Arbitration
(or how to merge knowledge bases)

Paolo Liberatore, Marco Schaerf
Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza",
via Salaria 113, 00198 Roma, Italia
Tel: +39-6-49918332, Fax: +39-6-85300849
email: {liberato,schaerf}@dis.uniroma1.it

A very preliminary and much shorter version of this work appeared in [1]. Work partially supported by ASI (Italian Space Agency) and MURST (Italian Ministery for University and Scientific Research)

DRAFT

Abstract

Knowledge-based systems must be able to "intelligently" manage a large amount of information coming from different sources and at different moments in time. Intelligent systems must be able to cope with a changing world by adopting a "principled" strategy. Many formalisms have been put forward in the AI and DB literature to address this problem. Among them, belief revision is one of the most successful frameworks to deal with dynamically changing worlds. Formal properties of belief revision have been investigated by Alchourron, Gärdenfors and Makinson, who put forward a set of postulates stating the properties that a belief revision operator should satisfy. Among these properties, a basic assumption of revision is that the new piece of information is totally reliable and, therefore, must be in the revised knowledge base.

Different principles must be applied when there are two different sources of information and each one has a different view of the situation, the two views contradicting each other. If we do not have any reason to consider any of the sources completely unreliable, the best we can do is to "merge" the two views in a new and consistent one, trying to preserve as much information as possible. We call this merging process *arbitration*. In this paper we investigate the properties that any arbitration operator should satisfy. In the style of Alchourron, Gärdenfors and Makinson we propose a set of postulates, analyze their properties and propose actual operators for arbitration.

Keywords

Knowledge Representation, Belief Revision, Merging of Knowledge Bases, Databases Integration, Arbitration Operators.

I. Introduction and Motivations

One of the main challenges of today's knowledge and databases systems is their ability to manage a large amount of information coming from different sources and at different moments in time. Advanced knowledge-based systems must cope with a changing world and with not completely reliable sources of information by adopting a "principled" strategy. In the context of logic-based representations of information, this problem has been analyzed in the fields of databases, artificial intelligence and logic.

In the field of databases, the necessity of being able to integrate databases from multiple organizations into a single database has been clearly pointed out by Silberschatz, Stonebraker and Ullman in [2]. In particular, they point out that the next generation of database systems should address the problems of incompleteness and inconsistency arising from the integration of distributed heterogeneous databases.

The main focus of this analysis has been to provide principles (postulates) that guide the process of accommodating new information into an existing logical theory. In the field of belief revision, Alchourron, Gärdenfors and Makinson in [3], [4] proposed a set of postulates (the AGM postulates from now on) that any operator revising a knowledge should satisfy. Katsuno and Mendelzon in [5] show that the AGM postulates correctly capture the process of acquiring new information about a given world, but are inadequate to model the process of acquiring new information about a changing scenario. For this purpose, they introduce a new set of postulates (KM postulates) to model the process of updating a knowledge base.

While the AGM and KM postulates define general properties, actual belief revision and update operators have been proposed by several authors (see, for example, [6], [7], [8], [9]).

Both update and revision deal with the problem of accommodating new, and completely reliable, information into an existing body of knowledge. A basic assumption of update and revision is that the new piece of information must be in the revised (or updated) knowledge base.

Suppose that we are in a completely different scenario. There are two different sources of information and each one has a different view of the situation, the two views (possibly) contradicting each other. If we do not have any reason to consider any of the sources completely unreliable, the best we can do is to "merge" the two views in a new and consistent one, trying to preserve as much information as possible. We call this merging process *arbitration*.

A motivating example is now in order:

Example 1: Mark has some savings and wants to invest them in the Stock Exchange. However, he is not an expert in economics. For this reason he asks two friends, John and Robert who have some knowledge of economics, which shares he should buy. John believes that the shares x will increase and that at least one of y and z will also increase. On the other hand, Robert believes that x is going to drop, while y is going up. Employing a propositional language we use the letters x, y and z to denote that the corresponding share will rise. The two suggestions can be represented as:

$$john \equiv x \wedge (y \vee z) \qquad robert \equiv \neg x \wedge y$$

Shares y seem the obvious choice, but Mark wants to be sure of his investment, so he is going to buy shares only if he believes that they are going to rise. In the process of making sense of these two suggestions, Mark has several choices:

1. He believes that John is more reliable than Robert and thus he is not sure that y will rise;
2. He believes that Robert is more reliable and thus concludes that y will definitely rise.
3. He considers the two suggestions equally trustworthy. Since he cannot believe both simultaneously ($john \wedge robert$ is inconsistent), he believes their disjunction. Notice that $john \vee robert$ does not imply y.
4. He considers the two forecasts equally trustworthy, but $john \vee robert$ is too indefinite. Thus he merges the two suggestions by believing a new formula, called the arbitration and denoted as $john \triangle robert$. It seems reasonable that the arbitration is more definite than the disjunction and that the result of arbitrating between $john$ and $robert$ implies y.

The idea of arbitration and a first formal study of its properties was first done by Revesz in [10], where he discusses the properties that an arbitration operator should satisfy. In this paper we propose a different characterization of arbitration. In the style of the work on belief revision and update, we propose a set of postulates for arbitration and argue for their plausibility. Furthermore, we analyze which properties can be imposed to our operators without trivializing them, how actual arbitration operators can be defined and then prove a representation theorem for arbitration operators. We then compare our proposal with Revesz's work showing that his original proposal does not satisfy most of our postulates.

The paper is organized as follows: In Section II we give some preliminaries on belief revision, update and arbitration as presented by Revesz. In Section III we propose the basic principles for arbitration and discuss their importance, while in Section IV we analyze in detail the properties enjoyed by an arbitration operator. In Section V we consider other principles desirable for arbitration and their impact. In Section VI we introduce specific arbitration operators, analyze whether they satisfy all principles and show their behavior by means of examples. In Section VII we compare our approach with related work which appeared in the literature. Finally, in Section VIII we draw some conclusions and discuss open problems and future work.

II. Preliminaries

Following Katsuno and Mendelzon [11], in this paper we restrict our attention to propositional languages over a finite alphabet V of propositional letters. A knowledge base k is a finite set of propositional formulae. An interpretation is a mapping from the variables in V into the set $\{\mathsf{true}, \mathsf{false}\}$. An interpretation M is a model of a formula a if and only if it makes it true (we denote this by $M \models a$). Interpretations and models of propositional formulae will be denoted as sets of letters (those which are mapped into true). If k and a are propositional formulae such that all models of k are models of a, i.e. if a is a logical consequence of k, we write $k \models a$.

Furthermore, a *literal* is either a propositional variable (letter) or its negation. A *clause* is a disjunction of literals. A *CNF* formula is a conjunction of clauses.

We use lower case letters (a, $b, \ldots$) to denote propositional formulae, while the corresponding capital ones (A, $B \ldots$) denote the set of models of the formula. Where this convention cannot be applied, following [11] we use $Mod(a)$ to denote the set of models of a and $form(A)$ to denote one formula whose models are A.

A. Principles of belief revision

In the literature, the first formal studies on the principles of belief revision have been presented by Alchourron, Gärdenfors and Makinson in [3], [4]. In these papers they present a set of postulates that all revision operators should satisfy. These postulates, known as the AGM postulates, assume that the revision operator applies to a deductively-closed set of formulae. In order to make the presentation more homogeneous, we present a reformulation of these postulates where the revision operator applies to propositional formulae. More precisely, we denote with k the knowledge base (that is the existing logical theory), with a the revising formula (that is the new information) and with $*$ the revision operator. This formulation has been presented by Katsuno and Mendelzon in [11], where they prove that this set of postulates is equivalent to the original one.

Thus, the AGM postulates for (finite) propositional knowledge bases are:

AGM1 $k * a$ implies a.

AGM2 If $k \wedge a$ is satisfiable then $k * a \equiv k \wedge a$.

AGM3 If a is satisfiable then $k * a$ is also satisfiable.

AGM4 If $k_1 \equiv k_2$ and $a_1 \equiv a_2$ then $k_1 * a_1 \equiv k_2 * a_2$.

AGM5 $(k * a) \wedge b$ implies $k * (a \wedge b)$.

AGM6 If $(k * a) \wedge b$ is satisfiable then $k * (a \wedge b)$ implies $(k * a) \wedge b$.

The intuitive meaning of the postulates is simple to understand. *AGM1* states that the new information a is always retained in the revision. When consistency is preserved, *AGM2* postulates that a is simply added to k. *AGM3* states that inconsistency cannot be introduced unless a is inconsistent. Furthermore, because of *AGM4* the revision operator obeys the Principle of Irrelevance of Syntax, and postulates *AGM5* and *AGM6* impose constraints on the behavior of revision in the presence of conjunctions.

Katsuno and Mendelzon in [11] have shown that to any revision operator satisfying *AGM1*-*AGM6* corresponds a family of reflexive, transitive and total orderings over the set of interpretations, one for each set of models F (in the sequel we denote $\preceq_F$ the ordering corresponding to F). Given a revision $*$ and its corresponding family of orderings, the following relation holds.

$$M(k * a) = Min(M(a), \preceq_{Mod(k)}) \tag{1}$$

where $Min(S, \leq)$ (S being a set and $\leq$ an ordering) is the set of the minimal elements of S wrt $\leq$.

Any given ordering $\preceq_F$ has the so-called property of *faithfulness*, that can be summarized as follows:

1. If $I \in F$ then $I \preceq_F J$ for any interpretation J.
2. If $I \in F$ and $J \notin F$ then $J \preceq_F I$ does not hold.

Roughly speaking, the models of F are exactly the minimal elements of $\preceq_F$.

The intuitive meaning of $\preceq_F$ is that it measures the distance between a formula represented by the set of models F and a model: given two models I and J, $I \preceq_F J$ holds if and only if I is considered more plausible to an agent believing F. In this sense, achieving the principle of minimal change, equation 1 means that the result of a revision is constituted by the models of a that are closer to k.

In the sequel we use $I \prec_F J$ to denote that $I \preceq_F J$ holds but $J \preceq_F I$ does not. Furthermore, $I \simeq_F J$ denotes that both $I \preceq_F J$ and $J \preceq_F I$ hold.

B. Revision operators

At the same time as the methodological work on principles of belief revision appeared, actual operators have been proposed in the literature. We now recall some of the revision operators, classifying them into formula-based and model-based ones. A more thorough exposition can be found in [12]. We use the following conventions: the symmetric difference between two sets S_1, S_2 is denoted by $diff(S_1, S_2)$ (that is $diff(S_1, S_2) = (S_1 - S_2) \cup (S_2 - S_1)$), while the expression $|S|$ denotes the number of elements of the set S. If S is a set of sets, $\cap S$ denotes the set formed intersecting all sets of S, and analogously $\cup S$ for union. To denote the minimal elements of a set S wrt a given ordering $\leq_*$ we write $Min(S, \leq_*)$, while $Max(S, \leq_*)$ denotes the maximal ones.

Formula-based approaches operate on the formulae which appear syntactically in the knowledge base k. Their goal is to preserve as many as possible of the formulae appearing in k. Obviously, they do not satisfy syntax-independence. To introduce them more precisely we need some notation. Let $C(k, a)$ be the set of the subsets of k which are consistent with the revising formula a:

$$C(k, a) = \{k' \subseteq k \ \mid \ k' \cup \{a\} \not\models \bot\}$$

and let $W(k, a)$ be the set of the maximal subsets of $C(k, a)$:

$$W(k, a) = Max(C(k, a), \subseteq)$$

Intuitively, the set $W(k, a)$ contains all the maximal subsets of k that we may retain when inserting a. Note that when $k \wedge a$ is consistent, the only maximal element of $C(k, a)$ is k itself. Therefore, $W(k, a) = k$.

