On the Semantics of Arbitration*

Peter Z. Revesz
Department of Computer Science and Engineering
University of Nebraska–Lincoln

Abstract: Revision and update operators add new information to some old information represented by a logical theory. Katsuno and Mendelzon show that both revision and update operators can be characterized as accomplishing a minimal change in the old information to accommodate the new information. Arbitration operators add two or more weighted informations together where the weights indicate the relative importance of the informations rather than a strict priority. This paper shows that arbitration operators can be also characterized as accomplishing a minimal change. The operator of model-fitting is also defined and analyzed in the paper.

1 Introduction

Arbitration is the process of settling a conflict between two or more persons. Arbitration occurs in many situations. For example, settling a labor dispute by an outsider, reaching a verdict in a trial, evaluating several alternative research hypotheses, negotiating an international peace agreement, or setting the price of a product in a competitive market, all can be viewed as cases of arbitration.

Arbitration is more general than selection and should not be confused with it. Selection means siding with one person in a conflict, while arbitration may mean siding with one person in some issues and another person in other issues. Arbitration yields a settlement that best satisfies several people's conflicting interests, subject to certain rules.

Arbitration is often done by one or more impartial persons, the arbitrators, but sometimes there is no clear arbitrator. For example, a product's price in the free market is settled by a process, not by well-defined arbitrators.

Arbitration is worthy of study on its own. This paper introduces a logical framework for the study of arbitration. This framework is a step towards making arbitration amenable to computer solutions. We will describe several cases of arbitration by appropriate sets of axioms and give sample arbitration operators as well.

The set of knowledge or belief that we have can be formally represented by a knowledge base K that consists of a set of logical formulas. If the set of formulas is deductively closed we also call the knowledge base a theory. Each formula is meaningfully kept separate within the knowledge base. For example, each may denote a different witness' testimony, the opinions of different newspaper editors, or the results of different scientific experiments. Knowledge bases can be constructed by successive set additions. For example, as each witness tells his or her story during a trial, the court clerk records the testimony. This can be represented by adding to the current knowledge base a formula describing the witness' testimony. The final knowledge base will be a set of formulas, with one formula for each distinct testimony. Using sets instead of ordered sequences as in [Ry91] reflects the intuition that the order of the witnesses' testimonies should not matter.

*Manuscript received November 2, 1993, revised September 5, 1995. Part of the research presented here was done while the author was visiting the University of Toronto. This research was supported by a grant from the Canadian Institute for Robotics and Intelligent Systems and an Oliver E. Bird Faculty Fellowship from the University of Nebraska-Lincoln. A preliminary version of this paper appeared in **[Rev93]**. For correspondence write to `revesz@cse.unl.edu`.

Arbitration may be used to test against the knowledge base several possible hypotheses. For example, a jury may test several hypothetical reconstructions of the crime based on a knowledge base containing the witnesses' testimonies. Similarly, a scientist may test several research hypotheses based on a knowledge base containing the results of different experiments. In neither case does arbitration change the knowledge base. Normally, the jury will not tamper with the records of the testimonies, nor will the scientist falsify the data. Arbitration in these cases resembles hypothetical querying [Gab85, Gin86, Bon90, Mey90, Gra91, GM91].

We consider two ways of restricting the set of hypotheses from which the arbitrators have to choose. In the first case the hypotheses are restricted by a single formula μ. For example, for a jury, μ may be a conjunction of two facts, one of which may state that no person can be at two places at the same time and the other may state that the heights and the fingerprints of adults do not change. In this case, the jury's verdict must be consistent with μ, that is, each hypothetical reconstruction of the crime will be a model of μ. Intuitively, the model that best fits the entire knowledge base is the one on which the jury is likely to reach a consensus. We call this special case of arbitration *model-fitting*. In the second case the hypotheses are restricted by a set of formulas. The arbitrators have to select from the set those formulas that best fit the whole knowledge base. For example, the jury may have to choose between the prosecution's argument and the defense's argument, both of which may be represented by complex formulas.

Arbitration is related to revision and update, which are fundamental concerns to databases [BS81, FUV83], to Artificial Intelligence [McC68, Rei92], and to belief revision [Mak85, Gär88]. Both revision and update are processes that change a knowledge base or theory when new information is presented that conflicts with it. As pointed out by Abiteboul and Grahne [AG85], by Keller and Winslett [KW85] and by Katsuno and Mendelzon [KM89], the nature of the new information is important in distinguishing between revision and update. In update the new information simply says that something we hold true is not true any more at the present time, while in revision the new information says that something we hold true was never true. In either case the new information is accommodated within the knowledge base but in different ways.

