing the right and left shifts which formalize the change of total pre-orders in revision operations.

Weighting

Let B be the set $\{0,1\}$, and $B[x]$, the set of polynomials which coefficients belong to $\{0,1\}$. Polynomials p(x) in $B[x]$ have the form: $\text{p}(x) = \sum_{k=1}^{n} p_{-k}\ x^{-k} + \sum_{i=0}^{m} p_i\ x^i$. The use of negative and positive indices are necessarily to facilating shifting steps. A right shift is simply obtained with a multiplication by x, while a left shift is obtained with a multiplication by x^{-1}. As we will see in next sections, these shift operations are the basis for the computing of the new weights in the revision process. The following definition introduces the lexicographical ordering used to compare polynomials.

Definition 1 *Let $p(x)$, $p'(x) \in B[x]$, $p(x) <_B p'(x)$ iff $\exists i \in \mathbb{Z}$ such that $\forall j,\ j < i$, $p_j = p'_j$ and $p_i < p'_i$.*

Definition 2 *A weighting distribution is a function which associates to each epistemic state Ψ and ω an interpretation, a polynomial of $B[x]$ denoted by $p^\omega(\Psi)(x)$.*

The smaller is $\text{p}^\omega(\Psi)(x)$, w.r.t $<_B$ the more plausible is ω. When the initial epistemic state Ψ has no specified ranking on $\mathcal{W}$, but only the belief set $Bel(\Psi)$, the weighting distribution is defined as follows:
$\forall\omega \in \mathcal{W}$, if $\omega \in Mod(Bel(\Psi))$ then $p_0 = 0$ else $p_0 = 1$. Namely, the weight of the interpretations which are countermodels of $Bel(\Psi)$ is supposed to be lower than the weight of the interpretations which are models of $Bel(\Psi)$.

Remark 1 *If the initial epistemic state has a ranking represented by ordinals, we encode this ranking with polynomials using binary decomposition of ordinals.*

Equivalences

Iterated revision requires a carefully definition of the equivalence between two epistemic states, more formally:

Definition 3 *Let Ψ and Φ be two epistemic states. Ψ and Φ are* weakly equivalent, *denoted $\Psi \equiv_w \Phi$, iff $Bel(\Psi) \equiv Bel(\Phi)$, and Ψ and Φ are* strongly equivalent, *denoted $\Psi \equiv_s \Phi$, iff $\Psi \equiv_w \Phi$ and $\forall\omega \in \mathcal{W}$, $p^\omega(\Psi)(x) = p^\omega(\Phi)(x)$.*

The weak equivalence expresses that two beliefs sets are logically equivalent. The strong equivalence expresses that two epistemic states have equivalent associated beliefs sets and also have same weighting on $\mathcal{W}$. One of the problems arising with revision when epistemic states only consist in belief sets is that two equivalent epistemic states are revised in the same way. The definition of a weighting on $\mathcal{W}$ allows for the introduction of two kinds of equivalences, that makes it possible to solve this problem. Two weakly equivalent epistemic states shall be differently revised, on the other hand two strongly equivalent epistemic states shall be equally revised. The notion of strong equivalence is analogous to Darwiche and Pearl's equality between epistemic states (Darwiche & Pearl 1997) and to semantic equivalence in possibility logic (Dubois & Prade 1992).

Note that the function that assigns to each epistemic state Ψ the total pre-order on $\mathcal{W}$, denoted $\leq_\Psi$, defined by: $\omega_1 \leq_\Psi \omega_2$ iff $p^{\omega_1}(\Psi)(x) \leq_B p^{\omega_2}(\Psi)(x)$ is a faithful assignment, (Katsuno & Mendelzon 1991), with respect to $Bel(\Psi)$. Namely it satisfies:
(1) If $\omega_1, \omega_2 \models Bel(\Psi)$ then $\omega_1 =_\Psi \omega_2$;
(2) if $\omega_1 \models Bel(\Psi)$ and $\omega_2 \not\models Bel(\Psi)$ then $\omega_1 <_\Psi \omega_2$;
(3) $\Psi \equiv_s \Phi$ iff $\leq_\Psi = \leq_\Phi$.

In (Katsuno & Mendelzon 1991) Condition (3) only requires the logical equivalence.

Preferring newer information

The revision operation, first defined here, prefers the last item of information. The general philosophy here is that an old assertion is less reliable than a new one. We believe that in many cases it seems reasonable to decrease the confidence that one has in an item of information, as time goes by. However, this revision operation attempts to satisfy as many previous observations as possible. That is, an old observation persists until it becomes contradictory with a more recent one. The

[1] As Katsuno and Mendelzon (Katsuno & Mendelzon 1991), we use a propositional formula instead of a belief set which is a deductively closed set of formulas.

revision operation uses the history of the sequence of previous observations to perform revision.