Fagin-Ullman-Vardi. In [6] (and independently by Ginsberg in [7]) the revised knowledge base is defined as a set of theories: $k *_{FUV} a \doteq \{k' \cup \{a\} \mid k' \in W(k, a)\}$. That is, the result of revising k is the set of all maximal subsets of k consistent with a, plus a. Logical consequence in the revised knowledge base is defined as logical consequence in each of the theories, i.e. $k *_{FUV} a \models q$ iff for all $k' \in W(k, a)$, $k' \cup \{a\} \models q$. In other words, Fagin, Ullman and Vardi consider all sets in $W(k, a)$ equally plausible and inference is defined skeptically, i.e. q must be a consequence of each set.

WIDTIO. Since there may be exponentially many new theories in $k *_{FUV} a$, a simpler (but somewhat drastic) approach is the so-called WIDTIO (When In Doubt Throw It Out), which is defined as $k *_W a \doteq (\cap W(k, a)) \cup \{a\}$. That is, only formulae which belong to all maximal subsets of k are retained.

As pointed out before, note that formula-based approaches are sensitive to the syntactic form of the theory. That is, the revision with the same formula P of two logically equivalent theories T_1 and T_2, may yield different results, depending on the syntactic form of T_1 and T_2. We illustrate this fact and the different behavior of the two revision operators through an example.

Example 2: Let k and a be defined as:

$$\begin{aligned} k &= \{x \vee y, \neg z \vee y, \neg x \vee z\} \\ a &= (\neg y) \end{aligned}$$

Then $W(k, a)$ is composed of three sets:

$$W(k, a) = \{(x \vee y, \neg z \vee y), (x \vee y, \neg x \vee z), (\neg z \vee y, \neg x \vee z)\}$$

Thus, $k *_{FUV} a = \{(x \vee y, \neg z \vee y, \neg y), (x \vee y, \neg x \vee z, \neg y), (\neg z \vee y, \neg x \vee z, \neg y)\}$, while $k *_W a = \{\neg y\}$. The consequences we can derive from $k *_{FUV} a$ are different from those derivable from $k *_W a$, since, $x \vee \neg z$ is a consequence of $k *_{FUV} a$, but not of $k *_W a$.

Now let $k' = \{y, \neg x \vee z\}$. It is easy to show that k and k' are logically equivalent. Nevertheless $k' *_{FUV} a = k' *_W a = \{\neg x \vee z, \neg y\}$, which is different from both $k *_{FUV} a$ and $k *_W a$. $\diamond$

Model-based approaches instead operate by selecting the models of a on the basis of some notion of proximity to the models of k. Model-based approaches assume k to be a single formula, and if k is a set of formulae it is implicitly interpreted as the conjunction of all the elements. We consider approaches where proximity between models of a and models of k is defined globally by considering *all* models of k. In other words, these approaches select the models of a that are closer to the models of k by considering at the same time all pairs of models $M \in Mod(k)$ and $N \in Mod(a)$ and find all the closest pairs. Thus, these methods only differ on how they measure distance between models.

To make the definitions more precise we need to introduce some notation. Let:

$$\delta(k, a) \doteq Min(diff(M, N) : M \in Mod(k), N \in Mod(a), \subseteq)$$

and

$$\mu(k, a) \doteq Min(|diff(M, N)| : M \in Mod(k), N \in Mod(a), \leq)$$

Note that $\delta(k, a)$ is a set of sets of propositional letters, while $\mu(k, a)$ is a non-negative number.

Satoh. In [9], the models of the revised knowledge base are defined as $Mod(k *_S a) \doteq \{N \in Mod(a) \mid \exists M \in Mod(k) : diff(N, M) \in \delta(k, a)\}$. That is, Satoh selects all models of a whose distance from a model of k is minimal (by set containment).

Dalal. This approach is similar to Satoh's one, but measures distance between models based on cardinality of the difference set. Let $d = \mu(k, a)$; in [8] the models of a revised theory are defined as $Mod(k *_D a) \doteq \{N \in Mod(a) \mid \exists M \in Mod(k) : |diff(N, M)| = d\}$. That is, Dalal selects all models of a whose distance from a model of k is minimal (by cardinality of the difference set).

We show their different behavior with an example:

Example 3: Let k and a be defined as:

$$\begin{aligned} k &= \{x, y, \neg z\} \\ a &= (\neg x \wedge \neg z) \vee (x \wedge \neg y \wedge z) \end{aligned}$$

Note that k has only one model $J = \{x, y\}$ while a has three models:

$$\begin{aligned} I_1 &= \emptyset \\ I_2 &= \{y\} \\ I_3 &= \{x, z\} \end{aligned}$$

The set differences between the model of k and each model of a are:

$diff(I, J)$	$I_1 = \emptyset$	$I_2 = \{y\}$	$I_3 = \{x, z\}$
$J = \{x, y\}$	$\{x, y\}$	$\{x\}$	$\{y, z\}$

Hence, the minimal differences between J and models of a are $\delta(k, a) = \{(x), (y, z)\}$, while $\mu(k, a) = 1$. Therefore, the models of $k *_S a$ are I_2, I_3, and the single model of $k *_D a$ is I_2. $\diamond$

C. Revesz's Definition of Arbitration

More recently, Revesz in [10] introduced the idea of arbitration of logical theories. The arbitration of two theories a and b (denoted as $a \triangle b$) is a new theory that retains "as much as possible" of the information contained in a and b.

This operation is very different from classical revision, because revision is not commutative. In fact, the second argument (a in the revision of $k * a$) is more important than the first one and it is always true in the result of the revision, as stated by postulate *AGM1*. In arbitration, the two operands are considered to be of equal importance, hence the postulate of commutativity will hold: $a \triangle b = b \triangle a$. We can say that arbitration adds two sets of (possibly incompatible) formulae where neither one is preferred with respect to the other one.

Revesz defines arbitration as a derived notion of *model-fitting*, where a model-fitting operator $\triangleright$ obeys the following postulates:

R1 $k \triangleright a$ implies a.

R2 If k is unsatisfiable then $k \triangleright a$ is unsatisfiable.

R3 If both k and a are satisfiable then $k \triangleright a$ is satisfiable.

R4 If $k_1 \equiv k_2$ and $a_1 \equiv a_2$ then $k_1 \triangleright a_1 \equiv k_2 \triangleright a_2$.

R5 $(k \triangleright a) \wedge b$ implies $k \triangleright (a \wedge b)$.

R6 If $(k \triangleright a) \wedge b$ is satisfiable then $k \triangleright (a \wedge b)$ implies $(k \triangleright a) \wedge b$.

R7 $(k_1 \triangleright a) \wedge (k_2 \triangleright a)$ implies $(k_1 \vee k_2) \triangleright a$.

R8 If $(k_1 \triangleright a) \wedge (k_2 \triangleright a)$ is satisfiable then $(k_1 \vee k_2) \triangleright a$ implies $(k_1 \triangleright a) \wedge (k_2 \triangleright a)$.

Arbitration is then defined as:

$$a \triangle b \doteq (a \vee b) \triangleright \mathsf{true} \tag{2}$$

Therefore, arbitration is the operation that fits the disjunction of a and b with respect to the set of all interpretations.

In the same paper Revesz introduces an actual model-fitting operator (that we denote as $\triangleright_R$) defined model-theoretically. Let $\leq_a$ be defined as: $I \leq_a J$ if and only if $odist(a, I) \leq odist(a, J)$, where

$odist(a, I) \doteq Max(\{|diff(I, J)| : J \in Mod(a)\}, \leq)$. The model-fitting operator is model-theoretically defined as

$$Mod(a \triangleright_R b) \doteq Min(Mod(b), \leq_a)$$

and the associated arbitration operator ($\triangle_R$) is defined as:

$$a \triangle_R b = (a \vee b) \triangleright_R \mathsf{true} \quad (3)$$

The model-fitting operator $\triangleright_R$ satisfies all postulates *R1* to *R8*. See [10] for more details.

While we believe that the idea of arbitration is very interesting, we do not consider the formalization presented by Revesz fully satisfactory. Before presenting our critique to Revesz's proposal, we introduce our postulates for arbitration and argue for their importance. Our criticism to Revesz's proposal is then discussed in Section VII.

III. Postulates

There are many properties that seem reasonable for an arbitration operator. We first introduce a set of 5 postulates stating the basic properties that all arbitration operators should satisfy. These postulates are discussed in detail and for each one we present compelling (in our opinion) motivations. Afterwards, we present 3 more postulates that enforce more specific properties. While we consider all postulates necessary, this second set of postulates imposes less compelling constraints.

In the following section we discuss the properties of the arbitration operators satisfying all our postulates.

Following the same conventions introduced in Section II, we introduce arbitration as an operator on propositional formulae. Thus, arbitration postulates refer to this representation. Nevertheless, as already noticed by other authors, proofs are easier if we assume that both the knowledge base and the revising formula are represented as sets of models. For this reason, we also present the arbitration postulates for a model-based representation.

A. Basic Postulates

The first and distinguishing property of arbitration is *commutativity*, enforced by the following postulate:

A1 $a \triangle b \equiv b \triangle a$.

The result of arbitrating between two knowledge bases should only contain information present in either a or b. That is, no information can be created by the arbitration. The following postulate enforces that the result of arbitration is less definite than the conjunction of the two knowledge bases:

A2 $a \wedge b$ implies $a \triangle b$.

If the two knowledge bases do not contradict each other we can assume both true and have as a result their conjunction. This formalizes the principle of retaining as much as possible of the information contained in the arbitrating formulae. In the example of the last section, if John's suggestion was consistent with Robert's one, we can suppose that both are truthful and come to believe $john \wedge robert$. This allows us to use all the information, and hence, to obtain more knowledge. Therefore, we assume the postulate:

A3 if $a \wedge b$ is satisfiable then $a \triangle b$ implies $a \wedge b$.

We want to preserve consistency of arbitration "as much as possible", and therefore we require that the result of an arbitration should be inconsistent if and only if both knowledge bases are inconsistent.

A4 $a \triangle b$ is unsatisfiable if and only if both a and b are unsatisfiable.

Another fundamental property we want to enforce on arbitration operators is syntax-independence.

A5 If $a_1 \equiv a_2$ and $b_1 \equiv b_2$ then $a_1 \triangle b_1 \equiv a_2 \triangle b_2$.

As already pointed out, proofs are simpler if we refer to a model-based representation of the knowledge base and the revising formula. For this reason we also present the corresponding postulates for a model-based representation, where both the operands and the result are sets of models:

M1 $A \triangle B = B \triangle A$.

M2 $(A \cap B) \subseteq (A \triangle B)$.

M3 If $A \cap B \neq \emptyset$ then $(A \triangle B) \subseteq (A \cap B)$.

M4 $A \triangle B = \emptyset$ if and only if $A = \emptyset$ and $B = \emptyset$.

M5 true.

Note that rephrasing postulates *A1*-*A4* in terms of sets of models is immediate, but is only possible if we have syntax-independence (Postulate *A5* holds). As a result, in terms of sets of models, postulate *A5* does not impose any other constraint on arbitration.