Several authors propose concrete operators for revision and update including Borgida [Bor85], Dalal [Dal88], Fagin, Kuper, Ullman and Vardi [FUV83, FKUV86], Grahne, Mendelzon and Revesz [GMR92], Hansson [Han91], Nebel [Neb90], Ryan [Ry91], Satoh [Sat88], Weber [Web86] and Winslett [Win88]. Instead of proposing a concrete operator Alchourrón, Gärdenfors and Makinson [AGM85, Gär88] give a set of axioms or postulates that every adequate revision operator should be expected to satisfy. Similarly, Katsuno and Mendelzon [KM92] present axioms for updates.

Arbitration may be used also as a knowledge base change operator [Rev93]. Suppose the old information is a knowledge base K and the new information is a sentence ϕ. Both revision and update resolve possible conflicts between the two by giving priority to the new information. However, if neither the old nor the new information is preferred, then arbitration over the knowledge base $K \cup \phi$ can resolve the conflicts. In this case the result of the arbitration may be preserved as the new knowledge base. Arbitration as a knowledge base change operator is useful for heterogeneous databases, which often require merging of large equally important sets of information to answer queries. This is also the motivation of the work on combination by Baral, Kraus, Minker and Subrahmanian [BKM91, BKMS92].

We believe that the three types of knowledge base operators complement each other. They can be even used alternately during a complex application. For example, consider again the example of the jury trial. If during a cross examination a witness changes his or her testimony (eg. "I was mistaken, the color of the car was probably red instead of blue"), then the testimony of this witness should be revised, while not changing any of the other testimonies. Later on we may use update as well, for example when a witness says "I was fired from my job after my testimony here last week". Then of course we may use arbitration in the end for finding the jury's verdict.

It is obvious that there are often significant differences in the trustworthiness or the importance of different sources of information. Therefore we are also going to consider a generalization of arbitration where the formulas within a knowledge base can have different weights. The weights we use are distinct from the "weights" of Fagin et al. [FUV83], which denote priority values, that is a strict ordering which is followed while trying to satisfy the most number of formulas. Our intuition for using weights is quite different. A formula within a knowledge base may have a large weight, but its consideration may be defeated by several formulas with lower weights. The weights in this paper are also distinct from the weights of Dalal [Dal88] which are assigned to propositional terms rather than to the set of models of formulas. They are also different from the possibility values of Zadeh [Zad78] which are restricted to range between zero and one instead of being arbitrary nonnegative real numbers.

Katsuno and Mendelzon [KM89, KM92] found an elegant model-theoretic characterization of revision and update when the knowledge base is a propositional theory T. They found that revision operators that satisfy the AGM postulates are exactly those that select from the models of the new information the closest models to any model of T. Update operators select for each model I of T the models of the new information that are closest to I. The new theory is the union of all such models. Grove [Gro88] also gives a characterization in terms of a system of spheres for revision operators that works for first-order knowledge bases as well. Analogously to these results model-fitting operators can be characterized as those operators that select from the models of an integrity constraint the overall closest models to the whole set of formulas of a knowledge base K.

The outline of the paper is the following. Section 2 lists some basic definitions in the case when the knowledge base is a set of propositional formulas. Section 3 defines by postulates the operation of model-fitting. Section 4 considers the generalization of model-fitting by adding weights. Section 5 defines arbitration by another set of postulates. Each of Sections 3 to 5 presents a model-theoretic characterization of the operator defined. Sections 3 to 5 consider only propositional knowledge bases. Section 6 discusses the case of arbitration of first-order knowledge bases. Section 7 compares arbitration operators with the decision making protocols of Borgida and Imielinski [BI84] and the combination operator of Baral et al. [BKM91, BKMS92]. Finally Section 8 lists some open problems.