Definition 4 *The revision of an epistemic state Ψ by a formula $\mu \in \mathcal{L}_{\mathcal{PC}}$, leads to a new epistemic state denoted $\Psi \circ_\triangleright \mu$. The modification of the weighting after revision by μ is:*

if $\omega \in Mod(\mu)$ then $p^\omega(\Psi \circ_\triangleright \mu)(x) = xp^\omega(\Psi)(x)$, otherwise $p^\omega(\Psi \circ_\triangleright \mu)(x) = xp^\omega(\Psi)(x) + 1$

The weights corresponding to the models of μ are right shifted and the weights corresponding to the countermodels of μ are right shifted and translated by 1.

Remark 2 *If $\mu = \perp$ then $Mod(Bel(\Psi \circ_\triangleright \mu)) = Mod(\Psi)$ and the total pre-order on $\mathcal{W}$ is preserved.*

Example 1 *Let Ψ be an epistemic state with a total pre-order $\leq_\Psi$ such that, $\omega_4 =_\Psi \omega_3 =_\Psi \omega_2 <_\Psi \omega_1$ and with associated belief set $Bel(\Psi) = a \vee b$. Let $\mu = \neg b$ and $\alpha = \neg a$ be propositional formulas. Let denote $r(x) = p^\omega(\Psi)(x)$, $q(x) = p^\omega(\Psi \circ_\triangleright \mu)(x)$ and $p(x) = p^\omega((\Psi \circ_\triangleright \mu) \circ_\triangleright \alpha)(x)$. The following array shows the changes of the total pre-order after a revision first by μ then by α:*

$\mathcal{W}$	a	b	$r(x)$	r_0	$q(x)$	q_0q_1	$p(x)$	$p_0p_1p_2$
ω_1	0	0	1	1	x	01	x^2	001
ω_2	0	1	0	0	1	10	x	010
ω_3	1	0	0	0	0	00	1	100
ω_4	1	1	0	0	1	10	$1+x$	110

In the above example, the columns $p_0\,p_1\,p_2$ give the total pre-order corresponding to the current epistemic state $(\Psi \circ_\triangleright \mu) \circ_\triangleright \alpha$, the columns $p_1\,p_2$ give the total pre-order corresponding to the previous epistemic state $\Psi \circ_\triangleright \mu$ and the columns of p_2 gives the total pre-order corresponding to the initial epistemic state Ψ. The values of coefficients of the polynomials show whether the interpretation satisfies (value 0) or not (value 1) the successive observations. For example, ω_2 satisfies α and $Bel(\Psi)$ but does not satisfy μ. Hence, the use of polynomials allows to come back to previous epistemic states. Indeed, let Ψ' be the actual epistemic state obtained after revising Ψ by μ. Then, μ can be recovered from Ψ' by defining models of μ in the following way: $\text{Mod}(\mu) = \{\omega : p^\omega(\Psi') <_B 1\}$, and the weighting distribution associated to the previous epistemic state Ψ can be recovered from Ψ' as follows: $p^\omega(\Psi) = x^{-1}p^\omega(\Psi')$ if $p^\omega(\Psi') <_B 1$ otherwise $p^\omega(\Psi) = x^{-1}(p^\omega(\Psi') - 1)$.

In our framework, the AGM postulates for epistemic states are rephrased as follows (Alchourron, Gärdenfors, & Makinson 1985), (Katsuno & Mendelzon 1991), (Papini & Rauzy 1995):

Modified AGM postulates for epistemic states

- $(R1\triangleright)$ $Bel(\Psi \circ_\triangleright \mu) \models \mu$.
- $(R2\triangleright)$ **If** $Bel(\Psi) \wedge \mu$ **is satisfiable, then** $\Psi \circ_\triangleright \mu \equiv_w \Psi \wedge \mu$.
- $(R3\triangleright)$ **If** μ **is satisfiable, then so is** $Bel(\Psi \circ_\triangleright \mu)$.
- $(R4\triangleright)$ **If** $\Psi_1 \equiv_s \Psi_2$ **and** $\mu_1 \equiv \mu_2$, **then,** $\Psi_1 \circ_\triangleright \mu_1 \equiv_s \Psi_2 \circ_\triangleright \mu_2$.
- $(R5\triangleright)$ $Bel(\Psi \circ_\triangleright \mu) \wedge \phi \models Bel(\Psi \circ_\triangleright (\mu \wedge \phi))$.
- $(R6\triangleright)$ **If** $Bel(\Psi \circ_\triangleright \mu) \wedge \phi$ **is satisfiable, then** $Bel(\Psi \circ_\triangleright (\mu \wedge \phi)) \models Bel(\Psi \circ_\triangleright \mu) \wedge \phi$.

