Prioritized Dynamic Retraction Function on Non-monotonic Information Updates

Giuseppe Primiero

ABSTRACT. In this paper a model for updates on prioritized belief sets and retractions thereof is introduced, using the standard format of Adaptive Logics. The core of the update retraction procedure is represented by abnormal expressions derivable in the language: they express updates with information that contradict previously derived contents. The adaptive strategy aims at restricting the validity of these formulas by focusing at each decreasing degree on the update which is the most rational to retract in order to restore consistency as soon as possible.

1 Introduction

The formalization of the dynamics of logical reasoning is nowadays a very extensive area of research, exploring various aspects of knowledge processes. Ongoing research focuses on the formulation of successful methodologies to revise a knowledge base (or belief set) in the light of new inconsistent information received, and to define an appropriate contraction operation in view of previously performed updates.

The basic requirement of consistency demands a process of resolution which, on the one hand, does not represent agents as completely irrational, so that they be able to maintain non-contradictory beliefs; on the other hand, the rationality principle at the basis of the retraction process might be related to different criteria: in some cases consisting in preserving older information, in other cases in accepting more recent data, or being dictated by trust ascribed to information sources. A less explored and more complex task is the representation of the agent's ability to choose among possible different updates. In general, the representation of inconsistencies related to a temporal order is needed for a more realistic formulation of such processes.

The present paper defines a logic called **AIUR** (*Adaptive Information Update Retraction*) which simulates the choice of the most rational retraction to perform on possible updates, in order to restore consistency on a belief set. The adaptive framework represents an alternative approach to other systems known from the literature: in particular, it allows the representation of inconsistent information processing without trivilizing belief sets, for which the AGM-paradigm from [1] provides no realistic modelling.

The standard set of doxastic actions for belief change in AGM contains *expansion* $(K + A)$, *revision* $(K * A)$ and *contraction* $(K \div A)$ for a belief set K and a belief A. The standard contraction operator satisfies the well-known AGM-postulates:

- K1 Closure: if K is logically closed, then so is $K \div A$ for all A ($K \div A = Cn(K \div A)$);
- K2 Inclusion: the result of a contraction-operation should always be included in the belief set prior to that operation ($(K \div A) \subseteq K$);
- K3 Vacuity: if some given content does not belong to the belief set, then contraction with respect to such content results in the belief set itself (If $A \notin K$ then $K \div A = K$);
- K4 Success: no logical consequence of the set of beliefs can be contracted – or else if new information comes in, then it must be incorporated into our belief set (If $\not\vDash A$ then $A \notin K \div A$);
- K5 Recovery: the removal of a belief A followed by the reintroduction of the same belief should lead to the original belief set ($K \subseteq Cn((K \div A) \cup \{A\})$);
- K6 Extensionality: logically equivalent sentences lead to equivalent contractions (If $\vDash A_1 \leftrightarrow A_2$ then $K \div A_1 = K \div A_2$).

The Inclusion Postulate is usually taken for granted in the AGM-paradigm. On the other hand, the Recovery Postulate is justified by an appeal to the following:

DEFINITION 1 (Principle of Informational Economy). Keep the loss of information to a minimum.

If this principle becomes overriding, the belief set resulting from contraction of K by A should be a maximal subset of K that does not imply A. This restriction is shown to lead to maxichoice contraction operations in [10], and usually it is considered too strong. In [21], the status of the Recovery postulate is analysed in view of the misapplication of maxichoice contraction operations to belief sets (theories): when contraction is applied to finite bases (rather than theories), recovery becomes vacuously true, and its body fails in general; on the other hand, when contraction is applied to sets closed under logical consequence, recovery becomes intuitive and morever useful.

Within the dominant AGM-paradigm, the difference between belief states assumed to be logically closed sets of sentences and non-closed sets (bases) is crucial in the treatment of inconsistent data. It has been argued at lenght in the literature that in real world situations knowing agents deal with inconsistent beliefs, without believing everything (explosion, see e.g. [16]). The case of inconsistencies in non-prioritized belief bases can be dealt with by enough applications of contraction on dispensable elements. One of the first formalization of this case is introduced in [15], where the operation of restoring consistency is called *consolidation* ($A!$), and it is performed by contraction on contradiction: $A! = A \div \bot$. The shortcoming is due to the underlying logic: consolidation cannot be used in a satisfactory way on inconsistent belief sets because all inconsistent sets are trivially equal in a classical setting, hence all distinctions are lost. A variation on the theme of inconsistent belief bases is given in [16], where a set of *local operators* for contraction, consolidation, revision and semi-revision is introduced. This is obtained by defining logical compartments in the belief base around a sentence, that is the subset of the base that is relevant to that sentence.[1] Consolidation is thus defined as an operation on the kernel for a belief base, determined by an inference operation which preserves consistency, inclusion and so-called core-retainment, corresponding to a form of compactness.[2]

A different approach to the localization of inconsistent data is introduced in [27]. The state of the beliefs on inconsistent information is considered a distinct element from the inconsistent data itself: even with respect to such data, beliefs are normally consistent. On the basis of this idea, the system introduced in [27] formulates operations to extract consistent beliefs form the represented information, thus differing from other treatments admitting inconsistencies by the use of a paraconsistent logic.[3] This holds – as in the logic AIUR here introduced – for beliefs closed under logical consequence, that is for belief sets, which eliminates the restriction from consolidation. To this aim, an inconsistency-tolerant 4-valued logic is used, the first degree entailment logic, which extends classical logic by admitting partial and inconsistent valuations keeping inconsistencies local. It defines methods to select only consistent three-valued valuations on finite information states. In particular, *extractors* are defined to obtain the consistent modulo of an inconsistent finite state: a consistency forcer function is defined to ignore contradictory valuation, another to take maximal consistent subvaluations of contradictory ones. The analogy with the adaptive logic introduced in this paper goes therefore at various levels: structurally, both use methods to select on the consequence relation to establish persistance of consistent data, and both admit these initial data to be inconsistent. The differences concern essentially the representation of such inconsistency in the language: whereas for the logic in [27] the valuations are the meta-theoretical means to represent inconsistent information, and minimization is the way to reduce it to relevant data; in the adaptive logic AIUR, inconsistencies will be given as derivable formulas in the language, and their form will implement the possible updates the agent is faced with. At the level of the selection method, the prioritized structure based on the temporal order of updates in AIUR allows to choose between the different consequence relations generated by an inconsistent premise set, whereas in the case of the extractors functions

[1] See also the related notion of *kernel* in [14] and the application of belief change strategies to model-based diagnosis in [31].

[2] See [16], pp. 60-62.

[3] See for example [23].

the resulting possibilities are treated on equal terms, and only some sort of meta-theoretical considerations can force to choose the "most consistent" valuations.[4]

It clearly appears from this analysis that the introduction of prioritized structure leads to a very different treatment of consistency-restore operations. In the AGM-paradigm, the formulation of priority relations is given by the introduction of the following principles:

DEFINITION 2 (Principle of Indifference). Contents held in equal regard should be treated equally.

DEFINITION 3 (Principle of Preference). Contents held in higher regard should be afforded a more favourable treatment.