B. More Specific Postulates

More specific properties can be derived from the corresponding revision properties, and from other weaker natural conditions of the intuitive meaning of arbitration.

First of all, we require a compositionality property of arbitration operators.

A6 $a \triangle (b \vee c) = \begin{cases} a \triangle b & \text{or} \\ a \triangle c & \text{or} \\ (a \triangle b) \vee (a \triangle c) \end{cases}$

This postulate guarantees that arbitration of composite formulae can be obtained via the composition of the arbitration of subformulae. It is also the natural extension to arbitration of the AGM postulates *AGM5* and *AGM6*. Gärdenfors in [4, property 3.16 pg. 57] states that the revision postulates *AGM5* and *AGM6* are equivalent (if the other postulates hold) to our postulate *A6* (with $*$ instead of $\triangle$).

While postulate *A2* forces arbitration to be less definite than conjunction, no constraints have been imposed so far to prevent arbitration from becoming very vague. For example, the following arbitration operator:

$$a \triangle b = \begin{cases} a \wedge b & \text{if consistent} \\ \mathsf{true} & \text{otherwise} \end{cases}$$

satisfies all postulates *A1*-*A6*. In order to avoid excessive loss of information, we require:

A7 $a \triangle b$ implies $a \vee b$.

Since arbitration is a merging process of two pieces of information, a very reasonable constraint is to require that both pieces contribute to the final result. This is obtained via postulate *A8*, which ensures that each knowledge base (if consistent) will be consistent with the result of the arbitration.

A8 If a is satisfiable then $a \wedge (a \triangle b)$ is also satisfiable.

These postulates can be rephrased in terms of a model-based representation, thus obtaining:

M6 $A \triangle (B \cup C) = \begin{cases} A \triangle B & \text{or} \\ A \triangle C & \text{or} \\ (A \triangle B) \cup (A \triangle C) & \end{cases}$

M7 $(A \triangle B) \subseteq (A \cup B)$.

M8 If $A \neq \emptyset$ then $A \cap (A \triangle B) \neq \emptyset$.

Coming back to the stock exchange example presented in Section I, we note that *john* can be rewritten as $(x \wedge y) \vee (x \wedge z)$. Therefore, postulate *A6* requires that the arbitration of *john* and *robert* can be obtained by arbitrating between *robert* with each disjunct ($x \wedge y$ and $x \wedge z$) of *john*, and then composing the results.

Together, postulates *A2* and *A7* state that the arbitration of sets corresponding to the formulae a and b is at most as definite as $a \wedge b$ (when it is not contradictory), and at least as definite as $a \vee b$. Notice that $john \vee robert$ is very indefinite (e.g. it does not imply y), while it seems reasonable that y is true in all models of the knowledge base $john \triangle robert$.

A reasonable constraint that we may require of arbitration is the following one:

M78 $\{I\} \triangle \{J\} = \{I, J\}$

where I and J are single interpretations, or its formula-based equivalent:

A78 If a and b are complete formulae, then $a \triangle b = a \vee b$

where a formula is complete if it has a single model. This is to say that when we arbitrate between two complete formulae we have no other choice than concluding their disjunction. In the model-based representation, *M78* expresses the fact that if the two formulae have a single model, we have no reason to choose one over the other and, thus, we accept both.

However, this postulate is not independent of the other ones, as shown by the following theorem:

Theorem 1: Every arbitration satisfying *M1*-*M6* (*A1*-*A6*) satisfies *M78* (*A78*) if and only if it satisfies *M7* (*A7*) and *M8* (*A8*).

Proof. If $\triangle$ satisfies *M7* it follows that $\{I\} \triangle \{J\} \subseteq \{I, J\}$ and *M8* implies that $\{I\} \cap \{I\} \triangle \{J\} \neq \emptyset$ and $\{J\} \cap \{I\} \triangle \{J\} \neq \emptyset$. Therefore, $\{I\} \triangle \{J\} = \{I, J\}$. Now assume that $\triangle$ satisfies *M1*-*M6* and *M78*.

By *M6* it follows that $A \triangle B \subseteq \bigcup_{I \in A, J \in B}(\{I\} \triangle \{J\})$. Thus, *M78* implies that $A \triangle B \subseteq A \cup B$ and that if $A \neq \emptyset$ then $A \cap (A \triangle B) \neq \emptyset$. ■

Arbitration postulates are very close in spirit to the AGM postulates, thus it is not surprising that revision can be obtained through arbitration.

Theorem 2: Let $\triangle$ be an arbitration operator satisfying postulates *A1*-*A8*, then the revision operator defined as $k *_{\triangle} a \doteq (k \triangle a) \wedge a$ satisfies postulates *AGM1*-*AGM6*.

Proof. Postulates *AGM1*-*AGM4* immediately follow. *AGM1* follows from the fact that $(k \triangle a) \wedge a$ implies a, *AGM2* follows from postulate *A3*, in fact, if $k \wedge a$ is satisfiable then $k *_{\triangle} a = (k \wedge a) \wedge a = k \wedge a$. Postulate *A8* implies that *AGM3* holds, while *AGM4* is trivial. As already pointed out, Gärdenfors in [4, property 3.16 pg. 57] states that revision postulates *AGM5* and *AGM6* are equivalent (if the other postulates hold) to:

$$AGM56 \;\; a * (b \vee c) = \begin{cases} a * b & \text{or} \\ a * c & \text{or} \\ (a * b) \vee (a * c) \end{cases}$$

If $a \wedge (b \vee c)$ is satisfiable the property trivially holds, so let us assume that $a \wedge (b \vee c)$ is unsatisfiable. We now prove that postulate *AGM56* holds for the revision operator $*_{\triangle}$. By definition, $a *_{\triangle} (b \vee c) = (a \triangle (b \vee c)) \wedge (b \vee c)$. With simple rewritings we obtain $((a \triangle (b \vee c)) \wedge b) \vee ((a \triangle (b \vee c)) \wedge c)$. By postulate *A6* we have that:

$$a *_{\triangle} (b \vee c) = \begin{cases} ((a \triangle b) \wedge b) \vee ((a \triangle b) \wedge c) & \text{or} \\ ((a \triangle c) \wedge b) \vee ((a \triangle c) \wedge c) & \text{or} \\ ((a \triangle b) \wedge b) \vee ((a \triangle b) \wedge c) \vee ((a \triangle c) \wedge b) \vee ((a \triangle c) \wedge c) \end{cases}$$

We now show that $((a \triangle b) \wedge c)$ implies $((a \triangle b) \wedge b)$. By *A7* we have that $a \triangle b$ implies $a \vee b$ and, therefore, $((a \triangle b) \wedge c) = ((a \triangle b) \wedge (a \vee b) \wedge c)$. With simple rewritings we obtain $(a \triangle b) \wedge ((a \wedge c) \vee (b \wedge c))$. Since we assumed that $a \wedge (b \vee c)$ is unsatisfiable, so is $a \wedge c$. As a consequence $((a \triangle b) \wedge c) = ((a \triangle b) \wedge b) \wedge c$. Hence, it implies $((a \triangle b) \wedge b)$. For the same reason we also have that $((a \triangle c) \wedge b)$ implies $((a \triangle c) \wedge c)$. We can now simplify $((a \triangle b) \wedge b) \vee ((a \triangle b) \wedge c)$ as $(a \triangle b) \wedge b)$ and $((a \triangle c) \wedge b) \vee ((a \triangle c) \wedge c)$ as $(a \triangle c) \wedge c)$. The above equation can be rewritten as:

$$a *_{\triangle} (b \vee c) = \begin{cases} (a \triangle b) \wedge b & \text{or} \\ (a \triangle c) \wedge c & \text{or} \\ ((a \triangle b) \wedge b) \vee ((a \triangle c) \wedge c) \end{cases}$$

Substituting $(a \triangle b) \wedge b$ with $a *_{\triangle} b$ and $(a \triangle c) \wedge c$ with $a *_{\triangle} c$ we obtain the postulate *AGM56*. ■

Reformulating arbitration in terms of belief revision is much more complex. For some specific arbitration and revision operators we show in Section VI that the following reduction holds:

$$a \triangle b = (a * b) \vee (b * a)$$

However, in general, choosing a revision operator satisfying *AGM1*-*AGM6* does not guarantee that the corresponding $\triangle$ satisfies all postulates *A1*-*A8*.

IV. Properties of Arbitration Operators

In this section, we study the properties of arbitrations satisfying *M1-M8*. A first observation comes from [5]: every revision satisfying AGM postulates *AGM1-AGM6* can be expressed with a suitable family of reflexive, transitive and total ordering $\leq_{Mod(k)}$ (one for each k) over the set of models of the given alphabet:

$$Mod(k * a) = Min(Mod(a), \leq_{Mod(k)})$$

This property expresses the result of revising a formula k via an ordering $\leq_{Mod(k)}$ that compares two models. The meaning of $I \leq_{Mod(k)} J$ is that, given the knowledge base k, we consider the model I "more or equally likely" than J, that is, the world represented by I is more (or equally) plausible than J.

As a result, this measure of plausibility is expected to vary on different bases k. For example, if $Mod(k) = \{\{a, b, c\}\}$, the model $I = \{b, c\}$ differs from $\{a, b, c\}$ only on the value of the variable a, thus it should be always considered more plausible than $J = \{\}$, that differs on all the variables, at least from a "syntactic" point of view.

However, the plausibility ordering may depend also on the meaning of the atoms in the considered alphabet. For example, there is no reason to consider I more plausible than $Z = \{a, c\}$, wrt the base k. As a result, in such cases most of the revision operators proposed in the literature, in such cases do not chose one model between I and J as the most likely world, so both of them are present in the revised knowledge base.

Merging the knowledge bases means that we have to take into account two different views of the world, and even two beliefs about plausibility of the models. For simplicity, we assume from now on a model-based representation of the arbitrating knowledge bases and of the result. Therefore, both the knowledge bases and the result will be a set of models.

We can suppose that any knowledge base A or B comes with its own plausibility ordering, say $\leq_A$ and $\leq_B$, each representing what an agent that knows A or B considers more or less likely.

We use $A <_F B$ to denote that $A \leq_F B$ holds but $B \leq_F A$ does not, and $A \cong_F B$ to denote that both $A \leq_F B$ and $B \leq_F A$ hold. Note that the ordering $\leq_F$ applies to *sets* of models, and not to single models as the ordering $\preceq_F$. Given a set of models A, we use the notation $\hat{A}$ to denote the set of sets $\{\{I\} | I \in A\}$. The following theorem shows how arbitrations can be expressed in terms of these orderings, and that the following properties of $\leq_F$ are equivalent to our postulates *M1-M8*.

L1 transitivity: if $A \leq_F B$ and $B \leq_F C$ then $A \leq_F C$

L2 if $A \subseteq B$ then $B \leq_F A$

L3 $A \leq_F A \cup B$ or $B \leq_F A \cup B$

L4 $B \leq_F C$ for every C iff $F \cap B \neq \emptyset$

L5 $A \leq_{C \cup D} B \Leftrightarrow \begin{cases} C \leq_{A \cup B} D \quad and \quad A \leq_C B & \text{or} \\ D \leq_{A \cup B} C \quad and \quad A \leq_D B \end{cases}$

Lemma 3: Any ordering $\leq_F$ satisfying *L1*-*L5* is reflexive and total.