2 Preliminaries

Let $\mathcal{T}$ be a finite set of propositional *terms*. We build propositional *formulas* from terms using the unary connective $\neg$ denoting *boolean negation*, and the binary connectives $\wedge$ and $\vee$ denoting *boolean and* and *boolean or*. We call each $I \subset \mathcal{T}$ an interpretation. Let $\mathcal{M}$ be the set of interpretations $\{I : I \subset \mathcal{T}\}$. The set of *models* of a formula ϕ denoted by $Mod(\phi)$ is defined as follows:

$$
\begin{aligned}
Mod(t) &= \{I \in \mathcal{M} : t \in I\} \\
Mod(\neg\phi) &= \mathcal{M} \setminus Mod(\phi) \\
Mod(\psi \vee \phi) &= Mod(\psi) \cup Mod(\phi) \\
Mod(\psi \wedge \phi) &= Mod(\psi) \cap Mod(\phi)
\end{aligned}
$$

In this paper we will use the expression $form(I_1, \ldots, I_k)$ to denote the formula that has exactly the models $I_1, \ldots, I_k$.

A knowledge base K is a set of formulas. A theory is a deductively closed set of formulas. If we have a consequence relation cn and K is any knowledge base, then $cn(K)$ is a theory. Let $\perp$ denote falsity, that is the formula with no models. We say that a theory T is consistent if and only if $\perp \notin T$. A knowledge base K is consistent if and only if the theory $cn(K)$ is consistent.

Gärdenfors considered the problem of theory change from an axiomatic point of view. In particular, Gärdenfors described the following axioms for revising a consistent theory.

Let $T^{+}\mu$ denote the smallest deductively closed set containing T and μ, and let $T_{\perp}$ denote the set of all formulas. If $*$ is any theory change operator, then for any consistent theory T and formulas μ and ϕ the following should hold:

(G1) $T^*\mu$ is a theory.

(G2) $\mu \in T^*\mu$.

(G3) $T^*\mu \subseteq T^+\mu$.

(G4) If $\neg\mu \notin T$, then $T^+\mu \subseteq T^*\mu$.

(G5) $T^*\mu = T_\perp$ only if μ is unsatisfiable.

(G6) If $\mu \equiv \phi$ then $T^*\mu = T^*\phi$.

(G7) $T^*(\mu \wedge \phi) \subseteq (T^*\mu)^+\phi$.

(G8) If $\neg\phi \notin T^*\mu$ then $(T^*\mu)^+\phi \subseteq T^*(\mu \wedge \phi)$.

Katsuno and Mendelzon were interested in studying propositional knowledge base revision. Since Katsuno and Mendelzon also assume that the knowledge base is consistent, they make the simplification of representing each knowledge base K by a single formula. This can be done because if the knowledge base is propositional and consistent, then finding the models that satisfy K means finding the models that satisfy the conjunction of the formulas in K, i.e., some propositional formula ψ. Considering this simplification leads to an interesting translation of Gärdenfors' axioms.

If $\circ$ is any propositional knowledge base change operator, then for any consistent knowledge bases represented by the propositional formulas ψ, ψ_1, ψ_2 and for any propositional formulas μ, μ_1, μ_2 and ϕ the following should hold:

(KM1) $\psi \circ \mu$ implies μ.

(KM2) If $\psi \wedge \mu$ is satisfiable then $\psi \circ \mu \leftrightarrow \psi \wedge \mu$.

(KM3) If μ is satisfiable then $\psi \circ \mu$ is also satisfiable.

(KM4) If $\psi_1 \leftrightarrow \psi_2$ and $\mu_1 \leftrightarrow \mu_2$ then $\psi_1 \circ \mu_1 \leftrightarrow \psi_2 \circ \mu_2$.

(KM5) $(\psi \circ \mu) \wedge \phi$ implies $\psi \circ (\mu \wedge \phi)$.

(KM6) If $(\psi \circ \mu) \wedge \phi$ is satisfiable then $\psi \circ (\mu \wedge \phi)$ implies $(\psi \circ \mu) \wedge \phi$.

Axiom (KM1) assures that the new knowledge will hold in the revised knowledge base. Axiom (KM2) assures that if the new information is consistent with the current knowledge base, then the new information will be simply inserted into the knowledge base. Axiom (KM3) assures that no unwarranted inconsistency will be introduced. Axiom (KM4) says that the result of a revision operation should depend only on the set of models of the sentences in the knowledge base, not on the particular syntax of those sentences. This rule is called Dalal's Principle of Irrelevance of Syntax. Axioms (KM5) and (KM6) assure that the set of the models of the new information that are closest to the knowledge base are chosen as the result of the revision. See [KM91] for more on the meaning and implications of these axioms, and for proofs that the operators of Dalal [Dal88] and Fagin et al. [FUV83] are true revision operators, that is, they satisfy all of the above axioms.