$(R1\triangleright)$, $(R2\triangleright)$, $(R3\triangleright)$, $(R5\triangleright)$ and $(R6\triangleright)$ are the straightforward translation of the corresponding original AGM postulates. For $(R5\triangleright)$ and $(R6\triangleright)$ we assume that $\not\models \neg\mu$. In contrast, $(R4\triangleright)$ is a weaker version of original AGM postulate; it requires that the epistemic states be strongly equivalent in order to be equally revised.

Theorem 1 *Let Ψ be an epistemic state, μ be a formula of $\mathcal{L}_{\mathcal{PC}}$ and let $\leq_\Psi$ be the total pre-order on $\mathcal{W}$ defined as in proposition 1. Then, the operator $\circ_\triangleright$ verifies the postulates $(R1)-(R6)$ and $Mod(Bel(\Psi \circ_\triangleright \mu)) = min(Mod(\mu), \leq_\Psi)$.* [2]

Darwiche and Pearl (Darwiche & Pearl 1997) formulated postulates which constrain the relationships between two successive epistemic states, the straightforward translation of the corresponding original DP postulates in our framework, are the following:

DP postulates for iterated revision

- $(C1\triangleright)$ **If** $\alpha \models \mu$ **then** $(\Psi \circ_\triangleright \mu) \circ_\triangleright \alpha \equiv_w \Psi \circ_\triangleright \alpha$.
- $(C2\triangleright)$ **If** $\alpha \models \neg\mu$ **then** $(\Psi \circ_\triangleright \mu) \circ_\triangleright \alpha \equiv_w \Psi \circ_\triangleright \alpha$.
- $(C3\triangleright)$ **If** $Bel(\Psi \circ_\triangleright \alpha) \models \mu$ **then** $Bel((\Psi \circ_\triangleright \mu) \circ_\triangleright \alpha) \models \mu$.
- $(C4\triangleright)$ **If** $Bel(\Psi \circ_\triangleright \alpha) \not\models \neg\mu$ **then** $Bel((\Psi \circ_\triangleright \mu) \circ_\triangleright \alpha) \not\models \neg\mu$.

The postulates $(C1\triangleright)$, $(C2\triangleright)$, $(C3\triangleright)$ and $(C4\triangleright)$ in relationship with total pre-orders associated to two successive epistemic states are the following:

- $(CR1\triangleright)$ **If** $\omega_1 \models \mu$ **and** $\omega_2 \models \mu$ **then** $\omega_1 \leq_\Psi \omega_2$ **iff** $\omega_1 \leq_{\Psi\circ_\triangleright\mu} \omega_2$.
- $(CR2\triangleright)$ **If** $\omega_1 \models \neg\mu$ **and** $\omega_2 \models \neg\mu$ **then** $\omega_1 \leq_\Psi \omega_2$ **iff** $\omega_1 \leq_{\Psi\circ_\triangleright\mu} \omega_2$.
- $(CR3\triangleright)$ **If** $\omega_1 \models \mu$ **and** $\omega_2 \models \neg\mu$ **then** $\omega_1 <_\Psi \omega_2$ **only if** $\omega_1 <_{\Psi\circ_\triangleright\mu} \omega_2$.
- $(CR4\triangleright)$ **If** $\omega_1 \models \mu$ **and** $\omega_2 \models \neg\mu$ **then** $\omega_1 \leq_\Psi \omega_2$ **only if** $\omega_1 \leq_{\Psi\circ_\triangleright\mu} \omega_2$.

Theorem 2 *The operator $\circ_\triangleright$ verifies $(C1\triangleright) - (C4\triangleright)$ and its corresponding faithful assignment verifies $(CR1\triangleright) - (CR4\triangleright)$.*

The defined $\circ_\triangleright$ revision operation preserves the relative ordering between the models of the added formula. Furthermore, the relative ordering between the countermodels of the added formula is preserved, and the ordering between models and countermodels of the added formula does not change. Moreover, this operation provides a stronger constraint because each model of the added formula is preferred to all its countermodels.

Preferring oldest information

Preferring the last item of information is not always desirable, it may, in certain cases, lead to unacceptable conclusions. We define a new revision operation which is a dual revision operation of the one previously introduced. The general philosophy here is that a new observation is less reliable than an old one. We increase the confidence that one has in an item of information, as time goes by. However, this revision operation attempt to satisfy as many new observations as possible.

[2] $min(Mod(\mu), \leq_\Psi)$ contains all models that are minimal in $Mod(\mu)$ according to the total pre-order $\leq_\Psi$.