Their combination with the Principle of Informational Economy requires the system to satisfy both the Inclusion and the Recovery Postulates. A contraction operator that satisfies the Inclusion Postulate but not the Recovery Postulate is usually called a *withdrawal* operator.[5] When retracting a belief A from a belief set K, there might be other beliefs in K that entail A, but one might not want to retrace all such beliefs back. By a result of Makinson which uses the Levi identity $K^*_A = (K^-_{\neg A})^+_A$, the set of withdrawal functions generates a unique revision function satisfying all the AGM-postulates.

Along with AGM contraction and recovery, the literature on similar functions or variant models for prioritized bases/sets is very extensive: the possible-worlds interpretation of spheres in [13]; the epistemic entrenchment relation from [11]; its variant in preferential bases presented in [24] and the related axiomatization from [9]; the application of withdrawal on belief bases in [32].[6] One way to satisfy the Principle of Informational Economy is to preserve the use of prioritized sets of beliefs (semantically represented by a preorder on their interpretations), and to allow updates and retractions on the basis of a temporal order. This is the structure of the logic **AIUR**.

The AGM functions on the contraction operator can be notoriously represented by total preorders on a finite set of propositions (their interpretations), see e.g. [13]. An agent keeping track of the information received, and updating his basis first with $\neg A$ and then with A, can infer in a prioritized way and therefore reject $\neg A$ (assuming she has a preference for the most recent information). The retraction of the information A will natu-

[4]It is interesting to notice that in [27] the possible development of, and comparison with adaptive semantics is foreseen; nonetheless, it is also stressed how these semantics "differs widely". We provide here an adaptive logic for prioritized bases, which therefore implements further particular cases that were not included in the basic ideas from [27]. For some more relation of adaptive logics with standard belief revision theory, see [8]. For another model of retraction on inconsistent bases on the basis of a logic for default reasoning, see [6].

[5]See [21],[12] and [5].

[6]For a complete overview of the definitions of contraction and withdrawal functions, see [25].

rally lead to increase support in the belief that $\neg A$, which depicts therefore an internal dynamics with respect to the reasons held against other contents. This dynamics, called *liberation* in [5], is especially intuitive when retraction is performed on belief sets updated by contradictory information.[7] Other extensions of the AGM-paradigm explore the formulation of the standard postulates using modal logics ([30], [26]) or extending normal modal operators by non-normal information operators to perform iterated belief revision ([4]). In both cases, connections can be drawn to the present work, where updates are interpreted in the framework of a modal logic with a non-standard consequence relation.[8]

The aim of the present work is to combine the temporally-based retraction on inconsistent updates, with the Principle of Information Economy: the former allows to preserve the Principles of Indifference (for simultaneously received data) and of Preference (for older data); the latter is satisfied by retracting data back to the first stage at which consistency is restored. This procedure refers to the *retraction* operator ($K \ominus A$) introduced in [18] which undoes the effect of a previous operation on a belief set. The operation defined in **AIUR** corresponds intuitively to a retraction on expansion, hence to a specific case of recovery of revision. A standard analysis of an update-retraction procedure is easily shown by the following diagram:

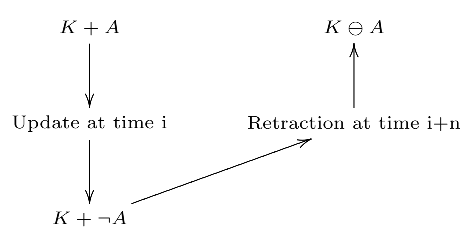

where the content of the retraction operation is based on a preference for the newer information. In [7] this form of update is implemented by the combination of two principles:

- *Primacy of new information*: the revised knowledge of the system should conform to the new information, which implies a complete relaince on its truth;

[7]The *retraction* operator from [5] does not satisfy Inclusion nor Recovery, but it satisfy Failure (if $\vDash A$, then $K \subseteq K \div A$).

[8]Another relevant paradigm for belief change is the Dynamic Epistemic Logic approach, [29]. In DEL an information update produces the elimination of irrelevant or redundant information; a Public Announcement $\phi!$ eliminates all the worlds which fail to satisfy ϕ; Knowledge after a Public Announcement $[\phi!K_jA]$ establishes that after the elimination in the current model of all the worlds that fail to satisfy ϕ, agent j knows that A; according to an Unsuccesful Update $\langle A \wedge \neg K_jA\rangle\neg(A \wedge \neg K_jA)$, the agent asserts that A and that agent j does not know that A, which becomes false because of the announcement. The extension to the temporal ordering is given by the framework of Epistemic Temporal Logic ([28]) and the introduction of past operators ([17]). Within this framework, neither a proper form of retraction or withdrawal on previously performed announcement, nor the mentioned combination of possible updates are considered.

- *Persistence of prior knowledge*: as much old knowledge as possible should be retained in the revised knowledge; the resulting state is obtained by some minimal change.

AIUR, which focuses on primacy of older data in view of new inconsistent information, models a specific case of the following further principle from [7]:

- *Fairness*: if there are many candidates for the revised knowledge that satisfy the above principles, then one of them should not be arbitrarily chosen.

It is the dynamics of selection on *possible updates*, whose combination is inconsistent in view of the starting belief set:

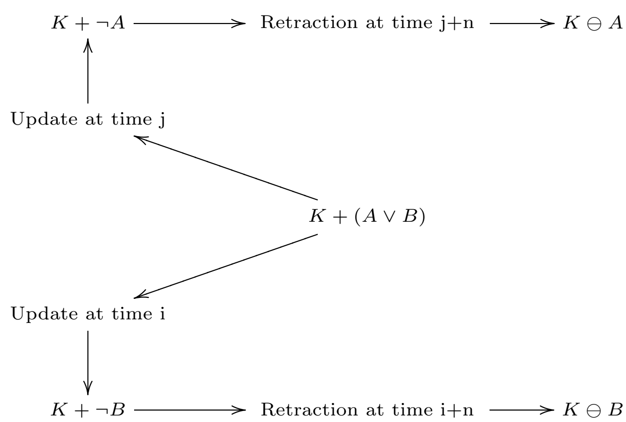

The temporally ordered update-retraction procedure induces a preorder on the models of the differently updated belief sets, revealing a relation of entrenchment: the older the update (that has not been yet retracted) the most entrenched becomes the related belief. A dynamic version of the principle of persistence of prior knowledge comes from recent literature in AI (see e.g. [33]):

DEFINITION 4 (Minimal Change Principle). During a state transition, the change between states should be as little as possible.

This principle becomes inappropriate when state change involves disjunctive descriptions of the resulting possible states. The present paper provides a resolution method for inconsistent outcomes involving the minimum amount of retracted information on disjunctive solutions.[9]

AIUR has the standard format of an Adaptive Logic (see [2, 3]):

DEFINITION 5. An Adaptive Logic (**AL**) in the standard format is defined as a triple $\mathbf{AL} = \langle LLL, \Omega, AS \rangle$, according to the following description:

[9] For an adaptive procedure that respect literally the principle from [7] on the primacy of newer information, see [22] and the comments in the conclusive section of this paper.