We say that a theory change operator * corresponds with a knowledge base revision operator $\circ$ if and only if for each propositional knowledge base ψ and formula μ, the condition $cn(\psi)^*\mu = cn(\psi \circ \mu)$. The following proposition is from [KM91].

Proposition 2.1 **[KM91]** Let * be a theory change operator and $\circ$ be its corresponding operator on propositional knowledge bases. Then * satisfies axioms (G1-G8) if and only if $\circ$ satisfies axioms (KM1-KM6). $\square$

A *pre-order* $\leq$ over $\mathcal{M}$ is a reflexive and transitive relation on $\mathcal{M}$. A pre-order is *total* if for every pair $I, J \in \mathcal{M}$, either $I \leq J$ or $J \leq I$ holds. We define the relation $<$ as $I < J$ if and only if $I \leq J$ and $J \not\leq I$.

The set of *minimal models* of a subset $\mathcal{S}$ of $\mathcal{M}$ with respect to a pre-order $\leq_\psi$ is defined as:

$$Min(\mathcal{S}, \leq_\psi) = \{I \in \mathcal{S} : \nexists I' \in \mathcal{S} \text{ where } I' <_\psi I\}$$

Katsuno and Mendelzon gave the following model-theoretic characterization of revision and update when the knowledge base represented by a single propositional formula. Let the symbol $\circ$ denote revision and the symbol $\diamond$ denote update operators.

Suppose we have for each knowledge base ψ a total pre-ordering $\leq_\psi$ of interpretations for closeness to ψ, where the pre-order $\leq_\psi$ satisfies certain conditions [KM91]. Revision operators that satisfy the AGM postulates are exactly those that select from the models of the new information ϕ the closest models to the propositional knowledge base ψ. That is,

$$Mod(\psi \circ \phi) = Min(Mod(\phi), \leq_\psi)$$

For updates assume for each I some partial pre-ordering $\leq_I$ of interpretations for closeness to I. Update operators select for each model I in $Mod(\psi)$ the set of models from $Mod(\phi)$ that are closest to I. The new theory is the union of all such models. That is,

$$Mod(\psi \diamond \phi) = \bigcup_{I \in Mod(\psi)} Min(Mod(\phi), \leq_I)$$

Katsuno and Mendelzon's characterization is often useful to give simple proofs that particular theory change operators are revision or update operators. As an example of this from [KM91] consider Dalal's operator.

Dalal uses the number of terms on which two interpretations I and J differ as a measure of distance between them. That is, $dist(I, J)$ is the cardinality of the set $(I \setminus J) \cup (J \setminus I)$. For example, if $I = \{A, B, C\}$ and $J = \{C, D, E\}$, then $dist(I, J) = 4$.

Dalal then defines the distance between a knowledge base ψ and an interpretation I as the minimum distance between any model in $Mod(\psi)$ and I. Now take the pre-order $\leq_\psi$ defined by $I \leq_\psi J$ if and only if $dist(\psi, I) \leq dist(\psi, J)$.

For the revision $\psi \circ \mu$, Dalal's operator always returns the set of $\leq_\psi$ minimal models of μ. Hence by Katsuno and Mendelzon's characterization above, Dalal's operator is a true revision operator.

It is clear that both Katsuno and Mendelzon [KM91] and Gärdenfors [Gär88] eliminate from consideration any case of changing inconsistent theories or knowledge bases.

3 Model-Fitting

This section gives first a formal definition of the set of model-fitting operations and a model-theoretic characterization of it. Model-fitting is more general than revision because it allows the knowledge base to be inconsistent.

We say that a knowledge base K is satisfiable (or consistent) if and only if the conjunction of all propositional formulas in K is satisfiable. The set $Mod(K)$ is the set of models of the conjunction of all the propositional formulas in K. We say that $K \wedge \mu$ is satisfiable if and only if there is an interpretation that satisfies all the formulas in K and also satisfies μ. The formula $K_1 \rightarrow K_2$ is true if and only if $\forall_{\psi_2 \in K_2} \exists_{\psi_1 \in K_1} \psi_1 \leftrightarrow \psi_2$ is true. The formula $K_1 \leftrightarrow K_2$ is true if and only if $K_1 \rightarrow K_2$ and $K_2 \rightarrow K_1$ both hold.