Definition 5 *The revision of Ψ by a formula μ, leads to a new epistemic state denoted $\Psi \circ_{\triangleleft} \mu$. The modification of the weighting after revision by μ is:*

if $\omega \in Mod(\mu)$ then $p^{\omega}(\Psi \circ_{\triangleleft} \mu)(x) = x^{-1}p^{\omega}(\Psi)(x)$,
otherwise $p^{\omega}(\Psi \circ_{\triangleleft} \mu)(x) = x^{-1}p^{\omega}(\Psi)(x) + 1$

This is different from the previous revision operation, since we now use left shifts (i. e. multiplication by x^{-1}).

Theorem 3 *The operator $\circ_{\triangleleft}$ verifies the postulates $(R2)$, $(R4)$ and $(R5)$.*

The postulates $(R1)$, $(R3)$ and $(R6)$[3] are not satisfied, because the new observation is not preferred in the next epistemic state. As older observations are preferred in the next epistemic state, the part of $Bel(\Psi)$ and the part of μ are inverted and hence new postulates can be formulated :

Proposition 1 *The operator $\circ_{\triangleleft}$ verifies:*

- $(R1\,b)$ $Bel(\Psi \circ_{\triangleleft} \mu) \models Bel(\Psi)$.
- $(R2\,b)$ **If** $Bel(\Psi) \wedge \mu$ **is not satisfiable, then** $\Psi \circ_{\triangleleft} \mu \equiv_w \Psi$.
- $(R3\,b)$ $Bel(\Psi)$ **is satisfiable iff** $Bel(\Psi \circ_{\triangleleft} \mu)$ **is satisfiable.**

Concerning iterated revision the following result holds:

Proposition 2 *The operator $\circ_{\triangleleft}$ verifies $(C3)$ and $(C4)$, and its corresponding faithful assignment verifies $(CR1)$, $(CR2)$, $(CR3)$ and $(CR4)$.*

Although the last observations are not preferred in the next epistemic state, the postulates $(C3)$ and $(C4)$ are satisfied. In contrast, the postulate $(C1)$ does not hold when $Bel(\Psi) \wedge \mu$ is satisfiable and $Bel(\Psi) \wedge \alpha$ is not satisfiable because $(\Psi \circ_{\triangleleft} \mu) \circ_{\triangleleft} \alpha \equiv_w \Psi \wedge \mu$ and $\Psi \circ_{\triangleleft} \alpha \equiv_w \Psi$. $(C2)$ does not hold when $Bel(\Psi) \wedge \mu$ is satisfiable and $Bel(\Psi) \wedge \alpha$ is satisfiable because $(\Psi \circ_{\triangleleft} \mu) \circ_{\triangleleft} \alpha \equiv_w \Psi \wedge \mu$ and $\Psi \circ_{\triangleleft} \alpha \equiv_w \Psi \wedge \alpha$. $(CR1) - (CR4)$ are satisfied since the defined $\circ_{\triangleleft}$ revision operation preserves the relative ordering between the models of the added formula.

Revision using $\circ_{\triangleright}$ and $\circ_{\triangleleft}$ together

In certain situations it seems reasonable to prefer a new observation and in others, the oldest observation has to be preferred. In order to use the two operations together, the definition of the weights associated to countermodels of μ are slightly modified in the following way: $p^{\omega}(\Psi \circ_{\triangleright} \mu)(x) = xp^{\omega}(\Psi)(x) + x^{MIN}$, and $p^{\omega}(\Psi \circ_{\triangleleft} \mu)(x) = x^{-1}p^{\omega}(\Psi)(x) + x^{MAX}$, where MIN and MAX are respectively the minimum and the maximum element of the set $\{i,\ i \in \mathbb{Z}$, such that $p_i^{\omega}(\Psi) \neq 0$, and $\omega \in \mathcal{W}\}$. When using only $\circ_{\triangleright}$ (resp. $\circ_{\triangleleft}$) then $MIN = 0$ (resp. $MAX = 0$) then we recover the previous definitions. Using these modifications, we have:

Proposition 3 *Let Ψ be an epistemic state, and two formulas μ, $\alpha \in \mathcal{L}_{\mathcal{PC}}$, $(\Psi \circ_{\triangleright} \mu) \circ_{\triangleleft} \alpha \equiv_s (\Psi \circ_{\triangleleft} \alpha) \circ_{\triangleright} \mu$.*

[3] $(R6)$ requires an additional condition to be satisfied that is $Bel(\Psi \circ_{\triangleleft} (\mu \wedge \phi)) \models \phi$.

This result shows that the order with wich the operators are used has no importance, since the $\circ_{\triangleright}$ revision operation uses right shifts and the $\circ_{\triangleleft}$ revision operation uses left shifts.