1. *LLL* is the lower limit logic, a monotonic, compact logic; for **AIUR** it is called $\mathbf{T}^+$ and consists of a set of superposed logics, each of them defined in terms of at least one temporally indexed abnormal formula; it corresponds to a multi-doxastic version of the modal logic **T**;

2. Ω is a set of abnormal formulas; for **AIUR** the abnormalities express conflicts between some $A \in K$ and some update with information $\neg A$ obtained at later time;

3. *AS* is the Adaptive Strategy; in this case *Minimal Abnormality Strategy*, which selects valid formulas at each stage, hence establishing retractions of updates on a temporal basis.

I shall proceed as follows. After some preliminaries in Section 2, an intuitive characterization of the adaptive interpretation of Information Updates and their Retractions will be provided in Section 3. In Section 4, the core of the logic **AIUR** is introduced, with the semantic and proof-theoretical characterizations of the Lower Limit Logic $\mathbf{T}^+$ and the set of abnormalities Ω. In Section 5 and 6, respectively the proof theory and the semantics of **AIUR** as well as some examples are formulated. The final Section gives some remarks on the distinction with the standard retraction operator and possible extensions of this research.

2 Preliminaries

Let $\mathcal{L}$ be the standard language of classical propositional logic (henceforth **CL**) that is formed from a finite set of atoms in the usual way. A **CL**-model is a function from the set of atoms to $\{0, 1\}$. I shall use $M, M_1, \ldots$ as metavariables for **CL**-models, $\mathcal{M}$ to denote the set of all **CL**-models. A model M is a model of a belief set K iff all the members of K are true in it; I shall use $\mathbf{M}$ as the subset of members of $\mathcal{M}$ in which all members of K are true. As usual, $M \vDash A$ will denote that M verifies A.

In the following, I define an operation of update in terms of an accessibility relation from the actual to some possible epistemic states, according to a temporal order. This relation preserves a defeasible interpretation of belief contents: the information held true by the agent in her actual epistemic state can still be rejected according to some other (later) state. It is in view of this property that we allow retractions on given updates. The frames of the possible worlds semantics of $\mathbf{T}^+$ are reflexive: according to this property each of the states is internally consistent, inconsistency arising by extensions to later states. The final belief set is obtained when temporal updates and related retractions have been completed.

The language $\mathcal{L}^I$ of $\mathbf{T}^+$ is the language $\mathcal{L}$ of **CL** extended with an appropriate set of indexed modalities for the operations of information update. The modal operator I behaves as a possibility operator. A belief set is intended as a finite set of sentences of $\mathcal{L}^I$, closed under logical consequence, i.e. such a set K corresponds to $Cn(K)$. K refers now to the (updated) belief set obtained by translating a finite set of sentences of $\mathcal{L}$ into the very same set of sentences in which each formula is prefixed by an occurence of the modal operator I_i. The index i is defined by membership to the set $\mathcal{T} = \{0, 1, \ldots, \}$ of temporal indices. An expression of the form I_iA holding in the belief set K has the meaning:

> I_iA: "the agent's belief set K is updated at time i with the information that A".

In view of the adaptive dynamic approach, at the end of the update-process the belief set will contain the result of finally accepted updates: to this aim the operator $\oplus I$ will be used, which behaves as a necessitation operator (validity in all states). By the formulation of indexed modalities, $\mathcal{L}^I$ is the language of a multi-modal version of **T**. In view of the meaning presently attached to such modalities, I will discard any expression formulating possible iterations of the form:

> $I_iI_{i-n}A$: "the agent's belief set K is updated at time i with the information that she has received at time $i-n$ the information that A".

In other words, the modal language will be restricted to first degree modalities and only modal formulas in which no nested belief operators occur will henceforth be considered as well-formed. This simplifies the language and it is enough for the present aim. I moreover restrict myself to a one-agent based language, and leave the possible extension to a multi-agent formulation to another occasion.

The set of atoms in $\mathcal{L}^I$ shall be denoted by $\mathcal{W}^P$ and $\mathcal{W}^{\pm}$ will be used to refer to the set of literals (an atom or negation of an atom). Simply $\mathcal{W}$ is used to refer to all well-formed formulas of $\mathcal{L}^I$, composed by $\mathcal{W}^{\pm}$ in terms of the standard propositional connectives plus the modal operator: $\neg, \wedge, \vee, \supset, I_i$. The symbols $\vDash$ and $\vdash$ have the usual meanings of consequence and derivability relations (with abbreviations for the related logical languages attached to them to distinguish among **AL** and the **LLL**). The abbreviation $\bigvee \Delta$ will stand for the disjunction of the members of Δ, where Δ is a set of formulas.

3 Intuitive Characterization of Updates and Retractions

The notion of update appeared in the literature of theory change in relation to the operation of belief revision defined by the AGM paradigm. The distinction between revision and update was introduced in [20] and it was later formalized in [19]. By revision one formalizes changes due to new information in a static world; update refers instead to modifications that a knowledge base undergoes when the world of reference changes. An obvious extension of the notion of update is the addition of inconsistent information.

Consider the following example. An agent is informed that an important event she wants to attend to will be organised in one of the two theaters of her city, the Blue or the Red Theater. She receives later the information that the event will not be organised in the Blue Theater, so she simply infers that the event will take place at the Red Theater. At this stage a standard dynamic of information update and revision takes place: an update incoming from the external world forces the agent to derive some new content. Assume that, by a yet later announcement, she becomes informed (by a seemingly equally reliable source) that the event will not be organised in the Red Theater.[10] In this case, the doxastic dynamics is more complex in view of the aimed consistency: it does not concern the deterministic change in a belief state due to a certain AGM-style change; nor it involves some non-deterministic but immediate change of belief. It rather requires the description of the agent's internal dynamics: she deals with a set of incoming data and is required to perform a rational choice among possibly contradictory (but equally reliable) informations.

The formalisation of the example in $\mathcal{L}^I$ is obtained by the following premise set:

(1) $K = \{I_1(p \vee q), I_2\neg p, I_3\neg q\}$.

This set expresses a (initially empty) belief set updated at the initial stage 1 with the information that $(p \vee q)$: "the event will take place either at the Blue or at the Red Theater"; at the next stage 2, the set K is updated with the information that $\neg p$: "it is not the Blue Theater the place where the event will take place"; and at the final stage 3 the agent becomes informed that $\neg q$: "it is not the Red Theater the place where the event will take place". To formulate a consistent consequence set on the basis of these informations means to establish which update is the most rational to retract. Let start by giving some basic definitions.

DEFINITION 6 (Non-monotonic Update). An update at time j with information A (I_jA) is non-monotonic with respect to K if $\neg A$ is $\mathbf{T}^+$-derivable from K by a previous update, i.e. if $K \vDash_{\mathbf{T}^+} I_i\neg A (i < j \in \mathcal{T})$. The situation obtained after the non-monotonic update at time j is denoted by a formula of the form $I_jA \wedge \neg A$.

DEFINITION 7 (Combined Non-monotonic Update). A rational agent is faced with a combined non-monotonic update if for a belief set K it holds that $K \nvDash_{\mathbf{T}^+} (I_iA \wedge \neg A), K \nvDash_{\mathbf{T}^+} (I_jB \wedge \neg B)(i, j \in \mathcal{T})$, but it holds that $K \vDash_{\mathbf{T}^+} (\bigvee \Delta)$ where formulas $(I_iA \wedge \neg A), (I_jB \wedge \neg B) \subseteq (\bigvee \Delta)$.