We say that a knowledge base change operator $\triangleright$ is a model fitting operator if and only if $\triangleright$ satisfies the following axioms for each (not necessarily consistent) propositional knowledge base K, and formulas μ and ϕ:

(M1) $K \triangleright \mu$ implies μ.

(M2) If $K \wedge \mu$ is satisfiable then $K \triangleright \mu \leftrightarrow K \wedge \mu$.

(M3) If μ is satisfiable then $K \triangleright \mu$ is also satisfiable.

(M4) If $K_1 \leftrightarrow K_2$ and $\mu \leftrightarrow \phi$ then $K_1 \triangleright \mu \leftrightarrow K_2 \triangleright \phi$.

(M5) $(K \triangleright \mu) \wedge \phi$ implies $K \triangleright (\mu \wedge \phi)$.

(M6) If $(K \triangleright \mu) \wedge \phi$ is satisfiable then $K \triangleright (\mu \wedge \phi)$ implies $(K \triangleright \mu) \wedge \phi$.

(M7) $(K_1 \triangleright \mu) \wedge (K_2 \triangleright \mu)$ implies $(K_1 \cup K_2) \triangleright \mu$.

Here axioms (M1-M6) are generalizations of axioms (KM1-KM6) for propositional knowledge bases that may be inconsistent. Axiom (M7) asserts that any model that is closest to both K_1 in μ and to K_2 in μ must also be a closest model to $K_1 \cup K_2$ in μ. The intuition behind axiom (M7) can be seen from the following example. Suppose that we have two committees of five people each. Suppose that the two committees both come up with a set of possible actions and that both sets of possible actions include some action A. Then it is reasonable to suppose that if the two committees were joined, then it would still come up with a set of possible actions that includes A.

The next theorem presents a model-theoretic characterization of model-fitting operators that satisfy axioms (M1-M7). At first we define for each knowledge base K a relation that orders interpretations in $\mathcal{M}$ with respect to their closeness to K. First we define a loyal assignment.

A *loyal* assignment is a function that assigns for each knowledge base K a pre-order $\leq_K$ such that the following four conditions hold. For each $I, J \in \mathcal{M}$ and knowledge bases K, K_1, K_2:

(1) If $I, J \in Mod(K)$ then $I <_K J$ does not hold.
(2) If $I \in Mod(K)$ and $J \notin Mod(K)$ then $I <_K J$.
(3) If $K_1 \leftrightarrow K_2$ then $\leq_{K_1} = \leq_{K_2}$.
(4) If $I \leq_{K_1} J$ and $I \leq_{K_2} J$ then $I \leq_{K_1 \cup K_2} J$.

In the definition of loyal the first three conditions are generalizations of the conditions of the definition of a faithful assignment in [KM92], while last condition is new. Using these definitions, the characterization theorem can now be stated as follows.

Theorem 3.1 A knowledge base operator satisfies axioms (M1-M7) if and only if there exists a loyal assignment that maps each knowledge base K to a total pre-order $\leq_K$ such that $Mod(K \triangleright \mu) = Min(Mod(\mu), \leq_K)$.

Proof Sketch: The first six axioms will depend on the first three conditions of loyalness, while the last axiom will depend on the fourth condition of loyalness. A detailed proof is given in the appendix. $\square$.

Theorem 3.1 is useful to prove in a simple way that particular theory change operators are model-fitting operators. As an example, consider the following operator.

Using Dalal's distance measure between interpretations (see Section 2), we define the overall distance *odist* between a knowledge base K and an interpretation I as follows. For any propositional sentence ψ we define the distance as:

$$odist(\psi, I) = \min_{J \in Mod(\psi)} dist(I, J)$$

Then for a knowledge base K we define:

$$odist(K, I) = \max_{\psi \in K} odist(\psi, I)$$

Then we assign to each knowledge base K the total pre-order $\leq_K$ defined by $I \leq_K J$ if and only if $odist(K, I) \leq odist(K, J)$. It is easy to see that this is a loyal assignment.