In previous sections we gave methods for constructing a new total pre-order on the interpretations. We now have to provide a syntactic counterpart of these semantical methods.

A syntactic representation of epistemic states

In this section, we give an alternative (but equivalent) representation of an epistemic state Ψ. Instead of explicitely specifying the total pre-order $\leq_{\Psi}$, the agent specifies a set of weighted formulas, called a weighted (or stratified) belief base and denoted by Σ_{Ψ}. Then we define a function κ which allows to recover $\leq_{\Psi}$ from Σ_{Ψ} by also associating to each interpretation ω a polynomial of B[x], that we denote by $\kappa_{\Sigma_{\Psi}}(\omega)(x)$. When this polynomial is equal to $\mathrm{p}^{\omega}(\Psi)(x)$ for each ω we say that Σ_{Ψ} is a compact (or syntactic) representation of $\leq_{\Psi}$.

Given this compact representation, we are interested in defining syntactic counterpart of $\circ_{\triangleright}$ (resp. $\circ_{\triangleleft}$), which syntactically transforms a weighted belief base Σ_{Ψ} and a new information μ, to a new weighted base, denoted by $\Sigma_{\Psi \circ_{\triangleright} \mu}$ corresponding to the new epistemic state $\Psi \circ_{\triangleright} \mu$. This new weighted base should be such that: $\forall \omega$ $\mathrm{p}^{\omega}(\Psi \circ_{\triangleright} \mu)(x) = \kappa_{\Sigma_{\Psi \circ_{\triangleright} \mu}}(\omega)(x)$.

In the following, we formally define the notion of weighted belief bases, and the function κ.

Definition 6 *A weighted belief base Σ_{Ψ} is a set of pairs $\{(\phi_i,\ p^{\phi_i}(\Psi)(x)) : i = 1, ..., n\}$ where ϕ_i is a propositional formula, and $p^{\phi_i}(\Psi)(x)$ is a non-null polynomial of B[x] (i.e., different from the polynomial 0).*

Polynomials associated to formulas are compared according to Definition 1. When $\mathrm{p}^{\phi}(\Psi)(x) >_B p^{\psi}(\Psi)(x)$, we say that ϕ is more important (certain, recent, has a higher priority etc) than the belief ψ. A weighted base Σ_{Ψ} is said to be consistent (resp. to entail ϕ) if its classical base (by forgetting the weights) is also consistent (resp. entails ϕ). Note that Σ_{Ψ} is not necessarily deductively closed. Moreover, nothing prevents Σ_{Ψ} from containing two weighted formulas $(\phi, p^{\phi}(\Psi)(x))$ and $(\psi, p^{\psi}(\Psi)(x))$ such that ϕ and ψ are classically equivalent, but having different weights $p^{\phi}(\Psi)(x) \neq p^{\psi}(\Psi)(x)$. In this case, we will see later that the least important belief can be removed from the weighted belief base.

Definition 7 *Let Σ_{Ψ} be a weighted belief base. The total pre-order $\leq_{\Psi}$, associated to Σ_{Ψ}, is obtained by attaching to each ω a weight $\kappa_{\Sigma_{\Psi}}(\omega)(x)$ defined by: $\kappa_{\Sigma_{\Psi}}(\omega)(x) = max\{p^{\phi_i}(\Psi)(x) : (\phi_i,\ p^{\phi_i}(\Psi)(x)) \in \Sigma_{\Psi}$ and $\omega \not\models \phi_i\}$, where by convention max($\emptyset$)=0.*

This semantics is basically the same as the one used in possibilistic logic (Dubois et al., 1994), in System Z (Pearl 1995) and for generating a complete epistemic entrenchment relation from a partial one (Williams 1994). Indeed, all these approches share the same idea, where they associate to each interpretation the weight of the most important formula falsified by the interpretation. The lowest is the weight of an interpretation, the most plausible it is and the preferred it is. In particular, models of Σ_Ψ (namely those with a weight equal to 0) are the most preferred ones.