Proof. Since $A \subseteq A$, from postulate *L2* it follows that $A \leq_F A$. Any relation satisfying *L1*-*L5* is total, in fact, given any set F and two sets of models A and B, we may have (postulate *L3*) either $A \leq_F A \cup B$ or $B \leq A_F \cup B$. In the first case, postulate *L2* gives $A \cup B \leq_F B$, and by transitivity $A \leq_F B$. In the second case $B \leq_F A \cup B$ plus $A \cup B \leq_F A$ gives $B \leq_F A$. ■

Lemma 4: Given an interpretation I, a set of models B and an ordering $\leq_F$, we have that $\{I\} \in Min(\hat{B}, \leq_F)$ iff $\{I\} \cong_F B$.

Proof. Suppose $B = \{B_1, \ldots, B_m\}$. From postulate *L3* it follows that either $\{B_1\} \leq_F B$ or $\{B_2, \ldots, B_m\} \leq_F B$. Recursively, we obtain the property that there exists a model B_i in B such that $\{B_i\} \leq_F B$. From *L2* it follows that $B \leq_F \{B_i\}$ and thus $\{B_i\} \cong_F B$. On the other direction, if $\{I\} \cong_F B$, we have that $B \leq_F \{J\}$ for any model J, since $\{J\} \subseteq B$. By transitivity $\{I\} \leq_F \{J\}$ for any $J \in B$. ■

We can show that any ordering satisfying *L1*-*L5* induces an arbitration operator satisfying *M1*-*M8*.

Theorem 5: If the family of orderings $\{\leq_F \mid F \text{ is a set of models}\}$ satisfies *L1*-*L5* then the arbitration operator defined as $A \triangle B = \{I \mid \{I\} \in Min(\hat{A}, \leq_B) \cup Min(\hat{B}, \leq_A)\}$ satisfies *M1*-*M8*.

Proof. Suppose that *L1*-*L5* are satisfied. We show that $\triangle$ satisfies *M1*-*M8*.

M1 By the definition of $\triangle$ it follows that

$$\begin{aligned} A \triangle B &= \{I \mid \{I\} \in Min(\hat{A}, \leq_B) \cup Min(\hat{B}, \leq_A)\} \\ B \triangle A &= \{I \mid \{I\} \in Min(\hat{A}, \leq_B) \cup Min(\hat{B}, \leq_A)\} \end{aligned}$$

so these sets are equal.

M2 Let I be a model of both A and B, that is $I \in A \cap B$. From property *L4* it follows that

$$\{I\} \leq_A C \qquad \{I\} \leq_B C \qquad \text{for any set of models } C$$

Since $\{I\}$ belongs to $\hat{A}$ and it is minimal in this set wrt $\leq_B$, it is in $A \triangle B$.

M3 Suppose $A \cap B \neq \emptyset$. Let $I \in A \triangle B$, so either $\{I\} \in Min(\hat{A}, \leq_B)$ or $\{I\} \in Min(\hat{B}, \leq_A)$ hold. Suppose $\{I\} \in Min(\hat{A}, \leq_B)$, thus $I \in A$. We will prove that $I \in B$, and therefore $I \in A \cap B$

By hypothesis, $A \cap B$ is not empty. Let J be a model in $A \cap B$. From *L4* it follows that

$$\{J\} \leq_B C \quad \text{for any set C}$$

but $\{I\} \in Min(\hat{A}, \leq_B)$, so $\{I\} \leq_B \{J\}$ and by transitivity

$$\{I\} \leq_B C \quad \text{for any set C}$$

Using property *L4* in the other direction we have $B \cap \{I\} \neq \emptyset$, thus $I \in B$.

M4 The ordering $\leq_F$ is transitive, and, as previously proved, reflexive and total. Hence, $Min(\hat{A}, \leq_B)$ and $Min(\hat{B}, \leq_A)$ are both empty iff A and B are empty too.

M5 Always satisfied.

M6 Being $\leq_A$ a total relation, it can either be $B <_A C$ or $B \cong_A C$ or $C <_A B$, where $B <_A C$ denotes as usual that $B \leq_A C$ holds but $C \leq_A B$ does not. We will prove that in each of these three cases, postulate *M6* holds.

By Lemma 4 $\{I\} \in Min(\hat{B}, \leq_A)$ if and only if $\{I\} \cong_A B$.

($B <_A C$) Let us suppose $B <_A C$. From the previous lemma, we have that if B_i is minimal in B wrt the ordering $\leq_A$, then $\{B_i\} \cong_A B$ thus $\{B_i\} \leq_A B$. Together with $B <_A C$ this implies by transitivity $\{B_i\} <_A C$. But if $C_j \in C$ we have $C \leq_A \{C_j\}$ and $\{B_i\} <_A \{C_j\}$. Hence, $Min(\hat{B} \cup \hat{C}, \leq_A) = Min(\hat{B}, \leq_A)$.

Now he have to prove that $Min(\hat{A}, \leq_{B\cup C}) = Min(\hat{A}, \leq_B)$. Let $A_i \in A$: we have already proved that $\{A_i\} \in Min(\hat{A}, \leq_{B\cup C})$ if and only if $\{A_i\} \leq_{B\cup C} A$. Note that $\{A_i\} \cup A = A$, thus we can use postulate *L5* in this form:

$$\{A_i\} \leq_{B\cup C} A \quad \Leftrightarrow \quad \begin{cases} B \leq_A C \quad \text{and} \quad \{A_i\} \leq_B A \\ \text{OR} \\ C \leq_A B \quad \text{and} \quad \{A_i\} \leq_C A \end{cases}$$

By hypothesis, $B \leq_A C$ is true while $C \leq_A B$ is false, so we have $\{A_i\} \leq_{B\cup C} A \Leftrightarrow \{A_i\} \leq_B A$, that implies $\{A_i\} \in Min(\hat{A}, \leq_{B\cup C}) \Leftrightarrow \{A_i\} \in Min(\hat{A}, \leq_B)$.

($B \cong_A C$) As said above, this means that both $B \leq_A C$ and $C \leq_A B$ hold. We prove first that $Min(\hat{B} \cup C, \leq_A) = Min(\hat{B}, \leq_A) \cup Min(\hat{C} \leq_A)$. Suppose that $\{B_i\} \in Min(\hat{B}, \leq_A)$ and $\{C_i\} \in Min(\hat{C}, \leq_A)$. These formulae imply $\{B_i\} \cong_A B$ and $\{C_i\} \cong_A C$. From $B \cong_A C$ it follows that $\{B_i\} \cong_A \{C_i\}$, thus the minimal models of B and C have the same plausibility.

In order to prove that $Min(\hat{A}, \leq_{B\cup C}) = Min(\hat{A}, \leq_A) \cup Min(\hat{B}, \leq_C)$, consider a model $A_i \in A$, and apply postulate *L5*.

$$\{A_i\} \leq_{B\cup C} A \quad \Leftrightarrow \quad \begin{cases} B \leq_A C \quad \text{and} \quad \{A_i\} \leq_B A \\ \text{OR} \\ C \leq_A B \quad \text{and} \quad \{A_i\} \leq_C A \end{cases}$$

But now both $B \leq_A C$ and $C \leq_A B$ holds, thus $\{A_i\} \leq_{B\cup C} A \Leftrightarrow \{A_i\} \leq_B$ or $\{A_i\} \leq_C A$. As a result, $\{A_i\} \in Min(\hat{A}, \leq_{B\cup C}) \Leftrightarrow \{A_i\} \in Min(\hat{A}, \leq_B)$ or $\{A_i\} \in Min(\hat{A}, \leq_C)$.

($C <_A B$) This case is similar to $B <_A C$ (just exchange B and C).

M7 By definition, the minimal elements of A are always in A, and the same happens for B.

M8 Since $\leq_F$ is transitive, reflexive and total, any set has at least one minimal element (we are considering propositional logic, so only finite sets of models are allowed). Hence, $Min(\hat{A}, \leq_B)$ is non-empty.

■

We can also show that any arbitration can be expressed with a family of orderings.

Theorem 6: If $\triangle$ satisfies *M1-M8*, then it can be associated to a family of orderings $\{\leq_F \mid F \text{ is a set of models}\}$ satisfying *L1-L5* such that

$$A \triangle B = \{I \mid \{I\} \in Min(\hat{A}, \leq_B) \cup Min(\hat{B}, \leq_A)\}$$

Proof. We divide the proof in three steps. First, we define the family of orderings associated to a given arbitration. In the second step, we prove that $A \triangle B$ is equivalent to $\{I \mid \{I\} \in Min(\hat{A}, \leq_B) \cup Min(\hat{B}, \leq_A)\}$. Finally, we prove that such a family of orderings satisfies postulates *L1-L5*.

We define the family of ordering. From Theorem 2 it follows that, given an arbitration $\triangle$ satisfying postulates *M1-M8*, the induced revision $*_\triangle$ defined by $A *_\triangle B = (A \triangle B) \cap B$ is a valid AGM revision. Katsuno and Mendelzon in [11] have shown that to any revision operator satisfying *AGM1-AGM6* corresponds a family of reflexive, transitive and total faithful orderings $\preceq_F$ (one for each set of models F) over the sets of the interpretations. Therefore, $*_\triangle$ can be defined in terms of this orderings as:

$$A *_\triangle B = Min(B, \preceq_A)$$

Now, we define a new ordering $\leq_B$, via the corresponding $\preceq_B$, as:

$$C \leq_B D \;\Leftrightarrow\; \forall J \in D \text{ there exists } I \in C \text{ such that } I \preceq_B J$$

This defines the family of orderings associated to an arbitration. Note that $\leq_B$ compares sets of models, rather than models and, by definition, $\leq_F$ is "equivalent" to $\preceq_F$ when C and D are singletons. Recall that $\hat{A} = \{\{I\} : I \in A\}$ and $\hat{B} = \{\{I\} : I \in B\}$, that is $\hat{A}$ is a set of sets, where the elements are the models of A and similarly for $\hat{B}$. Using some rewriting we can express arbitration in terms of orderings:

$$\begin{aligned} A \triangle B &= [(A \triangle B) \cap A] \cup [(A \triangle B) \cap B] \\ &= (B *_\triangle A) \cup (A *_\triangle B) \\ &= \{I | I \in Min(A, \preceq_B) \cup Min(B, \preceq_A)\} \\ &= \{I \mid \{I\} \in Min(\hat{A}, \leq_B) \cup Min(\hat{B}, \leq_A)\} \end{aligned}$$

This proves that given $\triangle$, we can find a family of orderings representing it. To prove the theorem, we have to show that this ordering satisfies *L1-L5*.

We want to point out that given a family of transitive, reflexive, total and faithful orderings $\preceq_F$, the corresponding family of orderings $\leq_F$ satisfies postulates *L1-L4* but not necessarily *L5*. In order to prove *L5* we need to directly use postulates *M1-M8*. As a consequence, this is not an exact reformulation of arbitration in terms of revision.

L1 Follows from the transitivity of $\preceq_F$.

L2 If $A \subseteq B$ then $\forall I \in A \exists J. \in B$ s. t. $(I \preceq J)$, therefore $B \leq_F A$. If $A = \emptyset$ the condition holds vacuosly.

L3 Since $\leq_F$ is reflexive, transitive and total (and we consider only finite sets of models), any nonempty set has at least one minimal element. Let $\{I\} \in Min(\hat{A} \cup \hat{B}, \leq_F)$: we have either $I \in A$ or $I \in B$. Assume that $I \in A$: we have that for any $J \in A \cup B$ it holds $\{I\} \leq_F \{J\}$. As a result $A \leq_F A \cup B$.