Condition (1) is satisfied because if $I, J \in Mod(K)$ then $odist(K, I) = odist(K, J) = 0$, hence $I <_K J$ does not hold. . Condition (2) is satisfied because if $I \in Mod(K)$ then $odist(K, I) = 0$ must be true. If $J \notin Mod(K)$ then $odist(\psi, J) > 0$ for some $\psi \in K$, hence $odist(K, J) > 0$ and $I <_K J$. Condition (3) is satisfied because K_1 and K_2 have the same set of models. Finally for condition (4), if $I \leq_{K_1} J$ and $I \leq_{K_2} J$ then $odist(K_1, I) \leq odist(K_1, J)$ and $odist(K_2, I) \leq odist(K_2, J)$ must be true. Then $odist(K_1 \cup K_2, I) = \max(odist(K_1, I), odist(K_2, I)) \leq \max(odist(K_1, J), odist(K_2, J)) = odist(K_1 \cup K_2, J)$. Hence $I \leq_{K_1 \cup K_2} J$ must hold.

We now define the result of $Mod(K \triangleright \mu)$ to be the minimal models of μ according to the ordering $\leq_K$. Then by Theorem 3.1 this operator satisfies axioms (M1-M7) and is a proper model-fitting operator.

Example 3.1 As an application of model-fitting consider a database class with three students. Assume that at the university where the class is held three database programming languages are available, namely SQL, Datalog and O_2. The instructor considers teaching either both SQL and Datalog or Datalog only. This can be represented as $\mu = (S \wedge D \wedge \neg O) \vee (\neg S \wedge D \wedge \neg O)$. The first student would like to learn either SQL or O_2, the second would like either Datalog or O_2 but not both, and the third would like to learn all three. That is the students suggest to the instructor to teach $\psi_1 = (S \vee O) \wedge \neg D$, $\psi_2 = (\neg S \wedge D \wedge \neg O) \vee (\neg S \wedge \neg D \wedge O)$ and $\psi_3 = S \wedge D \wedge O$.

Considering only the propositional terms S, D, and O, $Mod(\mu) = \{\{S, D\}, \{D\}\}$, while $Mod(\psi_1) = \{\{S\}, \{O\}, \{S, O\}\}$, and $Mod(\psi_2) = \{\{D\}, \{O\}\}$, and $Mod(\psi_3) = \{\{S, D, O\}\}$. We calculate that $odist(\psi_1, \{S, D\}) = 1$ and $odist(\psi_1, \{D\}) = 2$ and $odist(\psi_2, \{S, D\}) = 1$ and $odist(\psi_2, \{D\}) = 0$ and $odist(\psi_3, \{S, D\}) = 1$ and $odist(\psi_3, \{D\}) = 2$.

Our knowledge base $K = \{\psi_1, \psi_2, \psi_3\}$. Hence we find that $odist(K, \{S, D\}) = 1 < odist(K, \{D\}) = 2$. Therefore, $Mod(K \triangleright \mu) = \{\{S, D\}\}$. This indicates that the instructor could best satisfy the class by teaching both SQL and Datalog. □

Example 3.1 is a situation which calls for model-fitting instead of revision. Revision operators satisfying [KM92] would fail because they would find that the conjunction of the three propositions ψ_1, ψ_2, ψ_3 is unsatisfiable. On the other hand, model-fitting handles this case quite nicely. Note that if the instructor decides to teach Datalog only then one student will be very satisfied, but the other two may well drop the class. Clearly this is not what we want. The choice of $\{S, D\}$ is the model that best fits the whole class, and will keep all students reasonably satisfied.

4 Weighted Model-Fitting

In this section we generalize the results of the previous section by considering weighted knowledge bases. A weighted knowledge base is a function $\tilde{K}$ from model sets to nonnegative real numbers. We denote the weight of a model set M in weighted knowledge base $\tilde{K}$ as $\tilde{K}(M)$. The real numbers are intended to describe the relative degree of importance of the model sets within the weighted knowledge base. Clearly this is a generalization of the knowledge bases of the previous section, because we can translate a regular knowledge base $K = \{\psi_1, \ldots, \psi_n\}$ into a weighted knowledge base $\tilde{K}$ having for each model set M weight $\tilde{K}(M) = 1$ if $M = Mod(\psi_i)$ for some $1 \leq i \leq n$ and $\tilde{K}(M) = 0$ otherwise.

In this section we replace the union operation on regular knowledge bases by the weighted union operation on weighted knowledge bases. If $\tilde{K_1}$ and $\tilde{K_2}$ are two weighted knowledge bases, we take their weighted union, denoted as $\tilde{K_1} \uplus \tilde{K_2}$, to be the weighted knowledge base $\tilde{K}$ such that for each model set M, $\tilde{K}(M) = \tilde{K_1}(M) + \tilde{K_2}(M)$.