Example 2 *Let :* $\Sigma_\Psi = \{(\neg a \vee \neg b,\ x+1), (\neg a,\ 1), (\neg b \vee a,\ 1), (\neg a \vee \neg b,\ x), (\neg b,\ x), (a \vee b,\ x^2)\}$
Then: $\kappa_{\Sigma_\Psi}(ab) = max\{1, x, x+1, x\} = x+1$. $\kappa_{\Sigma_\Psi}(a\neg b) = max\{1\} = 1$. $\kappa_{\Sigma_\Psi}(\neg ab) = max\{x, 1\} = 1$. $\kappa_{\Sigma_\Psi}(\neg a \neg b) = max\{x^2\} = x^2$. *Using Definition 1,* $\neg a \neg b$ *is the preferred one, then* $\neg ab$ *and* $\neg ba$ *are less preferred than* $\neg a \neg b$ *and lastly* ab *is the least preferred one. Note that there is no* ω *such that* $\kappa_{\Sigma_\Psi}(\omega)(x) = 0$. *This expresses the fact that the belief base* Σ_Ψ *is inconsistent.*

Computing Bel(Ψ) syntactically

Given Σ_Ψ as a compact representation of $\leq_\Psi$, we propose to compute Bel(Ψ) directly from Σ_Ψ, such that:

$\text{Mod}(\text{Bel}(\Psi)) = \{\omega : \nexists \omega' \ s.\ t.\ \kappa_\Sigma(\omega') < \kappa_\Sigma(\omega)(x)\}$.

But first, we proceed to some pre-processing steps which make the computation easier, and which reduce the size of revised knowledge bases. These preprocessing steps consist in removing useless (or redundant) formulas. These formulas are tautologies and the so-called subsumed beliefs are defined by:

Definition 8 *A formula* $(\phi,\ p^\phi(\Psi)(x))$ *is subsumed in* Σ_Ψ *if it can be classically entailed from formulas of* Σ_Ψ *having a weight greater than* $p^\phi(\Psi)(x)$.

Theorem 4 *Let* Σ_Ψ *be a weighted base. Let* Σ'_Ψ *be a new base obtained from* Σ_Ψ *by removing tautologies and subsumed formulas. Then* Σ_Ψ *and* Σ'_Ψ *are equivalent, in the sense that* $\forall\omega$ *we have:* $\kappa_{\Sigma_\Psi}(\omega)(x) = \kappa_{\Sigma'_\Psi}(\omega)(x)$.

We denote by Σ^*_Ψ the weighted subbase obtained by removing tautologies and subsumed formulas from Σ_Ψ. Σ^*_Ψ is a partial epistemic entrenchment in the sense of (Williams 1995). It is easy to imagine an algorithm that computes Σ^*_Ψ from Σ_Ψ

Example 3 *Let us consider the weighted belief base of the previous example. The only subsumed formulas are* $(\neg a \vee \neg b,\ x)$ *(which is entailed by* $(\neg a \vee \neg b,\ x+1)$*) and* $(\neg b,\ x)$ *(which is entailed by* $(\neg b \vee a,\ 1)$*). The previous algorithm returns the final subbase:*
$\Sigma^*_\Psi = \{(\neg a \vee \neg b,\ x+1), (\neg a,\ 1), (\neg b \vee a,\ 1), (a \vee b,\ x^2)\}$.

The removing of tautologies and subsumed formulas allows us a direct computation of Bel(Ψ).

Theorem 5 *If* Σ^*_Ψ *is consistent, then Bel(*Ψ*) is the classical base (i.e., without weights) associated to* Σ_Ψ. *If* Σ^*_Ψ *is not consistent, then let Minweight be the set of beliefs in* Σ^* *having minimal weights. Then Bel(*Ψ*) is the classical base of* Σ^*_Ψ *- Minweight.*

Example 4 *(continued) Since* Σ^*_Ψ *is inconsistent, then: Minweight* $= \{(a \vee b,\ x^2)\}$. *Therefore, Bel(*Ψ*) is the classical base (by forgetting weights) of:* $\Sigma^*_\Psi -$ $Minweight = \{(\neg a \vee \neg b,\ x+1), (\neg a,\ 1), (\neg b \vee a,\ 1)\}$. *Clearly, Bel(*$\Psi$*) has exactly one model which is* $\neg a \neg b$. *Moreover, it is easy to check that* $\neg a \neg b$ *has the minimal weight in* κ_{Σ_Ψ} *computed previously in Example 2.*