L4 One of the properties of $\preceq_F$ is that $Min(M, \preceq_F) = F$, where M is the set of all the interpretations. If $F \cap B \neq \emptyset$ then any $I \in F \cap B$ is minimal in M, so $\forall J \in C \;\; \exists I \in B$ s.t. $I \preceq_F J$. Conversely, if $F \cap B = \emptyset$, then the models of F are less than the models of B, so $F <_F B$.

L5 Given $\preceq_F$, it is possible that the corresponding $\leq_F$ does not satisfy property *L5*. This is why *L5* has a much longer proof: we must show that it holds directly from *M1*-*M8*, without using the properties of $\preceq_F$.

We will prove first that, if $A \neq \emptyset$

$$A \leq_C B \;\Leftrightarrow\; A \triangle C \subseteq (A \cup B) \triangle C$$

Let us suppose $A \cap C \neq \emptyset$. We have that $A \leq_C B$ by the postulate *L4* proved above. Furthermore $(A \cup B) \cap C \neq \emptyset$, and thus $(A \cup B) \triangle C = (A \cup B) \cap C = (A \cap C) \cup (B \cap C)$, which implies $A \cap C \subseteq (A \cup B) \triangle C$. Since $A \triangle C = A \cap C$, the claim holds.

Let us suppose $A \cap C = \emptyset$, and assume $A \leq_C B$. Since $B - A \subseteq B$ we have $B \leq_C B - A$, and by transitivity $A \leq_C B - A$. Let us assume $B - A \neq \emptyset$ (on the converse, if $B - A = \emptyset$ then $B \subseteq A$ and $(A \cup B) \triangle C = A \triangle C$ and the claim holds trivially).

Now, from postulate *M6* it follows that

$$(A \cup B) \triangle C = [A \cup (B - A)] \triangle C = \begin{cases} A \triangle C & \text{or} \\ (B - A) \triangle C & \text{or} \\ (A \triangle C) \cup [(B - A) \triangle C] & \end{cases}$$

We can prove that the second case (the only one in which $A \triangle C$ is not contained in $(A \cup B) \triangle C$) is impossible. Let us assume $J \in (B - A) \triangle C$ and $J \in (A \cup B) \triangle C$. The second formula is equivalent to $\{J\} \in min(\hat{A} \cup \hat{B}, \leq_C)$. From the assumption $A \leq_C B$ and the definition of $\leq_C$ it follows that for any $J \in B$ there exists a model $I \in A$ such that $I \preceq_C J$. This implies $\{I\} \leq_C \{J\}$ (by definition of $\leq_C$) and thus $\{I\} \in min(\hat{A} \cup \hat{B}, \leq_C)$. As a result, $I \in (A \cup B) \triangle C$. Since I belongs to A but neither to $B - A$ nor to C, it is impossible that $(A \cup B) \triangle C = (B - A) \triangle C$.

Suppose $A \triangle C \subseteq (A \cup B) \triangle C$: we show that $A \leq_C B$. From the hypothesis it follows

$$(A \triangle C) \cap A \subseteq [(A \cup B) \triangle C] \cap (A \cup B)$$

that implies $C *_\triangle A \subseteq C *_\triangle (A \cup B)$. This means that there are models in A that are less than or equal to all the models of B, so $\forall J \in B. \exists I \in A$ s.t. $I \preceq_C J$, that implies $A \leq_C B$.

This concludes the proof that, if $A \neq \emptyset$ then

$$A \leq_C B \;\Leftrightarrow\; A \triangle C \subseteq (A \cup B) \triangle C$$

Notice that $A \cap B \neq \emptyset$ implies $A \leq_{C \cup D} B$ and $C \leq_{A \cup B} D$ and $A \leq_C B$. The same happens if $A \cap D \neq \emptyset$. In both cases *L5* holds trivially. In the sequel we assume that $A \cap (C \cup D) = \emptyset$ and, for similar reasons, $B \cap (C \cup D) = \emptyset$.

We prove now the first implication of *L5*, that is,

$$A \leq_{C \cup D} B \Rightarrow \begin{cases} C \leq_{A \cup B} D \quad \textit{and} \quad A \leq_C B \quad \text{or} \\ D \leq_{A \cup B} C \quad \textit{and} \quad A \leq_D B \end{cases}$$

We can do it in three steps:

1. Prove that $A \leq_{C \cup D} B$ implies either $A \leq_C B$ or $A \leq_D B$.
2. Prove that $A \leq_{C \cup D} B$ and $C <_{A \cup B} D$ imply $A \leq_C B$.
3. Prove that $A \leq_{C \cup D} B$ and $D <_{A \cup B} C$ imply $A \leq_D B$.

We start proving that $A \leq_{C \cup D} B$ implies either $A \leq_C B$ or $A \leq_D B$. Note that $A = \emptyset$ implies $B = \emptyset$ and the property holds. Recall that $A \leq_{C \cup D} B$ is equivalent to (since $A \neq \emptyset$).

$$A \triangle (C \cup D) \subseteq (A \cup B) \triangle (C \cup D)$$

But postulate *M6* implies

$$(A \cup B) \triangle (C \cup D) = \begin{cases} (A \cup B) \triangle C & = \begin{cases} A \triangle C \\ (B - A) \triangle C \\ (A \triangle C) \cup [(B - A) \triangle C] \end{cases} \\ (A \cup B) \triangle D & = \begin{cases} A \triangle D \\ (B - A) \triangle D \\ (A \triangle D) \cup [(B - A) \triangle D] \end{cases} \\ [(A \cup B) \triangle C] \cup [(A \cup B) \triangle D] & = \begin{cases} (A \triangle C) \cup (A \triangle D) \\ ([B - A] \triangle C) \cup (A \triangle D) \\ (A \triangle C) \cup ([B - A] \triangle C) \cup (A \triangle D) \\ (A \triangle C) \cup ([B - A] \triangle D) \\ ([B - A] \triangle C) \cup ([B - A] \triangle D) \\ (A \triangle C) \cup ([B - A] \triangle C) \cup ([B - A] \triangle D) \\ (A \triangle C) \cup (A \triangle D) \cup ([B - A] \triangle D) \\ ([B - A] \triangle C) \cup (A \triangle D) \cup ([B - A] \triangle D) \\ (A \triangle C) \cup ([B - A] \triangle C) \cup (A \triangle D) \cup ([B - A] \triangle D) \end{cases} \end{cases}$$

From postulate *M8* it follows that $A \triangle (C \cup D)$ contains at least a model of A, and hence the same happens to $(A \cup B) \triangle (C \cup D)$. Since neither $(B - A) \triangle C$ nor $(B - A) \triangle D$ contain models of A, it holds either $A \triangle C \subseteq (A \cup B) \triangle C$ or $A \triangle D \subseteq (A \cup B) \triangle D$. In other words, either $A \leq_C B$ or $A \leq_D B$ holds.

Now we show that $A \leq_{C \cup D} B$ and $C <_{A \cup B} D$ imply $A \leq_C B$. Since $A \neq \emptyset$ it follows that

$$A \triangle (C \cup D) \subseteq (A \cup B) \triangle (C \cup D)$$

In a similar manner, from *M6* and the fact that $C <_{A \cup B} D$ it follows

$$(A \cup B) \triangle (C \cup D) = (A \cup B) \triangle C$$

But $[A \triangle (C \cup D)] \cap A \neq \emptyset$ from *M8*, hence $[(A \cup B) \triangle C] \cap A \neq \emptyset$. From *M6* it follows

$$(A \cup B) \triangle C = [A \cup (B - A)] \triangle C = \begin{cases} A \triangle C \\ (B - A) \triangle C \\ (A \triangle C) \cup ((B - A) \triangle C) \end{cases}$$

This implies $A \triangle C \subseteq (A \cup B) \triangle C$ simply because the result contains at least a model of A, while $(B - A) \triangle C$ does not. This completes the proof, since $A \triangle C \subseteq (A \cup B) \triangle C$ is equivalent to $A \leq_C B$.

The third statement, $A \leq_{C \cup D} B$ and $D <_{A \cup B} C$ imply $A \leq_D B$, has a proof similar to that of the second one, where we replace C with D and vice versa.

Now we have to prove that

$$A \leq_{C \cup D} B \Leftarrow \begin{cases} C \leq_{A \cup B} D \quad \textit{and} \quad A \leq_C B \quad \text{or} \\ D \leq_{A \cup B} C \quad \textit{and} \quad A \leq_D B \end{cases}$$

consider the implication

$$C \leq_{A \cup B} D \text{ and } A \leq_C B \quad \Rightarrow \quad A \leq_{C \cup D} B$$

Suppose that $A \cap (C \cup D) \neq \emptyset$. From postulate *L4* proved above it follows that $A \leq_{C \cup D} B$, and the claim holds.

On the converse, suppose that $A \cap (C \cup D) = \emptyset$. Using the property previously proved, the following equivalences hold:

$$\begin{aligned} C \leq_{A \cup B} D \quad &\text{iff} \quad (A \cup B) \triangle C \subseteq (A \cup B) \triangle (C \cup D) \\ A \leq_C B \quad &\text{iff} \quad A \triangle C \subseteq (A \cup B) \triangle C \\ A \leq_{C \cup D} B \quad &\text{iff} \quad A \triangle (C \cup D) \subseteq (A \cup B) \triangle (C \cup D) \end{aligned}$$

(if A or C are empty the hypothesis implies that $D = B = \emptyset$, and the claim holds).

The first two lines imply that $A \triangle C \subseteq (A \cup B) \triangle (C \cup D)$. As a result, $(A \cup B) \triangle (C \cup D)$ contains models of A. Now observe that *M6* implies that

$$(A \cup B) \triangle (C \cup D) = (A \cup (B - A)) \triangle (C \cup D) = \begin{cases} A \triangle (C \cup D) & \text{or} \\ (B - A) \triangle (C \cup D) & \text{or} \\ A \triangle (C \cup D) \cup (B - A) \triangle (C \cup D) \end{cases}$$

Since neither $B - A$ nor $C \cup D$ contains models of A, we have that $A \triangle (C \cup D) \subseteq (A \cup B) \triangle (C \cup D)$, and hence $A \leq_{C \cup D} B$.

■

Now, another question arises: how much information is required in order to determine the behavior of a valid arbitration? In other words, is it necessary to know $A \triangle B$ for any pair of sets A and B ?

The following theorem shows that the knowledge of $X \triangle Y$, where X and Y have cardinality less than or equal to 2, is enough to determine the result of $A \triangle B$ for any A and B.

Theorem 7: Suppose that:

1. $\triangle$ satisfies *M1-M8*
2. We know $X \triangle Y$ for any X, Y such that $|X| \leq 2$ and $|Y| \leq 2$.

Then we can calculate $A \triangle B$ for any pair of sets A, B.