We use the function *Form* to map weighted knowledge bases into regular knowledge bases. More precisely, if $S_1, \ldots S_n$ are the model sets with nonzero weights in $\tilde{K}$, then the expression $Form(\tilde{K})$ denotes the regular knowledge base K that consists of the set of formulas $form(S_i)$ for each $1 \leq i \leq n$. As in the previous section, if μ is a formula we say that $K \wedge \mu$ is satisfiable if and only if there is an interpretation that satisfies all the formulas in K and also satisfies μ.

We say that a knowledge base operator is a weighted model-fitting operator if and only if it satisfies the following axioms for all weighted propositional knowledge bases $\tilde{K}, \tilde{K_1}, \tilde{K_2}$ and propositional formulas μ and ϕ:

(W1) $\tilde{K} \triangleright \mu$ implies μ.

(W2) If $Form(\tilde{K}) \wedge \mu$ is satisfiable then $\tilde{K} \triangleright \mu \leftrightarrow Form(\tilde{K}) \wedge \mu$.

(W3) If μ is satisfiable then $\tilde{K} \triangleright \mu$ is also satisfiable.

(W4) If $\mu \leftrightarrow \phi$ then $\tilde{K} \triangleright \mu \leftrightarrow \tilde{K} \triangleright \phi$.

(W5) $(\tilde{K} \triangleright \mu) \wedge \phi$ implies $\tilde{K} \triangleright (\mu \wedge \phi)$.

(W6) If $(\tilde{K} \triangleright \mu) \wedge \phi$ is satisfiable then $\tilde{K} \triangleright (\mu \wedge \phi)$ implies $(\tilde{K} \triangleright \mu) \wedge \phi$.

(W7) $(\tilde{K_1} \triangleright \mu) \wedge (\tilde{K_2} \triangleright \mu)$ implies $(\tilde{K_1} \uplus \tilde{K_2}) \triangleright \mu$.

(W8) If $(\tilde{K_1} \triangleright \mu) \wedge (\tilde{K_2} \triangleright \mu)$ is satisfiable then $(\tilde{K_1} \uplus \tilde{K_2}) \triangleright \mu$ implies $(\tilde{K_1} \triangleright \mu) \wedge (\tilde{K_2} \triangleright \mu)$.

Here axioms (W1-W6) are generalizations of axioms (M1-M6) to weighted knowledge bases. Furthermore, we have generalized (M7) into two axioms (W7) and (W8) that together express the following condition: the closest models to $\tilde{K_1} \uplus \tilde{K_2}$ in μ are exactly the intersection of the closest models to $\tilde{K_1}$ in μ and the closest models to $\tilde{K_2}$ in μ if the intersection is nonempty.

We say that a weighted knowledge base $\tilde{K}$ is satisfiable if and only if the intersection of all model sets with nonzero weights in $\tilde{K}$ is nonempty. A weighted knowledge base is unsatisfiable if and only if it is not satisfiable. We say that an interpretation I is a model of a weighted knowledge base $\tilde{K}$, written as $I \in Mod(\tilde{K})$, if and only if I is an element of each model set with nonzero weight in $\tilde{K}$.

A *weighted loyal* assignment is a function that assigns for each weighted knowledge base $\tilde{K}$ a pre-order $\leq_{\tilde{K}}$ such that the following four conditions hold. For each $I, J \in \mathcal{M}$ and weighted knowledge bases $\tilde{K}, \tilde{K_1}, \tilde{K_2}$:

(1) If $I, J \in Mod(\tilde{K})$ then $I <_{\tilde{K}} J$ does not hold.
(2) If $I \in Mod(\tilde{K})$ and $J \notin Mod(\tilde{K})$ then $I <_{\tilde{K}} J$.

(3) If $I \leq_{\tilde{K}_1} J$ and $I \leq_{\tilde{K}_2} J$ then $I \leq_{\tilde{K}_1 \uplus \tilde{K}_2} J$.
(4) If $I <_{\tilde{K}_1} J$ and $I \leq_{\tilde{K}_2} J$ then $I <_{\tilde{K}_1 \uplus \tilde{K}_2} J$.

The following theorem is a model-theoretic characterization of weighted model-fitting.