Syntactic counterpart of $\circ_\triangleright$ and $\circ_\triangleleft$

This section gives the syntactic counterpart of $\circ_\triangleright$. Let us illustrate the construction of $\Sigma_{\Psi\circ_\triangleright\mu}$ when Σ_Ψ only contains one formula $\{(\phi,\ 1)\}$ representing the initial belief set. Given a new observation μ, let the reader check that the new ordering $\leq_{\Psi\circ_\triangleright\mu}$, in the semantical construction, is encoded by: $p^\omega(\Psi \circ_\triangleright \mu)(x) = 0$ if $\omega \models \phi \wedge \mu$; $p^\omega(\Psi \circ_\triangleright \mu)(x) = x$ if $\omega \models \neg\phi \wedge \mu$; $p^\omega(\Psi \circ_\triangleright \mu)(x) = 1$ if $\omega \models \phi \wedge \neg\mu$; $p^\omega(\Psi \circ_\triangleright \mu)(x) = x+1$ if $\omega \models \neg\phi \wedge \neg\mu$. Now, let us see how to recover this ordering by building a weighted belief base $\Sigma_{\Psi\circ_\triangleright\mu}$ and using the function κ. Recall that κ is defined with respect to the strongest falsified belief. Therefore, in order to recover $\leq_{\Psi\circ_\triangleright\mu}$, one should assign to $\phi \vee \mu$ the highest weight (therefore each countermodel of $\phi \vee \mu$, namely each model of $\neg\phi \wedge \neg\mu$, gets the highest weight hence the least preferred interpretation), then $\neg\phi \vee \mu$ will get a smaller weight, which has a weight greater than $\phi \vee \neg\mu$. Therefore, we get the following weighted base:

$\Sigma_{\Psi\circ_\triangleright\mu} = \{(\phi \vee \mu,\ x+1), (\neg\phi \vee \mu, 1), (\phi \vee \neg\mu, x)\}$

It is easy, to check that: $\forall\omega,\ \kappa_{\Psi\circ_\triangleright\mu}(\omega)(x) = p^\omega(\Psi \circ_\triangleright \mu)(x)$.

Note that models of this weighted base are those which satisfy $\phi \wedge \mu$, and by definition of the function κ they get lowest rank and hence there is no need for additional beliefs. Besides, we can easily check that the above base can be simplified into an equivalent one (in the sense that they have the same κ) which is:
$\Sigma_{\Psi\circ_\triangleright\mu} = \{(\phi \vee \mu,\ x+1), (\mu, 1), (\phi \vee \neg\mu, x)\}$
Intuitively, replacing $(\neg\phi \vee \mu, 1)$ by $(\mu, 1)$ is justified by the fact that we already have $\phi \vee \mu$ with the highest rank, and from $(\phi \vee \mu, x+1)$ and $(\neg\phi \vee \mu, 1)$ we deduce $(\mu, 1)$, hence adding $(\neg\phi \vee \mu, 1)$ is equivalent to only add $(\mu, 1)$. Moreover, from $(\mu, 1)$ and $(\phi \vee \neg\mu, x)$ we also deduce (ϕ, x) To summarize, the new weighted base contains three parts: first, the old formula ϕ, with the smallest weight equal to x $p^\phi(\Psi)(x) =$ x, then the new formula μ, with a weight equal to 1, and finally, the disjunction $\phi \vee \mu$, with the highest weight equal to x $p^\phi(\Psi)(x) + 1$=x+1.

The following definition generalizes the previous result to the case where the weighted belief base Σ_Ψ contains more than one belief.

Definition 9 *The weighted base* $\Sigma_{\Psi\circ_\triangleright\mu}$ *associated to the epistemic state* $\Psi \circ_\triangleright \mu$ *is composed of:*

- *the new observation* μ *with a rank:* $p^\mu(\Psi \circ_\triangleright \mu)(x) = 1$.
- *all the pieces of information of* Σ_Ψ *however with the new rank, namely for each belief* ϕ *in* Σ_Ψ*:*
  $p^\phi(\Psi \circ_\triangleright \mu)(x) =x\ p^\phi(\Psi)(x)$
- *all the possible disjunctions between beliefs* ϕ *of* Σ_Ψ *and* μ *with the following weights:*
  $p^{\phi\vee\mu}(\Psi \circ_\triangleright \mu)(x) =x\ p^\phi(\Psi)(x) + 1$

Once $\Sigma_{\Psi\circ_\triangleright\mu}$ is built we remove tautologies and subsumed beliefs. The following proposition shows that $\Sigma_{\Psi\circ_\triangleright\mu}$ allows us to recover the total pre-order associated to the epistemic state $\Psi \circ_\triangleright \mu$ syntactically.

Theorem 6 *Let Σ_Ψ be the weighted base associated to an epistemic state Ψ, such that $\forall\omega$, $p^\omega(\Psi)(x) = \kappa_{\Sigma_\Psi}(\omega)(x)$. Let μ be a new formula. Then for each ω:*

$$p^\omega(\Psi \circ_\triangleright \mu)(x) = \kappa_{\Sigma_{\Psi\circ_\triangleright\mu}}(\omega)(x)$$

The following definition gives the syntactic counterpart of $\Psi \circ_\triangleleft \mu$ and characterizes the structure of $\Sigma_{\Psi\circ_\triangleleft\mu}$, which is exactly the same as the one of $\Sigma_{\Psi\circ_\triangleright\mu}$ except that the weighting is not the same.