Proof. From the previous theorem it follows that any $\triangle$ satisfying *M1-M8* is represented by the corresponding family of orderings, and any of these orderings $\leq_F$ can be determined by $\{I\} \leq_F \{J\}$ iff $I \in \{I, J\} \triangle F$. Suppose $I, J \notin F$ (on the converse, the result is trivial). Thus

$$\{I, J\} \triangle F = \{Z \mid \{Z\} \in Min(\hat{F}, \leq_{\{I,J\}}) \cup Min(\widehat{\{I, J\}}, \leq_F)\}$$

Let $F = \{F_1, \dots, F_n\}$. We already know $\leq_{\{I,J\}}$ simply because $L \leq_{\{I,J\}} M$ iff $L \in \{I, J\} \triangle \{L, M\}$, thus we are able to determine $Min(\hat{F}, \leq_{\{I,J\}}) = \{\{F_{i_1}\}, \dots, \{F_{i_k}\}\}$. From postulate *M6* we have

$$\{I, J\} \triangle \{F_1, \dots, F_n\} = \begin{cases} \{I, J\} \triangle \{F_1\} \\ \vdots \\ \{I, J\} \triangle \{F_n\} \\ \text{union of the above formulae} \end{cases} \tag{4}$$

But $F_{i_z} \in \{I, J\} \triangle F$ for any $1 \leq z \leq k$: from *M7* it follows that $\{I, J\} \triangle \{F_i\} \subseteq \{I, J\} \triangle F$ for any $1 \leq z \leq k$. On the converse, $F_j \notin \{I, J\} \triangle F$ for any $j \notin \{i_1, \dots, i_k\}$. Since $F_j \in \{I, J\} \triangle \{F_j\}$, it follows that $\{I, J\} \triangle \{F_j\} \nsubseteq \{I, J\} \triangle F$.

Essentially, from $\leq_{\{I,J\}}$ we are able to determine which models of F are in the result (i.e. the value of the indexes i_z). Using 4 and property *M8*, we deduce that $\{I, J\} \triangle F$ is composed exactly by the "sub-arbitrations" $\{I, J\} \triangle F_{i_z}$, where F_{i_z} is one of the models of F that is in the result. Formally,

$$\{I, J\} \triangle F = (\{I, J\} \triangle \{F_{i_1}\}) \cup \dots \cup (\{I, J\} \triangle \{F_{i_k}\})$$

But $\{I, J\} \triangle \{F_i\}$ is one of the known results of $\triangle$, so we can determine whether $I \in \{I, J\} \triangle F$ by verifying if $I \in \{I, J\} \triangle \{F_i\}$ for at least one $F_i \in \{F_1, \dots, F_k\}$.

■

V. Alternative Arbitrations

In the previous sections we introduced a set of postulates for arbitration. Here we investigate on alternative choices for the specific postulates. Two interesting properties for arbitration, not yet considered, are the following ones:

A9' if a implies c and b implies d then $a \triangle b$ implies $c \triangle d$.

A10 $a \triangle (b \triangle c) = (a \triangle b) \triangle c$ if a, b and c are pairwise inconsistent.

Note that *A10* only applies to formulae that are pairwise inconsistent. The corresponding rules in a model-based representation are:

M9' if $A \subseteq C$ and $B \subseteq D$ then $A \triangle B \subseteq C \triangle D$

M10 $A \triangle (B \triangle C) = (A \triangle B) \triangle C$ if A, B and C are pairwise disjoint.

Monotonicity (*M9'*) ensures that arbitration of larger sets obtains as result a larger set, while associativity (*M10*) is a natural property to impose on the arbitration of three or more theories, making the arbitration independent of the order of execution.

A. Monotonicity

While *M9'* seems a very reasonable constraint, it nevertheless has a dramatic impact, as shown by the following negative result:

Theorem 8: There is no arbitration that satisfies *M1-M4* and *M9'*.

Proof. Let A and B be two sets of models such that $A \cap B = \emptyset$. We prove that if one of these has at least two models and arbitration satisfies *M9'*, then $A \triangle B = \emptyset$, so that arbitration violates postulate *M4*. Let $I, J \in A$ be two models of A. If arbitration satisfies *M2* and *M3* then $A \triangle (B \cup \{I\}) = \{I\}$, and $A \triangle (B \cup \{J\}) = \{J\}$. By *M9'* it follows that: $A \triangle B \subseteq A \triangle (B \cup \{I\}) = \{I\}$ and $A \triangle B \subseteq A \triangle (B \cup \{J\}) = \{J\}$. Hence $A \triangle B = \emptyset$. ■

As a consequence, we can assume *M9'* but not *M4*, or assume *M9'* only when $C \cap D = \emptyset$. This leads to a revised form for it:

A9 if $c \wedge d$ is unsatisfiable, a implies c and b implies d then $a \triangle b$ implies $c \triangle d$.

M9 if $C \cap D = \emptyset$, $A \subseteq C$ and $B \subseteq D$ then $A \triangle B \subseteq C \triangle D$

Note that in the proof of Theorem 8 it is crucial that $C \cap D$ is non empty, therefore, the theorem does not apply to this revised version.

While arbitrations satisfying *M1-M8* have an elegant representation theorem, when we assume postulates *M1-M6* and *M9* we can only find two approximating formulae.

Let

$$A \triangle_h B = \begin{cases} A \cap B & \text{if } A \cap B \neq \emptyset \\ \bigcup_{I \in A, J \in B} h(I, J) & \text{otherwise} \end{cases}$$

where h is a function with two interpretations as operands, and a set of interpretations as result. This arbitration is monotone (and satisfies *M6*).

Theorem 9: If $\triangle$ is monotone, then there exist two functions, f and g, such that $A \triangle_f B \subseteq A \triangle B \subseteq A \triangle_g B$ for any sets of models A and B, where:

$$f(I, J) = \{I\} \triangle \{J\}$$

$$g(I, J) = \bigcup_{\{R | I \in R, J \notin R\}} R \triangle \neg R$$

Proof. For any pair of I, J with $I \in A$ and $J \in B$, we have $\{I\} \subseteq A$ and $\{J\} \subseteq B$, and hence by monotonicity

$$\{I\} \triangle \{J\} \subseteq A \triangle B.$$

Hence, we have

$$\cup_{I \in A,\ J \in B} \{I\} \triangle \{J\} \subseteq A \triangle B$$

For the second part, note that for any A and B with $A \cap B = \emptyset$ there exists a set R such that $A \subseteq R$ and $B \subseteq \neg R$. By monotony it follows

$$A \triangle B \subseteq R \triangle \neg R$$

and therefore

$$A \triangle B \subseteq \bigcup_{I \in A,\ J \in B,\ I \in R,\ J \notin R} R \triangle \neg R$$

■

Another monotonic arbitration can be defined as:

$$A \triangle^h B = \begin{cases} A \cap B & \text{if } A \cap B \neq \emptyset \\ \bigcap_{I \notin A, J \notin B} h(I, J) & \text{otherwise} \end{cases}$$

A property of this arbitration is given by the following theorem:

Theorem 10: For any monotonic $\triangle$, there exists a function g s. t.

$$A \triangle B \subseteq A \triangle^g B$$

for any sets of models A and B where g is defined as:

$$g(I, J) = \begin{cases} \top & \text{if } I = J \\ \bigcup_{\{M | I \notin M, J \in M\}} M \triangle \neg M & \text{otherwise} \end{cases}$$

Proof. Since for $I \notin A$, $J \notin B$ and $A \cap B = \emptyset$, there exists an M such that $I \notin M$, $J \in M$ and $A \subseteq M$ and $B \subseteq \neg M$. Hence

$$A \triangle B \subseteq M \triangle \neg M$$

thus

$$A \triangle B \subseteq \bigcup_{\{M | \exists I, J\ .\ I \notin A,\ J \notin B,\ I \in M,\ J \in \neg M\}} M \triangle \neg M$$

■

B. Associativity

In this section, we consider arbitrations satisfying the associativity postulate. We show that the only associative arbitration satisfying condition *M7* and *M8* is $A \triangle B = A \cup B$.

Theorem 11: If $\triangle$ satisfies *M1*, *M7*, *M8* and *M10* then $\triangle = \cup$.

Proof. First of all, we show that the repeated arbitration of singletons becomes the union:

$$\{I_1\} \triangle \cdots \triangle \{I_n\} = \{I_1, \ldots, I_n\}$$

By postulate *M7* it immediately follows that $\{I_1\} \triangle \cdots \triangle \{I_n\} \subseteq \{I_1, \ldots, I_n\}$. We now show that for all elements $I_j \in \{I_1, \ldots, I_n\}$ we have that $I_j \in \{I_1\} \triangle \cdots \triangle \{I_n\}$. By associativity and commutativity of the arbitration operator we can rewrite $\{I_1\} \triangle \cdots \triangle \{I_n\}$ as $\{I_j\} \triangle (\{I_1\} \triangle \cdots \triangle \{I_{j-1}\} \triangle \{I_{j+1}\} \cdots \triangle \{I_n\})$. Therefore, postulate *M8* implies that $I_j \in \{I_1\} \triangle \cdots \triangle \{I_n\}$. Thus, we obtain that $\{I_1\} \triangle \cdots \triangle \{I_n\} = \{I_1, \ldots, I_n\}$.

The result now easily follows, in fact $A \triangle B$ can be rewritten as the arbitration of the union of all interpretations of A with the union of all interpretations of B: $\{A_1, \ldots, A_n\} \triangle \{B_1, \ldots, B_m\}$. With simple rewritings we obtain:

$$(\{A_1\} \cup \cdots \cup \{A_n\}) \triangle (\{B_1\} \cup \cdots \cup \{B_m\}) =$$
$$= \{A_1\} \triangle \cdots \triangle \{A_n\} \triangle \{B_1\} \triangle \cdots \triangle \{B_m\}$$

Since $\{A_i\}$ and $\{B_i\}$ are singletons, it follows that $\{A_1\} \triangle \cdots \triangle \{A_n\} \triangle \{B_1\} \triangle \cdots \triangle \{B_m\} = A \cup B$.

■

Since $\cup$ does not satisfies postulate *M3*, the following theorem holds as well.

Theorem 12: No arbitration can satisfy *M1-M8* and *M10* at the same time.

Another interesting property of associative arbitrations satisfying *M1-M7* concerns the arbitration of singletons (sets with only one element).

Theorem 13: Let $\leq$ be the relation such that $I \leq J$ iff $I \in \{I\} \triangle \{J\}$. As a consequence, whenever $\triangle$ satisfies *M1-M7* and *M10*, the relation $\leq$ is transitive (and, obviously, reflexive and total).

Proof. Suppose $I \leq J$ and $J \leq Z$: these expressions are equivalent to $I \in \{I\} \triangle \{J\}$ and $J \in \{J\} \triangle \{Z\}$. From associativity it follows that

$$\begin{aligned} \{I\} \triangle (\{J\} \triangle \{Z\}) &= (\{I\} \triangle \{J\}) \triangle \{Z\} \\ &= (\{J\} \triangle \{I\}) \triangle \{Z\} \\ &= \{J\} \triangle (\{I\} \triangle \{Z\}) \end{aligned}$$

But $I \in \{I\} \triangle (\{J\} \triangle \{Z\})$, so $I \in \{J\} \triangle (\{I\} \triangle \{Z\})$. This is possible only if $I \in \{I\} \triangle \{Z\}$, thus $I \leq Z$. ■

An example of associative arbitration is:

$$A \triangle B = \begin{cases} A \cap B & \text{if } A \cap B \neq \emptyset \\ Min(A \cup B, \leq) & \text{otherwise} \end{cases}$$

where $\leq$ is a reflexive, transitive, and total order on interpretations.