[10] The assumption on the equal reliability of the information sources is functional to performing the selection procedure uniquely on a temporal order. Different priority relations can be designed varying the relevance of sources.

In the logic $\mathbf{T}^+$, neither the formula $I_2\neg p \wedge p$ nor $I_3\neg q \wedge q$ are derivable from K, but their disjunction is. This means that the agent has to select either the validity of $\neg p$ provided that it is an update on p (and thus on $(p \vee q)$), *or* the validity of $\neg q$ provided it is an update on q (and thus on $(p \vee q)$). The retraction of one of these updates is necessary in those cases – as in our initial example – where the agent considers reliable some initial information (viz. that the event will take place, either in the Blue or in the Red Theater). On this basis, some later information is wrong or inaccurate.

Semantically, all three premises are verified in each of their models, and all models verify the disjunction $(I_2\neg p \wedge p) \vee (I_3\neg q \wedge q)$. In view of some crucial notions introduced in a later section, it will appear that **AIUR** selects the information that – if retracted – first makes it possible to come back to a consistent belief state. In the example this is obtained by selecting those models that verify $(I_3\neg q \wedge q)$ and falsify $(I_2\neg p \wedge p)$. In these models, also p and $\neg q$ are falsified, which means the updates $I_1 p$ and $I_3\neg q$ are retracted, whereas the formula $(q \wedge \neg p)$ is verified.[11]

The dynamic proof-theory for the retraction procedure is also based on the validity of updates depending on the order of time. This is translated by the typical procedure of derivation in the adaptive proof-theory, where contents are derivable on conditions. The disjunction $\bigvee(\Delta)$ of combined non-monotonic updates is $\mathbf{T}^+$-derivable. A derivation performed by the adaptive logic **AIUR** assumes that one disjunct $I_3\neg q \wedge q$ holds on condition that the older update $I_2\neg p \wedge p$ be false. This implies that any content assuming $I_3\neg q \wedge q$ being false is no longer derivable, so that the update with $\neg q$ is retracted.

4 A Logic for Non-monotonic Information Updates

The present section provides the semantic and syntactic characterizations of the Lower Limit Logic $\mathbf{T}^+$ and the definition of the derivable abnormal formulas that describe the non-monotonic updates.

As for any Adaptive Logic in standard format, $\mathbf{T}^+$ is a Tarski Logic (monotonic and compact). In this case, one is dealing with a prioritized multi-doxastic logic, whose consequence set for a premise set K is built from models verifying the various indexed abnormal formulas for each degree. The adaptive logic **AIUR** selects on the $\mathbf{T}^+$-consequence set according to the procedure defined by the Minimal Abnormality Strategy, in order to establish which disjunct(s) of combined non-monotonic updates hold(s): in turn this will tell which updates persist, and which retraction is needed.

[11] The result is thus an adaptive selection applied on a prioritized set of models. This recalls the models of *liberation* defined in [5]. For the model of σ-liberation, both the agent's set of beliefs and the way to remove beliefs are formed on the basis of the information received in "the course of its intellectual career" - the priority of the most recent information allows to extract consistent subsets out of the possibly inconsistent set of beliefs; for the linear liberation, different candidate belief sets are ordered, with the current belief set being the most preferred ones in the ordering.

$\mathbf{T}^+$ is characterized by a standard possible-world semantics. A $\mathbf{T}^+$-model is a quadruple $M = \langle W, w_0, \mathcal{R}, v\rangle$, where W is a set of possible worlds corresponding to epistemic states, in which formulas from the language $\mathcal{L}^I$ are valuated; w_0 is the actual state of knowledge of the agent; $\mathcal{R}$ is a set of accessibility relations $R_i : w_0 \rightarrow W (i \in \mathcal{T})$ from the actual to the set of possible states; $v : \mathcal{W}^P \times W \rightarrow \{1, 0\}$ is the valuation function. Each possible state is time-indexed: the accessibility relation from one state to another simulates the update; reflexive relations allow for any such state only an atom or its negation to hold (viz. no inconsistent update is allowed in *one* world), whereas accessing various consecutive states the agent may face inconsistent beliefs.

The valuation of formulas in a model M is standardly characterized for logical connectives and with the update operator I_i defined as a possibility operator:

C1 where $A \in \mathcal{W}^P$, $v_M(A, w) = v(A, w)$
C2 $v_M(\neg A, w) = 1$ iff $v_M(A, w) = 0$
C3 $v_M(A \vee B, w) = 1$ iff $v_M(A, w) = 1$ or $v_M(B, w) = 1$
C4 $v_M(A \wedge B, w) = 1$ iff $v_M(A, w) = 1$ and $v_M(B, w) = 1$
C5 $v_M(A \supset B, w) = 1$ iff $v_M(A, w) = 0$ or $v_M(B, w) = 1$
C6 $v_M(I_i A, w) = 1$ iff $v_M(A, w') = 1$ for some w' such that $R_i ww'$

The standard semantic notions are defined as usual. A model M verifies A iff $v_M(A, w_0) = 1$; A is valid in $\mathbf{T}^+$ ($\vDash_{\mathbf{T}^+} A$) if it is verified by all its models; A is a consequence of a set of premises in $\mathbf{T}^+$ ($K \vDash_{\mathbf{T}^+} A$) if A is true in every model of K.

Proof-theoretically, the weakest characterization of $\mathbf{T}^+$ is given by the axioms of the modal logic $\mathbf{T}$. In the following, the description of the logic for non-monotonic updates will make use of the possibility fragment of $\mathbf{T}$, represented by the update operator I_i. In the syntactic formulation, the notions of Final Information Update and Final Information Update Retraction will be defined: these express the adaptive counterparts of update and retraction after any possible dynamics in the derivation has been performed. The abbreviation $\oplus I$ shall be used for the contents updated and no longer retractable after time n, where for all $I_i A \in K, i \in \mathcal{T} = \{1, \ldots, n\}$. In this way, the neccessitation fragment of $\mathbf{T}$ can be reconstructed so that $\mathbf{T}^+$ contains the axioms of **CL** plus:

Distribution $\oplus I(\varphi \rightarrow \psi)) \rightarrow (\oplus I\varphi \rightarrow \oplus I\psi)$;
Reflexivity $\oplus I\varphi \rightarrow \varphi$;
Modus Ponens $\varphi \rightarrow \psi; \varphi$, then ψ;
Necessitation $\varphi \rightarrow \oplus I\varphi$.

Non-monotonic updates derivable from a premise set K represent the second element in the definition of the adaptive logic **AIUR**. For each index $i \in \mathcal{T}$, a formula of the form

DEFINITION 8 (Set of Abnormalities). $\Omega_i = \{I_i A \wedge \neg A \mid A \in \mathcal{W}^{\pm}\}$

is an abnormality. Disjunctions of abnormalities formally define expressions for combined non-monotonic updates:

DEFINITION 9 (Dab-Formula). $Dab(\Delta)$ stands for $\bigvee(\Delta)$ where $\Delta \subseteq \Omega$.