Definition 10 *The weighted base $\Sigma_{\Psi\circ_\triangleleft\mu}$ associated to the epistemic state $\Psi \circ_\triangleleft \mu$ is composed of:*

- *the new observation μ with a rank: $p^\mu(\Psi\circ_\triangleleft\mu)(x) = 1$.*
- *all the pieces of information ϕ of Σ_Ψ with the rank: $p^\phi(\Psi \circ_\triangleleft \mu)(x) = x^{-1}p^\phi(\Psi)(x)$*
- *all the possible disjunctions between beliefs ϕ of Σ_Ψ and μ with the weights: $p^{\phi\vee\mu}(\Psi \circ_\triangleleft \mu)(x) = x^{-1}p^\phi(\Psi)(x) + 1$*

The following proposition shows that $\Sigma_{\Psi\circ_\triangleleft\mu}$ also allows us to recover the total pre-order associated to the epistemic state $\Psi \circ_\triangleleft \mu$ syntactically.

Theorem 7 *Let Σ_Ψ be the weighted base associated to an Ψ, and μ be a formula. Then for each ω:*
$p^\omega(\Psi \circ_\triangleleft \mu)(x) = \kappa_{\Sigma_{\Psi\circ_\triangleleft\mu}}(\omega)(x)$

Concluding discussions

This paper proposes a revision rule, with a practical syntactic counterpart, that operates a maximal shift of all models of a new observation, while retaining their relative ordering. This type of belief change rule is the most drastic one can think of, in accordance with Darwiche and Pearl's principles, while Boutilier's natural revision is the least refined change rule, whereby only the best models of μ are made maximally plausible. Our changing rule considers that the input is strongly believed (but its strength decreases as time goes by), since each model of μ is preferred to all countermodels of μ. In contrast, Boutilier's natural revision the input is only weakly believed, since only the best models of μ are preferred to μ. Milder changing rules have been proposed by Williams (1994) and (Dubois and Prade, 1997) where they explicitely specify the level of acceptance of the input.

The proposed change rule can be viewed as based on the ceteris-paribus principle, that has been applied to representing preferences (C. Boutilier & al. 1997), whereby preferring μ means that any choice satisfying μ is better than any choice satisfying $\neg\mu$. It suggests that our revision approach might be relevant as well, not only for encoding a sequence of timed obsevations, but also for the update preference.

The use of polynomials allows to solve the drawback addressed by Spohn (Spohn 1988) concerning the reversibility of the changing rules. Indeed, we have shown that it is always possible to go back to previous epistemic states, while this is not be possible if WOP (well-ordered partitions) or ordinals have been used. It might be interesting to see if other revision tools can be efficiently encoded by means of simple operations on polynomials. Boutilier's natural belief revision can be easily encoded by shift operations on polynomials and syntactic counterpart to Boutilier's natural belief revision can be provided.

References

Alchourron, C.; Gärdenfors, P.; and Makinson, D. 1985. On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision. *Journal of Symbolic Logic* 50:510–530.

Boutilier, C. 1993. Revision Sequences and Nested Conditionals. In *Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence (IJCAI'93)*, 519–531.

C. Boutilier, R. B., and al. 1997. A constraint-based approach to preference elicitation and decision making. In *Working notes of the AAAI'97 Spring Symp. Series on Qualitative Preferences in Deliberation and Practical Reasoning. Stanford. Mar. 24-26*, 19–28.

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

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

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

Katsuno, H., and Mendelzon, A. 1991. Propositional Knowledge Base Revision and Minimal Change. *Artificial Intelligence* 52:263–294.

Papini, O., and Rauzy, A. 1995. Revision in Extended Propositional Calculus. In Froidevaux, C., and Kholas, J., eds., *Proceedings of ECSQARU'95*, volume 946, 328–335. LNAI.

Papini, O. 1999. Iterated revision operations stemming from the history of an agent's observations. In Williams, M. A., and Rott, H., eds., *Frontiers of Belief Revision*, (to appear). Kluwer Acad. Publisher.

Pearl, J. 1995. System z: A natural ordering of defaults with tractableapplications to default reasoning. In Parikh, R., ed., *Proc. of the 3rd Conf. on Theoretical Aspects of Reasoning about Knowledge (TARK'90)*, 121–135. Morgan Kaufmann.

Spohn, W. 1988. Ordinal conditiona functions: a dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, 105–134. Kluwer. Acad. Publisher.

Williams, M. A. 1994. Transmutations of Knowledge Systems. In et al., J. D., ed., *Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'94)*, 619–629. Morgan Kaufmann.

Williams, M. A. 1995. Iterated Theory Base Change:A Computational Model. In *Proc. of 14th Int. Joint Conf. on Artificial Intelligence (IJCAI'95)*, 1541–1547.