C. Distributivity

Postulate *A6* enforces a weak form of distributivity of union wrt arbitration. If we fully enforce distributivity (i.e. $A \triangle (B \cup C) = (A \triangle B) \cup (A \triangle C)$) then arbitration collapses to disjunction, that is $A \triangle B \equiv A \cup B$.

Theorem 14: A fully distributive arbitration satisfying also *A78* is equivalent to the union.

Proof. Let $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_m\}$. We have

$$\begin{aligned} A \triangle B &= (A \triangle \{B_1\}) \cup \cdots \cup (A \triangle \{B_m\}) \\ &= \bigcup_{i=1..n \quad j=1..m} \{A_i\} \triangle \{B_j\} = A \cup B \end{aligned}$$

■

VI. Arbitration Operators and Examples

In this section we introduce actual arbitration operators, check whether they satisfy the postulates and show their behavior on examples. In order to simplify definitions and proofs, we assume in this section that the two revising formulae a and b are satisfiable and therefore, that the corresponding sets of models A and B are non empty. An arbitration operator can be defined as:

$$A \triangle_{\leq} B = \{I, J | \langle I, J \rangle \in Min(A \times B, \leq)\}$$

where $\leq$ is a reflexive, transitive and total relation over *ordered* pairs of interpretations, such that $\langle I, J \rangle \cong \langle J, I \rangle$ for any I and J.

The intuitive meaning of $\leq$ is the measure of closeness (or distance) between models. The ordering over pairs of models formalize our valuation of the distance: $\langle I, J \rangle \leq \langle L, M \rangle$ means that I and J are closer than L and M, or that the distance between I and J is less than the distance between L and M. This ordering depends on our valuation of the distance. However, it is clear that a model is the closest to itself. Hence, the pairs $\langle I, I \rangle$ are the minimal elements of the $\leq$ ordering. We remind that the notation $x \cong y$ is a shorthand for $x \leq y$ and $y \leq x$, while $x < y$ is a shorthand for $x \leq y$ and not $y \leq x$.

Theorem 15: If $\leq$ is reflexive, transitive, total and satisfies the condition $\langle J, J \rangle \cong \langle I, I \rangle$ and $\langle I, I \rangle < \langle I, J \rangle$ for every I, J (if $I \neq J$) then $\triangle_{\leq}$ satisfies *M1-M8*.

Proof. Property *M1* is trivially satisfied. For *M2* and *M3*, note that $A \cap B \neq \emptyset$ implies that given $I \in A \cap B$, the pair $\langle I, I \rangle \in A \times B$ is strictly less than any other pair $\langle I, J \rangle \in A \times B$ with $I \neq J$ but equal to any other $\langle Z, Z \rangle \in A \times B$. Thus, $A \triangle_{\leq} B = \{I, J | \langle I, J \rangle \in Min(A \times B, \leq)\} = A \cap B$.

Property *M4* follows from the assumption that both A and B are non empty and therefore, $A \times B$ is non-empty too. Since $\leq$ is reflexive and transitive, any finite set has at least one minimal element.

Property *M6* follows from the transitivity of $\leq$: since $A \times (B \cup C) = (A \times B) \cup (A \times C)$, it follows that $Min(A \times (B \cup C), \leq)$ is equal either to $Min(A \times B, \leq)$ or $Min(A \times C, \leq)$ or the union of these two sets.

Postulate *M7* holds because $A \triangle_{\leq} B = \cup\{\{I,J\} | \langle I,J\rangle \in Min(A \times B, \leq) \subset (A \cup B)$. Since $\langle I,J\rangle \in Min(A \times B, \leq)$ contains pairs of elements, where the first element is a model of A and the second is a model of B, if $A \neq \emptyset$ then $A \cap (A \triangle B) \neq \emptyset$, thus postulate *M8* holds.

■

This is a very simple definition, but, unfortunately, there are arbitrations satisfying *M1*-*M8* such that no relation $\leq$ can represent them.

Since the intuitive meaning of $\leq$ is the measure of closeness (or distance) between models, when arbitrating between two formulae with sets of models A and B, it is natural to take the closest pairs of models. Incidentally, we note that the revision operators of Borgida [13], Dalal [8], Forbus [14] and Satoh [9] can be represented in a very similar fashion:

$$Mod(k * a) = \begin{cases} Mod(k \wedge a) & \text{if non empty} \\ \{J | \exists I. \langle I,J\rangle \in Min(K \times A, \leq)\} & \text{otherwise} \end{cases}$$

where $\leq$ is defined as follows:

Dalal	$\langle I,J\rangle \leq_{Dalal} \langle L,M\rangle$	iff	$\lvert diff(I,J)\rvert \leq \lvert diff(L,M)\rvert$
Borgida	$\langle I,J\rangle \leq_{Borgida} \langle L,M\rangle$	iff	$I = L$ and $diff(I,J) \subseteq diff(L,M)$
Satoh	$\langle I,J\rangle \leq_{Satoh} \langle L,M\rangle$	iff	$diff(I,J) \subseteq diff(L,M)$
Forbus	$\langle I,J\rangle \leq_{Forbus} \langle L,M\rangle$	iff	$I = L$ and $\lvert diff(I,J)\rvert \leq \lvert diff(L,M)\rvert$

Note that only Dalal's relation is total. The others are reflexive and transitive but not total.

We can give an example of arbitration choosing one of these relation. Forbus's and Borgida's ones are relations over *ordered* pairs of models, hence they aren't suitable for our application. Since only Dalal's relation is total it seems the better choice. However, we can also take Satoh's relation, but since it is not total, the corresponding arbitration does not satisfy *M6*. The corresponding arbitration operators are defined as follows:

$$\begin{aligned} A \triangle_D B &= \{I,J | \langle I,J\rangle \in Min(A \times B, \leq_{Dalal})\} \\ A \triangle_S B &= \{I,J | \langle I,J\rangle \in Min(A \times B, \leq_{Satoh})\} \end{aligned}$$

From Theorem 15 it follows that Dalal-like arbitration ($\triangle_D$) satisfies postulates *M1*-*M8* (under the assumption that both A and B are non-empty). In this arbitration, distance between interpretations I and J is the cardinality of $diff(I,J)$, that is the number of literals that are assigned different values.

Consider again the running example of the stock exchange. There are two pairs of closest models, that is $I_1 = \{x,y\} \in Mod(john)$ and $J_1 = \{y\} \in Mod(robert)$ and $I_2 = \{x,y,z\} \in Mod(john)$ and $J_2 = \{y,z\} \in$

$Mod(robert)$ whose distance is 1. It follows that $Mod(john) \triangle_D Mod(robert) = \{(x, y), (y), (x, y, z), (y, z)\}$ and therefore it implies y. Notice that y is not a consequence of $john \vee robert$.

Satoh's relation is not total, hence the corresponding arbitration ($\triangle_S$) does not satisfy *M6*. However, it satisfies the others (*M4* only under the assumption that both A and B are non-empty).

These definitions of arbitration are very close to the corresponding revisions. In fact, it is possible to transform Dalal's (Satoh's) revision into Dalal-like (Satoh-like) arbitration and vice versa, using the following relations:

$$k * a = (k \triangle a) \wedge a$$

$$a \triangle b = (a * b) \vee (b * a)$$

All the arbitrations considered above are operators over pairs of sets of models. Some authors feel that in some cases a revision that is insensitive on the syntax of the knowledge base is inappropriate.

Indeed, as already discussed in Section II, operators of this kind depend on the form in which the bases are expressed and thus they may not obey the principle of irrelevance of the syntax (formulated, for example, in [8]): two equivalent bases may lead to a different dynamic behavior (that is, the result of revision may depend on the syntax of the operands).

Given a set of formulae c, the set $W(c, \mathsf{true})$ is the set of the maximal consistent subsets of c

$$W(c, \mathsf{true}) = \max(\{d \subseteq c \mid d \text{ is consistent}\}, \subseteq)$$

Any element of $W(a \cup b, \mathsf{true})$ can be viewed as an arbitration that preserves as much as possible of the information contained in a and b.

However, without more specific knowledge, we have no criterium to choose one of these maximal sets. We introduce two arbitration operators based on the belief revision operators $*_{FUV}$ and $*_W$ introduced in the preliminaries. The first one considers all maximal subsets of $a \cup b$:

$$a \triangle_{FUV} b = W(a \cup b, \mathsf{true})$$

while the second one is more drastic and takes only the formulae that are in the intersection of all these sets:

$$a \triangle_W b = \bigcap W(a \cup b, \mathsf{true})$$

These arbitrations can be reduced to their corresponding revisions via

$$\begin{aligned} \{f_1, \ldots, f_m\} * P &= \{f_1 \wedge P, \ldots, f_m \wedge P\} \triangle \{P\} \\ \{a_1, \ldots, a_m\} \triangle \{b_1, \ldots, b_r\} &= \{a_1, \ldots, a_m, b_1, \ldots, b_r\} * \{\mathsf{true}\} \end{aligned}$$

Note that $\triangle_{FUV}$ satisfies postulates *A1*-*A4*, *A7* and *A8* but it violates *A5* (syntax-independence) and *A6*, while $\triangle_W$ satisfies postulates *A1*-*A4* and *A8* but it violates *A5* (syntax-independence), *A6* and *A7*.

We illustrate through an example the different behavior of the four arbitration operators introduced.

Example 4: Let $a = \{w, x \rightarrow (y \wedge z), y \rightarrow (x \wedge z), z \rightarrow (x \wedge y)\}$ and $b = \{\neg x \wedge y \wedge z\}$. Note that a has two models $I_1 = \{w\}$ and $I_2 = \{x, y, z, w\}$, while b has models $J_1 = \{y, z\}$ and $J_2 = \{y, z, w\}$. As a consequence, $a \triangle_D b$ has models I_2 and J_2, while $a \triangle_S b$ has models I_1, I_2 and J_2.

Coming to the formula-based operators, the set of maximal consistent sets is:

$$W(a \cup b, \mathsf{true}) = \{\{w, x \rightarrow (y \wedge z), y \rightarrow (x \wedge z), z \rightarrow (x \wedge y)\}, \{w, \neg x \wedge y \wedge z, x \rightarrow (y \wedge z)\}\}$$

Therefore $a \triangle_{FUV} b = \{\{w, x \rightarrow (y \wedge z), y \rightarrow (x \wedge z), z \rightarrow (x \wedge y)\}, \{w, \neg x \wedge y \wedge z, x \rightarrow (y \wedge z)\}\}$, while $a \triangle_W b = \{w, x \rightarrow (y \wedge z)\}$. $\diamond$

A. Complexity

We briefly address the computational complexity of deciding the following problem: Given three propositional formulae a, b and c, what is the complexity of deciding whether $a \triangle b \models c$ holds.

We first prove a general result that holds for all arbitration operators satisfying postulates *A1-A4*.

Theorem 16: If $\triangle$ satisfies *A1-A4*, then the problem of deciding whether $a \triangle b \models c$ holds, is both NP and coNP-hard.

Proof. Let a be a propositional formula. We prove NP-hardness by showing a reduction from SAT. We have that $\mathsf{true} \triangle a \models a$ iff a is satisfiable. In fact, if a is satisfiable then $\mathsf{true} \triangle a = a$ and, obviously, it implies a. If a is not satisfiable nevertheless $\mathsf{true} \triangle a$ is satisfiable because of *A4*. Since a satisfiable formula does not imply an unsatisfiable one we have that $\mathsf{true} \triangle a \not\models a$.