If Δ is a singleton, $Dab(\Delta)$ is simply an abnormality ($A \vee Dab(\emptyset)$), i.e. a member of Ω; if Δ is empty, $Dab(\Delta)$ is empty as well. In view of the indexed set of abnormalities, by the following definition one determines the position in time of combined non-monotonic updates:

DEFINITION 10 (Dab-formula at degree). A Dab-formula $Dab(\Delta)$ is said to be of degree i ($Dab^i(\Delta)$) iff $\Delta \subseteq \Omega_i$, and for any i' such that $Dab(\Delta')$ and $\Delta' \subseteq \Omega_{i'}$, then $i' < i \in \mathcal{T}$.

According to definitions 8 and 10, the temporal indices of the update operators are used to order derivable disjunctions of abnormalities by increasing degrees. The set of abnormalities at degree i is the union of sets $\Omega_1 \cup \Omega_2 \cup \ldots \cup \Omega_i$:

DEFINITION 11 (The Set of all indexed Abnormalities).

$$\Omega_n = \bigcup_{i=1}^{n} \Omega_i$$

In the case of our example $K = \{I_1(p \vee q), I_2\neg p \text{ and } I_3\neg q\}$, the set of abnormalities at degree 3 is given by the union of all disjuncts of abnormal formulas at previous degrees:

(2) $\Omega_3 = \{\{I_3\neg q \wedge q\} \cup \{I_2\neg p \wedge p\} \cup \{(I_1 p \wedge \neg p), (I_1 q \wedge \neg q)\}\}$.

In the following sections I shall first consider the proof-theoretical and next the semantic selection performed according to the third element of the adaptive logic **AIUR**, namely the Minimal Abnormality Strategy. The adaptive strategy defines a way to establish which disjuncts of $\mathbf{T}^+$-derivable Dab-formulas can be considered valid: this in turn means to accept some update and to allow some retraction.

5 Syntactic Selection of Retractions

In this section, the adaptive strategy of the logic **AIUR** is given in its proof-theoretical formulation. The Minimal Abnormality Strategy is the final element in the definition of the adaptive logic for update retractions.

The standard structure of a line in an adaptive derivation contains the following elements: (i) a line number; (ii) the derived formula; (iii) the line numbers of the formulas from which the element in (ii) is derived; (iv) the name of the rule(s) applied to derive the formula from previous lines; (v) the condition on which the second element is derived. In a **AIUR**-derivation the second element of a line is derived provided the elements of the condition are assumed to be false on the premise set. The adaptive rules for the derivation of a new line from previous ones are:

PREM at any stage of a proof, for any $A \in K$, one may add to the proof a line consisting of:
(i) an appropriate line number;
(ii) A;
(iii) a dash;
(iv) PREM;
(v) $\emptyset$;
by the *premise rule* premises are introduced in a line on the empty condition;

RU at any stage of a proof, for any $B \in \mathcal{W}^P$, if $A_1, \ldots A_n \vdash_{\mathbf{T}^+} B$, and $\Delta_1, \ldots, \Delta_n$ are the conditions for $A_1, \ldots A_n$, a line may be added consisting of:
(i) an appropriate line number;
(ii) B;
(iii) the line numbers of the $A_1, \ldots A_n$;
(iv) RU;
(v) $\Delta_1 \cup \ldots \cup \Delta_n$;
by the *unconditional rule* a $\mathbf{T}^+$-derivable formula can be added to the proof, without any new condition but (if any) the conditions of the formulas to which the rule is applied;

RC at any stage of a proof, for any $B \in \mathcal{W}^P$, if $A_1, \ldots A_n \vdash_{\mathbf{T}^+} B \vee Dab(\Delta)$, and $\Delta_1, \ldots, \Delta_n$ are the conditions for $A_1, \ldots A_n$, a line may be added consisting of:
(i) an appropriate line number;
(ii) B;
(iii) the line numbers of the $A_1, \ldots A_n$;
(iv) RC;
(v) $\Delta_1 \cup \ldots \cup \Delta_n \cup \Delta$;
by the *conditional rule* the derivability by $\mathbf{T}^+$ of a formula of the form $B \vee Dab(\Delta)$ allows the derivation of B on the assumption that all members of $Dab(\Delta)$ are false; in the new line they will be introduced as a new condition.

The derivation of a *Dab*-formula of a given degree is restricted by an extra condition. Provided that a disjunction of abnormalities expresses the possible non-monotonic updates the agent is faced with, the extra-condition establishes that one of the update is accepted assuming that updates indexed at earlier times are false. I shall call this the *Regularity Condition*: its intuitive meaning is that the analysis of a belief set which has become inconsistent requires each new update to make older ones false. The Regularity Condition for *Dab*-formulas is defined as follows:

DEFINITION 12 (Regular **AIUR** *Dab*-formulas). Given a **AIUR** proof, $Dab^i(\Delta)$ is a regular *Dab*-formula of degree i iff

(i) $Dab^i(\Delta)$ is derived on condition Θ;

(ii) $\Delta \subseteq \Omega_i$ and

(iii) $\Theta \subseteq \Omega_1, \ldots, \Omega_{i-1}$.

If there is no other Δ for which $Dab^i(\Delta)$ is regular, then this *Dab*-formula is a called a *Minimal* **AIUR** *Dab-formula* of degree i:

DEFINITION 13 (Minimal **AIUR** *Dab*-formula). Given a **AIUR** proof, $Dab^i(\Delta)$ is a minimal *Dab*-formula of degree i iff $Dab^i(\Delta)$ is regular according to Definition 12, and there is no $\Delta' \subset \Delta$ for which $Dab^i(\Delta')$ is regular.

According to the Regularity Condition (which is a restriction on the Conditional Rule), the derivability of a minimal information update is accepted on condition of the falsity of previously performed updates. Informally, the minimal *Dab*-formulas for an **AIUR** premise-set explicitate all the updates that the agent has undergone at each consecutive temporal stage from her initially empty belief set. The adaptive selection can now be performed on these updates to show which one needs to be retracted.

The selection on minimal *Dab*-formulas has effect on the adaptive notion of derivability. Derivability in an adaptive logic is dynamic, which means that a formula derived at one stage of the proof can be later withdrawn. This implies first the definition of derivation at stage:

DEFINITION 14 (Derivability at Stage). A content A is derived from K at stage s of an **AIUR**-proof iff A is the second element of a line l which is not marked at stage s.

A *Dab*-formula of degree i can be derived at stage s of a proof iff it is regular. Accordingly, one defines the notion of minimal *Dab*-formula at stage :

DEFINITION 15 (Minimal *Dab*-formula at stage). Given a **AIUR** proof, $Dab^i_s(\Delta)$ is a minimal *Dab*-formula of degree i at stage s of the proof iff at that stage $Dab^i(\Delta)$ is regular and not marked, and there is no $\Delta' \subset \Delta$ for which $Dab^i(\Delta')$ is regular.

The marking procedure will now define the withdrawing of formulas previously derived at some stage of a derivation. A minimal *Dab*-formula derived at a stage s causes the marking of any previous line in which its conditions (at stage s assumed to be false) were derived. This is obtained by the following procedure. Let a choice set of $\Sigma = \{\Delta_1, \Delta_2, ...\}$ be a set that contains an element out of each member of Σ, and let a minimal choice set of Σ be a choice set of Σ of which no proper subset is a choice set of Σ. Consider now the minimal choice set $\Phi^i_s(K)$ of the minimal *Dab*-formulas at degree i of a premise set K at stage s of an **AIUR**-proof. Then a marked line is one whose content is retracted and considered (at that stage) no longer derived:

DEFINITION 16 (Marking for Minimal Abnormality). A line l of a **AIUR**-proof where a formula A is derived on condition Δ is marked at stage s iff

(i) there is no $\phi \in \Phi^i_s$ such that $\phi \cap \Delta = \emptyset$, or

(ii) for some $\phi \in \Phi^i_s$ there is no line at which A is derived on condition Θ and at that line $\phi \cap \Theta = \emptyset$.

The procedure goes on - if needed - stepwise at any next stage by looking at the minimal choice set of lower degree Φ^{i-1}_{s+i}, next at Φ^{i-2}_{s+j} and so on.

The dynamic notion of derivability at stage has its counterpart at the end of the marking procedure in the notion of final derivability: a line which is (and stays) *unmarked* is considered finally derived. This becomes for **AIUR** a specific definition of valid information update $\oplus I$:

DEFINITION 17 (Final Information Update $\oplus I$). An information update $I_i A$ for K is valid according to a **AIUR**-proof if A is finally derived from K in a **AIUR**-proof, i.e. iff

(i) A is the second element of a line l of the proof;

(ii) line l is not marked at that stage s and any extension of the proof in which line l is marked can be further extended such that l is unmarked, i.e. for any later stage $s+i$ at which A is derived on condition Θ, for any $\phi \in \Phi^{i-n}_{s+i}$, it holds that $\phi \cap \Theta = \emptyset$.

For lines that stays unmarked, there will be lines whose content shall remain *marked*. With respect to these contents, the dynamics of derivability of information updates allows the following definition of final update retraction $\ominus I$:

DEFINITION 18 (Final Information Update Retraction $\ominus I$). A non-monotonic information update $I_j B$ of K is retracted according to a **AIUR**-proof if B is marked at some stage of the derivation, i.e. iff

(i) B is the second element of a line l of the proof on the condition that an update $I_i A \wedge \neg A$ is false;

(ii) at some next stage of the proof, the update $I_i A \wedge \neg A$ is derived as a regular one;

(iii) at any next stage of the very same derivation $I_i A \wedge \neg A$ stays unmarked, which makes line l finally marked.

5.1 Examples

Let us consider a derivation from the mentioned premise set $K = \{I_1(p \vee q), I_2\neg p, I_3\neg q\}$:

1	$I_1(p \vee q)$	PREM	$\emptyset$	
2	$I_2\neg p$	PREM	$\emptyset$	
3	$I_3\neg q$	PREM	$\emptyset$	
4	$p \vee q$	1; RC	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q\}$	
5	$\neg p$	2; RC	$\{I_2\neg p \wedge p\}$	
6	q	4, 5; RU	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q, I_2\neg p \wedge p\}$	
7	$\neg q$	3; RC	$\{I_3\neg q \wedge q\}$	$\checkmark^{10}$
8	p	4, 7; RU	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q, I_3\neg q \wedge q\}$	$\checkmark^{10}$
9	$(I_2\neg p \wedge p) \vee (I_3\neg q \wedge q)$	1, 2, 3; RC	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q\}$	
10	$I_3\neg q \wedge q$	9; RC	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q, I_2\neg p \wedge p\}$	

Lines 7 and 8 are marked in view of line 10 so that their contents $\neg q$ and p are considered no longer derived, and in particular the update $I_3\neg q$ is retracted. On the other hand, lines 5 and 6 stay unmarked, which means that in view of this **AIUR**-proof, content q holds and the update $I_2\neg p$ is accepted. In the following line

11	$(I_1 p \wedge \neg p) \vee (I_1 q \wedge \neg q)$	1, 2, 3; RC	$\{I_2\neg p \wedge p, I_3\neg q \wedge q\}$

the Dab^1_{11}-formula is not minimal according to Definition 15 because its condition is of a higher degree and thus it is not regular. Hence, none of its disjuncts will be selected in a minimal choice set Φ^1_{11} (in fact, there is no such minimal choice set of degree 1, because no *Dab*-formula of degree 1 can be minimal), and therefore no further marking is possible. This makes the derivability of the content at line 10 and the related retractions final.

Consider a different example, namely the premise set $K = \{I_1\neg p, I_2(p \vee q), I_3\neg q\}$; the following is a valid **AIUR**-derivation:

1	$I_1\neg p$	PREM	$\emptyset$	
2	$I_2(p \vee q)$	PREM	$\emptyset$	
3	$I_3\neg q$	PREM	$\emptyset$	
4	$\neg p$	1; RC	$\{I_1\neg p \wedge p\}$	
5	$\neg q$	3; RC	$\{I_2 p \wedge \neg p, I_3\neg q \wedge q\}$	$\checkmark^{10}$
6	$p \vee q$	2; RC	$\{I_2 p \wedge \neg p, I_3\neg q \wedge q\}$	$\checkmark^{10}$
7	p	5; 6; RU	$\{I_2 p \wedge \neg p, I_3\neg q \wedge q\}$	$\checkmark^{10}$
8	q	4, 6; RC	$\{I_1\neg p \wedge p, I_2 p \wedge \neg p, I_2 q \wedge \neg q\}$	
9	$(I_2 p \wedge \neg p) \vee (I_3\neg q \wedge q)$	1, 2, 3; RC	$\{I_1\neg p \wedge p\}$	
10	$I_3\neg q \wedge q$	1, 2, 3; 9 RC	$\{I_1\neg p \wedge p, I_2 p \wedge \neg p\}$	

According to this derivation the contents at lines 5, 6 and 7 are marked: this means that the update $I_3\neg q$ is retracted, because the content q stays unmarked at line 8. On the other hand, $\neg p$ and obviously $p \vee q$ are derivable, so that the updates with these contents are preserved. Hence, in this case, the quickest way to restore consistency is to retract the update at the last stage and to preserve the initial ones. As for the previous example, also in this derivation the following *Dab*-formula:

11	$(I_1\neg p \wedge p) \vee (I_2 q \wedge \neg q)$	1, 2, 3; RC	$\{I_2 p \wedge \neg p, I_3\neg q \wedge q\}$

is not regular and therefore does not allow for any further marking.