We prove coNP-hardness by showing a reduction from the problem of deciding whether a propositional formula is valid. We have that a is valid iff $\mathsf{true} \triangle \mathsf{true} \models a$. In fact, $\mathsf{true} \triangle \mathsf{true} = \mathsf{true}$ and $\mathsf{true} \models a$ if and only if a is valid. ■

This is a result that holds for all revision operators, we now focus on the four specific ones defined above. To obtain the complexity results we use the reductions to belief revision outlined above and the complexity results for belief revision operators presented by Eiter and Gottlob in [12]. In fact, the arbitration operator has exactly the same complexity of the corresponding belief revision one.

Proposition 17: Let a, b and c be propositional formulae, then the problem of deciding whether $a \triangle b \models c$ holds, is $P^{NP[O(log\ n)]}$-complete for $\triangle = \triangle_D$, Π_2^p-complete for $\triangle = \{\triangle_S, \triangle_{FUV}\}$ and Π_2^p-hard and in $P^{\Sigma_2^p[O(log\ n)]}$ for $\triangle = \triangle_W$.

VII. Related Work

As already pointed out in the introduction, Revesz was the first author to introduce the idea of arbitration and formally define it. We want now to compare our proposal with the one he put forward in [10].

First of all, Revesz did not provide a set of postulates for arbitration, but he introduces the model-fitting operator ($\triangleright$) and then defines arbitration as $a \triangle b \doteq (a \vee b) \triangleright \mathsf{true}$. This definition only depends on the disjunction of a and b, therefore, there is no way to distinguish a's contribution to the result from b's one.

Considering all our postulates, it is easy to show that Revesz's arbitration operators only satisfy the basic property of commutativity, represented by postulate *A1*, consistency preservation (*A4*) and syntax-independence (*A5*). More precisely:

Theorem 18: Any model-fitting operator $\triangleright$ satisfying postulates *R1-R8* defines an arbitration operator $a \triangle b \doteq (a \vee b) \triangleright \mathsf{true}$ that satisfies postulates *A1*, *A4* and *A5*.

Proof. By definition, $a \triangle b = (a \vee b) \triangleright \mathsf{true}$ while $b \triangle a = (b \vee a) \triangleright \mathsf{true}$. Since disjunction is commutative it follows that $a \triangle b = b \triangle a$. Since $\triangleright$ satisfies postulate *R2* we have that if $a \vee b$ is unsatisfiable so is $a \triangle b$, therefore if a and b are unsatisfiable so is $a \triangle b$. Furthermore, by postulate *R3* it follows that if $a \vee b$ and true are satisfiable so is $a \triangle b$. Since true is always satisfiable we have that postulate *A4* holds. Postulate *A5* holds because model fitting satisfies syntax-independence (*R4*). ■

On the other hand, all other postulates are not satisfied by all arbitration operators as defined by Revesz.

Theorem 19: There exist model-fitting operators $\triangleright$ satisfying postulates *R1-R8*, defining an arbitration operator $a \triangle b \doteq (a \vee b) \triangleright \mathsf{true}$ that does not satisfy postulates *A2*, *A3*, *A6*, *A7* and *A8*.

Proof. Let $a = (\neg x \wedge \neg y)$ and $b = (x \wedge y) \vee (\neg x \wedge \neg y)$. Using the arbitration operator $\triangle_R$ defined by Equation 3 in Section II-C, it turns out that $a \triangle_R b = (\neg x \wedge y) \vee (x \wedge \neg y)$. Therefore, since $(\neg x \wedge y) \vee (x \wedge \neg y)$ is not a consequence of $a \wedge b$ the operator $\triangle_R$ does not satisfy *A2*. Now notice that $a \wedge b$ is satisfiable but $a \triangle_R b$ does not imply $a \wedge b$. Thus $\triangle_R$ does not satisfy *A3* either. Furthermore, $a \triangle_R b$ does not imply $a \vee b$ and, even though a is satisfiable $a \wedge (a \triangle_R b)$ is unsatisfiable. Thus violating *A7* and *A8*.

Now let $a = (\neg x \wedge \neg y)$, $b = (x \wedge y)$ and $c = (x \wedge \neg y)$. We have that $a \triangle_R (b \vee c) = (x \wedge \neg y)$, while $a \triangle_R b = (x \wedge \neg y) \vee (\neg x \wedge y)$ and $a \triangle_R c = x$. Therefore, $a \triangle_R (b \vee c)$ is different from $a \triangle_R b$, $a \triangle_R c$ and $(a \triangle_R b) \vee (a \triangle_R c)$. Thus violating *A6*. ■

However, it is possible to slightly modify Revesz's definition in order to satisfy *A2* and *A3*.

$$a \triangle_n b = \begin{cases} a \wedge b & \text{if } a \wedge b \text{ is satisfiable} \\ \text{Revesz's definition} & \text{otherwise} \end{cases}$$

Nevertheless, even $\triangle_n$ does not satisfy *A6*, *A7*, *A8* monotonicity (*M9*) and associativity (*M10*).

Theorem 20: There exist model-fitting operators $\triangleright$ satisfying postulates *R1-R8*, defining an arbitration operator $a \triangle_n b$ that does not satisfy postulates *A6*, *A7*, *A8*, monotonicity (*A9*) and associativity (*A10*).

Proof. Violation of *A6* is proven in Theorem 19. Let $a = (\neg x \wedge \neg y)$ and $b = (x \wedge y)$. Using the arbitration operator $\triangle_R$, it turns out that $a \triangle_R b = (\neg x \wedge y) \vee (x \wedge \neg y)$. Therefore, $a \triangle_R b$ does not imply $a \vee b$ and, even though a is satisfiable $a \wedge (a \triangle_R b)$ is unsatisfiable. Thus violating *A7* and *A8*.

Let $a = (\neg x \wedge \neg y)$, $b = (x \wedge y)$ and $c = (\neg x \wedge \neg y) \vee (x \wedge \neg y)$. Arbitrating Between a, b and c, b we obtain

$a\triangle_R b = (\neg x \wedge y) \vee (x \wedge \neg y)$ and $c\triangle_R b = x$. Even though $a \wedge b$ is unsatisfiable and a implies c, $a\triangle_R b$ does not imply $c\triangle_R b$, thus violating *A9*.

Let $a = x \wedge y$, $b = \neg x \wedge y$ and $c = (x \wedge \neg y) \vee (\neg x \wedge \neg y)$. Using the arbitration operator $\triangle_R$ presented in Section II-C, it turns out that $(a\triangle_R b)\triangle_R c = (x \wedge y) \vee (x \wedge \neg y) \vee (\neg x \wedge y) \vee (\neg x \wedge \neg y)$ while $a\triangle_R(b\triangle_R c) = (x \wedge \neg y) \vee (\neg x \wedge y)$. Therefore, *A10* is not satisfied. ■

Summing up our comparison with Revesz's work, we believe that arbitration is a primitive notion and that it is better formalized by directly giving properties for it, rather than defining it as a derived notion of model-fitting. To back up our claim we showed that Revesz's arbitration does not satisfy many properties that, in our opinion, are crucial.

Other authors have considered the issue of amalgamating possibly inconsistent databases. In fact, as already pointed out in the introduction, this problem has been mentioned by Silberschatz, Stonebraker and Ullman as a very important challenge for the designers of the next generation of databases systems.

The most interesting proposal is the one presented by Subrahmanian in [15]. In this work he defines a model where (possibly mutually contradictory) databases are integrated via a *supervisory* database, that handles the resolution of possible conflicts arising from contradictory information contained in the various databases. Knowledge in the databases is expressed using a multivalued logic called "annotated logic". Our approach is not comparable with Subrahmanian's one because goals and scenarios of applications are distinct. Subrahmanian's approach is best suited to situations where the original databases cannot (or should not) be modified. Therefore, he defines a supervisor that makes us feel that we have a consistent set of information, by mediating and choosing one info over the other when contradiction arises. This approach implicitly assumes the existence of a conflict resolution strategy that can help us in making sense of contradictory information.

In our approach we want to generate a new database which merges the original ones retaining as much information as possible while preserving consistency. Furthermore, we do not assume the existence of a specific conflict resolution strategy which enables us to solve inconsistency.

VIII. Concluding Remarks

In this paper we presented a new approach to the problem of combining possibly inconsistent information coming from different sources each of them equally reliable. We have defined new principles that guide the process of merging (possibly contradictory) knowledge bases into a new and consistent one. The proposed operator of arbitration captures some aspects of the belief revision operators as defined by the AGM postulates, but it also has some very distinct features. Our notion of arbitration is then compared with the notion presented by Revesz in [10].

Some issues deserve further investigation. Among the most prominent ones there are:

1. In this work we assume that all the information contained in the two knowledge bases has the same

importance. It might well be the case that some sentences (or models) are more important than others. Our framework should be generalized to be able to take these differences (if present) into account.

2. Subrahmanian in [15] requires the definition of conflict resolution strategies to solve contradictions. We are presently investigating how our system can be generalized so that it can represent and use conflict resolution strategies, if present.

Acknowledgments

The authors thank the reviewers who gave important advice on how to improve the paper. We also thank Peter Revesz for comments on a previous draft of this paper.

References

[1] P. Liberatore and M. Schaerf, "Arbitration: A commutative operator for belief revision", in *Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence (WOCFAI-95)*, 1995, pp. 217–228.

[2] A. Silberschatz, M. Stonebraker, and J. D. Ullman, "Database systems: Achievements and opportunities", *Communications of the ACM*, vol. 34, no. 10, pp. 110–120, 1991.

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

[4] P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press, Cambridge, MA, 1988.

[5] H. Katsuno and A. O. Mendelzon, "On the difference between updating a knowledge base and revising it", in *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91)*, 1991, pp. 387–394.

[6] R. Fagin, J. D. Ullman, and M. Y. Vardi, "On the semantics of updates in databases", in *Proceedings of the Second ACM SIGACT SIGMOD Symposium on Principles of Database Systems (PODS-83)*, 1983, pp. 352–365.

[7] M. L. Ginsberg, "Conterfactuals", *Artificial Intelligence Journal*, vol. 30, pp. 35–79, 1986.

[8] M. Dalal, "Investigations into a theory of knowledge base revision: Preliminary report", in *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, 1988, pp. 475–479.

[9] K. Satoh, "Nonmonotonic reasoning by minimal belief revision", in *Proceedings of the International Conference on Fifth Generation Computer Systems (FGCS-88)*, 1988, pp. 455–462.

[10] P. Z. Revesz, "On the semantics of theory change: Arbitration between old and new information", in *Proceedings of the Twelfth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS-93)*, 1993, pp. 71–82.

[11] H. Katsuno and A. O. Mendelzon, "Propositional knowledge base revision and minimal change", *Artificial Intelligence Journal*, vol. 52, pp. 263–294, 1991.

[12] T. Eiter and G. Gottlob, "On the complexity of propositional knowledge base revision, updates and conterfactuals", *Artificial Intelligence Journal*, vol. 57, pp. 227–270, 1992.

[13] A. Borgida, "Language features for flexible handling of exceptions in information systems", *ACM Transactions on Database Systems*, vol. 10, pp. 563–603, 1985.

[14] K. D. Forbus, "Introducing actions into qualitative simulation", in *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, 1989, pp. 1273–1278.

[15] V. S. Subrahmanian, "Amalgamating knowledge bases", *ACM Transactions on Database Systems*, vol. 19, no. 2, pp. 291–331, 1994.