Consider now the following derivation from the premise set $K = \{I_1(p \vee q), I_2 \sim p, I_3 \sim q, I_4 p\}$:

1	$I_1(p \vee q)$	PREM	$\emptyset$	
2	$I_2 \neg p$	PREM	$\emptyset$	
3	$I_3 \neg q$	PREM	$\emptyset$	
4	$I_4 p$	PREM	$\emptyset$	
5	$p \vee q$	1; RC	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q\}$	
6	$\neg p$	2; RC	$\{I_2 \neg p \wedge p\}$	
7	q	5, 6; RU	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q, I_2 \neg p \wedge p\}$	
8	$\neg q$	3; RC	$\{I_3 \neg q \wedge q\}$	$\surd^{13}$
9	p	5, 8; RU	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q, I_3 \neg q \wedge q\}$	$\surd^{13}$
10	$(I_4 p \wedge \neg p) \vee (I_3 \neg q \wedge q)$	1, 2, 3, 4; RC	$\{I_2 \neg p \wedge p, I_1 \neg q \wedge q\}$	
11	$I_4 p \wedge \neg p$	10; RC	$\{I_3 \neg q \wedge q, I_2 \neg p \wedge p, I_1 \neg q \wedge q\}$	$\surd^{13}$
12	$(I_3 \neg q \wedge q) \vee (I_2 \neg p \wedge p)$	1, 2, 3; RC	$\{I_1 p \wedge \neg p, I_1 q \wedge \neg q\}$	
13	$I_3 \neg q \wedge q$	12; RC	$\{I_2 \neg p \wedge p, I_1 p \wedge \neg p, I_1 q \wedge \neg q\}$	

By this derivation the selection given by the minimal choice set Φ^4_{11} does not have any marking effect, which means no retraction is performed and no consistency is reached. The next selection by Φ^3_{13} leads to the marking of lines 8, 9 and 11, so that the updates $I_4 p$ and $I_3 \neg q$ are retracted, whereas the update $I_2 \neg p$ is valid and q stays derivable in view of update at time 1. This example shows that it is not always the latest update that must be retracted in order to restore consistency.

6 Semantic Selection of Retractions

In the present section the semantic selection of updates and retractions for the adaptive logic **AIUR** is introduced, consisting in a selection procedure on the indexed sets of models of the logic $\mathbf{T}^+$. According to the Minimal Abnormality Strategy, those $\mathbf{T}^+$-models are selected which are not more abnormal than what is required by the premises. The degree of abnormality is established again in terms of the temporal order of the abnormal formulas, directly imported in the models that verify them.

The consequence set of a premise set K in the combined adaptive logic **AIUR** is obtained by the super-position of increasing **LLL**-consequence sets (where i is the maximal number for which I_i occurs in the premise set):

$$Cn_{AIUR}(K) = Cn_{\mathbf{T^+}_i}(Cn_{\mathbf{T^+}_{i-1}}(\ldots(Cn_{\mathbf{T^+}_1}))).$$

To each index i it corresponds a set of abnormalities. The semantics describes a selection on the related models, in which at least one abnormal formula is validated. The selection of valid abnormal formulas determines the valid retractions.

The selection is based on the indexing procedure for *Dab*-formulas. On this basis, a semantic definition for establishing the minimal abnormal degree of a model with respect to the premise set is formulated. Let us start by the definition of a *Dab*-consequence:

DEFINITION 19 (*Dab*-Consequence). $Dab^i(\Delta)$ is a *Dab*-consequence of a premise set K at degree i iff $K \vDash_{\mathbf{T}^+} Dab(\Delta)$, and $\Delta \subseteq \Omega_i$.

If $Dab(\Delta)$ is a *Dab*-consequence of a set K, then so is any $Dab(\Delta')$ such that $\Delta' \subset \Delta$. This is why a further definition is needed:

DEFINITION 20 (Minimal *Dab*-Consequence). $Dab^i(\Delta)$ is a *minimal* *Dab*-consequence of K at degree i iff $K \vDash_{\mathbf{T}^+} Dab(\Delta)$, $\Delta \subseteq \Omega_i$, and there is no $\Delta' \subset \Delta$ such that $K \vDash_{\mathbf{T}^+} Dab(\Delta')$ at that degree.

Where $Dab^1(\Delta), \ldots, Dab^i(\Delta)$, are the minimal *Dab*-consequences of a set of premises K at the various degrees, any $\mathbf{T}^+$-model of K will verify at least one member out of each disjunction of abnormalities. For any $\mathbf{T}^+$-model, its abnormal part is defined as follows:

DEFINITION 21 (Abnormal model). Provided M is a $\mathbf{T}^+$-model, $Ab^i(M) = \{A \in \Omega_i \mid M \vDash A\}$.

The set $\mathbf{M}$ of all $\mathbf{T}^+$-models of K at degree i contains abnormalities at the various degrees up to i: $\Delta \subseteq \Omega_1, \ldots, \Delta \subseteq \Omega_i$. At each degree a set of minimal *Dab*-consequences is verified. In order to perform a selection that resolves inconsistencies as soon as possible, abnormalities of degree i are in general preferred to those of degree $i-1$. The selection starts therefore by taking into account all the abnormal models at degree i; then it considers those at degree $i-1$, and so on.

DEFINITION 22 (Selected Models). For each degree i, the minimal Dab^i-formulas verified at that degree induce a selection on the set $\mathbf{M}$ of all models of K, performed stepwise as follows:

$$sel_0(\mathbf{M}) = \{M \mid M \vDash K\};$$
$$sel_{n+1}(\mathbf{M}) = \{M \in sel_n(\mathbf{M}) \mid \text{ for no } M' \in sel_n(\mathbf{M}) \text{ such that } Ab^{i-(n+1)}(M') \subset Ab^{i-(n+1)}(M)\}.$$

At step zero, one considers the entire set $\mathbf{M}$ of models of K and restricts this set stepwise going to each abnormal model of a lower degree. The selection stops when the selected models verify some abnormal formula but consistency is restored on the premise set. The selected models are called *minimally abnormal*:

DEFINITION 23 (Minimally Abnormal Model). A $\mathbf{T}^+$-model M of K is minimally abnormal iff it is included in all possible selection steps: $M \in sel_0(\mathbf{M}) \cap sel_1(\mathbf{M}) \cap \ldots$.

The minimally abnormal models establish the **AIUR**-consequence set of the premises:

DEFINITION 24 (**AIUR**-Consequence). $K \vDash_{\mathbf{AIUR}} A$ iff A is verified by all minimally abnormal models of K.

Consider, as an example, the premise set $K = \{I_1(p \vee q), I_2 \sim p, I_3 \sim q, I_4 p\}$. The consequence set $\mathbf{M}(K)$ contains four types of abnormal models:

$$Ab^4(M) = \{(I_4p \wedge \neg p), (I_3\neg q \wedge q), (I_2\neg p \wedge p), (I_1p \wedge \neg p), (I_1q \wedge \neg q)\};$$
$$Ab^3(M) = \{(I_3\neg q \wedge q), (I_2\neg p \wedge p), (I_1q \wedge \neg q)\};$$
$$Ab^2(M) = \{(I_2\neg p \wedge p), (I_1p \wedge \neg p)\};$$
$$Ab^1(M) = \{\emptyset\}.$$

The selection proceeds from the highest to the lower degree of abnormality. By the first step one obtains the entire collection of abnormal models $sel_0(\mathbf{M}) = \{\{Ab^4\}, \{Ab^3\}, \{Ab^2\}, \{Ab^1\}\}$, so that $K \vDash_{\mathbf{T}^+_4} (p, \neg q, \neg p, (p \vee q))$. The next selection step $sel_1(\mathbf{M}) = \{\{Ab^3\}, \{Ab^2\}, \{Ab^1\}\}$ will reduce the degree of abnormality by one, but the retraction $\ominus I_4p$ is uneffective in view of the aimed consistency. Hence, it goes on with the next selection step $sel_2(\mathbf{M}) = \{\{Ab^2\}, \{Ab^1\}\}$ at degree 2 retracting $\ominus I_3\neg q$. As a result $K \vDash_{\mathbf{T}^+_2} (\neg p, (p \vee q))$ and the selection stop with $K \vDash_{\mathbf{AIUR}} (\neg p, q)$, where consistency is restored.

7 Conclusion

In the logic **AIUR** a retraction operation for the adaptive logic of information updates has been defined. The given procedure restores the belief set at the time before the retracted update was performed. The adaptive retraction satisfies the Inclusion Postulate in a preferential structure; moreover, when one assumes logical closure of final updates and retractions, the Recovery Postulate is also satisfied. The related postulates and conditions are formulated with due modifications concerning the priority of the core (initial update), which is always the most safe from retractions. The most intuitive application for this selection procedure seems to be the dynamic form of reasoning proper of intelligent agents interacting with external information sources and able to retract previously accepted contents.

The very first advantage of using an adaptive frame for dealing with updates and retractions of inconsistent information is the possibility of generalizing to the predicative case. By the definition of a restricted form of abnormalities, one can moreover formulate retraction operations that are not temporally based, rather contentually restricted: this can lead to a definition of an *elimination* operator that erases from a propositional formula all the knowledge that involves a particular fact, i.e. all previous preconceptions on a set of propositional letters (see [18]).

In [22] an Adaptive Logic for simple Information Update, called **AIU**, is presented. It defines an update operation which simulates simple expansion with inconsistent information and preserves the most recent information when updates turn a belief state into inconsistency. In order to obtain this aim, it defines (both semantically and syntactically) operations of elimination of older information which has become unreliable in view of more recent updates. The analysis of the simple update procedure is apt to model mechanical procedures and it is presented for the specific application to database theory.

Acknowledgements

Research for this paper was supported by subventions from Ghent University. The author wishes to thank Joke Meheus for helpful discussions at early stages of this research, and a referee for comments that led to improvements of its presentation.

BIBLIOGRAPHY

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

[2] D. Batens. A general characterization of adaptive logics. *Logique & Analyse*, 173–175:45–68, 2001.

[3] D. Batens. A universal logic approach to adaptive logics. *Logica universalis*, 1:221–242, 2007.

[4] G. Bonanno. Temporal interaction of information and belief. *Studia Logica*, 86:381–407, 2007.

[5] R. Booth, S. Chopra, A. Ghose, and T. Meyer. Belief liberation (and retraction). *Studia Logica*, Special issue on Reasoning about Action and Change(79):47–72, 2005.

[6] G. Brewka. Belief revision in a framework for default reasoning. In *The Logic of Theory Change Workshop, Konstanz*, volume 465 of *Lecture Notes In Computer Science*. Springer Verlag, 1991.

[7] M. Dalal. Investigations into theory of knowledge base revision. In *Proc. AAAI-88*, pages 475–479. St. Paul, MN, 1988.

[8] K. De Clerq. Maxichoice contraction and revision generalized to include the inconsistent case. *Logic Journal of the IGPL*, 2008. to appear.

[9] E. Ferme' and R.O. Rodriguez. A brief note about rott contraction. *Logic Journal of the IGPL*, 6(6):835–842, 1998.

[10] P. Gärdenfors. *Knowledge in Flux*. Cambridge University Press, 1988.

[11] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In M.Y. Vardi, editor, *Proceedings TARKII*, pages 83–95, 1988.

[12] S.M. Glaster. Recovery recovered. *Journal of Philosophical Logic*, 29:171–206, 2000.

[13] A. Grove. Two modelings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[14] S.O. Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59:845–859, 1994.

[15] S.O. Hansson. Semi-revision. *Journal of Applied Non-Classical Logics*, 7(2):151–175, 1997.

[16] S.O. Hansson and R. Wassermann. Local change. *Studia Logica*, 70:49–76, 2002.

[17] T. Hoshi and A. Yap. ETL, DEL and past operators. In J. van Benthem and E. Pacuit, editors, *Proceedings of the Workshop on Logic and Intelligent Interaction*, ESSLLI08, pages 132–142, 2008.

[18] H. Katsuno and A.O. Mendelzon. A unified view of propositional knowledge base updates. In *IJCAI 1989*, pages 1413–1419, 1989.

[19] H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In Allen, Fikes, and Sandewall, editors, *Principles of Knowledge Representation and Reasoning*, pages 387–394. Morgan Kaufmann, 1991.

[20] A.M. Keller and M. Winslett Wilkins. On the use of an extended relational model to handle changing incomplete information. In *IEEE Transactions on Software Engineering*, volume SE-11, pages 620–633, 1985.

[21] D. Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosohical Logic*, 16:383–394, 1987.

[22] G. Primiero. A model for processing updates with inconsistent information on propositional databases. Submitted for the Proceedings of the IV World Congress on Paraconsistency.

[23] G. Restall and J. Slaney. Realistic belief revision. Technical Report TR-ARP-2-95, Research School of Information Sciences and Engineering and Centre for Information Science Research, Australian National University, Canberra, Australia, 1995.

[24] H. Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1(1):45–78, 1992.
[25] H. Rott and M. Pagnucco. Severe withdrawal (and recovery). *Journal of Philosophical Logic*, 28:501–547, 1999. Corrected complete reprint in issue February 2000.
[26] K. Segerberg. Belief revision from the point of view of doxastic logic. *Bulletin of the Interest Group in Pure and Applied Logics*, pages 535–553, 1995.
[27] A.M. Tamminga. *Belief Dynamics - (Epistemo)logical Investigations*. PhD thesis, Universiteit van Amsterdam, 2001.
[28] J. van Benthem, J. Gerbrandy, and E. Pacuit. Merging frameworks for interaction: DEL and ETL. In *Theoretical Aspects Of Rationality And Knowledge*, Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge, pages 72–81, 2007.
[29] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2006.
[30] B. van Linder, W. van der Hoek, and J-J Ch. Meyer. Actions that make you change your mind. In *KI'95*, Proceedings of the 19th Annual German Conference on Artificial Intelligence, pages 185–196. Springer Verlag, 1995.
[31] R. Wassermann. Local diagnosis. In *Proceedings of the Eighth International Workshop on Nonmonotonic Reasoning*, 2000.
[32] M.A. Williams. On the logic of theory base change. In C. MacNish, D. Pearce, and L.M. Pereira, editors, *Logics in Artificial Intelligence*, volume 838 of *Lectures Notes in Artificial Intelligence*, pages 86–105. Springer Verlag, 1994.
[33] Y. Zhang and N.Y. Foo. On propositional knowledge base updates. *Australian Journal of Intelligent Information Processing Systems*, 2:20–29, 1995.

Giuseppe Primiero
Centre for Logic and Philosophy of Science
Philosophy and Moral Science Department
Ghent University, Belgium
E-mail: Giuseppe.Primiero@UGent.be