Theorem 4.1 A knowledge base operator satisfies axioms (W1-W8) if and only if there exists a weighted loyal assignment that maps each weighted knowledge base $\tilde{K}$ to a total pre-order $\leq_{\tilde{K}}$ such that $Mod(\tilde{K} \triangleright \mu) = Min(Mod(\mu), \leq_{\tilde{K}})$.

Proof Sketch: The first seven axioms will depend on the first three conditions of weighted loyal assignments and vice versa, while the last axiom and the fourth condition of weighted loyal assignments will depend on each other. A detailed proof is given in the appendix. □

Next we see an example of a weighted model-fitting operator. We define the weighted distance *wdist* between a weighted knowledge base $\tilde{K}$ and an interpretation I as:

$$wdist(\tilde{K}, I) = \sum_{M \in 2^{\mathcal{M}}} \tilde{K}(M) * wdist(M, I)$$

where the distance between a model set and a model is defined as:

$$wdist(M, I) = \min_{J \in M} dist(I, J)$$

Next we define for each weighted knowledge base $\tilde{K}$ the total pre-order $\leq_{\tilde{K}}$ such that $I \leq_{\tilde{K}} J$ if and only if $wdist(\tilde{K}, I) \leq wdist(\tilde{K}, J)$. It is easy to show that this is a weighted loyal assignment.

Condition (1) is satisfied because if $I, J \in Mod(\tilde{K})$ then $wdist(M, I) = wdist(M, I) = 0$ for each model set M with nonzero weight in $\tilde{K}$. Hence $wdist(\tilde{K}, I) = wdist(\tilde{K}, J) = 0$ and $I <_K J$ does not hold. Condition (2) is satisfied because if $I \in Mod(\tilde{K})$ then $wdist(\tilde{K}, I) = 0$ must be true. If $J \notin Mod(K)$ then $wdist(M, J) > 0$ for some M with nonzero weight in $\tilde{K}$, hence $wdist(\tilde{K}, J) > 0$ and $I <_{\tilde{K}} J$. Condition (3) is satisfied because if $I \leq_{\tilde{K}_1} J$ and $I \leq_{\tilde{K}_2} J$ then $wdist(\tilde{K}_1, I) \leq wdist(\tilde{K}_1, J)$ and $wdist(\tilde{K}_2, I) \leq wdist(\tilde{K}_2, J)$ must be true. Then $wdist(\tilde{K}_1 \uplus \tilde{K}_2, I) = \sum_{M \in 2^{\mathcal{M}}} \tilde{K}(M) * wdist(M, I) = \sum_{M \in 2^{\mathcal{M}}} (\tilde{K}_1(M) + \tilde{K}_2(M)) * wdist(M, I) = \sum_{M \in 2^{\mathcal{M}}} \tilde{K}_1(M) * wdist(M, I) + \sum_{M \in 2^{\mathcal{M}}} \tilde{K}_2(M) * wdist(M, I) = wdist(\tilde{K}_1, I) + wdist(\tilde{K}_2, I) \leq wdist(\tilde{K}_1, J) + wdist(\tilde{K}_2, J) = wdist(\tilde{K}_1 \uplus \tilde{K}_2, J)$. Hence $I \leq_{\tilde{K}_1 \uplus \tilde{K}_2} J$ must hold. Condition (4) can be shown similarly to condition (3).

We now define the result of $Mod(\tilde{K} \triangleright \mu)$ to be the minimal models of μ according to the ordering $\leq_{\tilde{K}}$. Then by Theorem 4.1 the operator is a weighted model-fitting operator.

Example 4.1 As an example of weighted model-fitting consider a database class with the same instructor as in Example 3.1 but with 35 students. The instructor's offering can be represented by the same sentence μ as in Example 3.1.

Suppose that 10 students would like to learn either SQL or O_2, 20 would like either Datalog or O_2 but not both, and 5 would like to learn all three languages. The students' requests can be represented by a weighted knowledge base $\tilde{K}$ in which the only model sets with nonzero weights are $Mod(\psi_1)$ with weight 10, $Mod(\psi_2)$ with weight 20 and $Mod(\psi_3)$ with weight 5, where the formulas ψ_1, ψ_2, ψ_3 are as in Example 3.1. Now we calculate that $wdist(\tilde{K}, \{S, D\}) = 35$ and $wdist(\tilde{K}, \{D\}) = 30$. Hence $Mod(\tilde{K} \triangleright \mu) = \{\{D\}\}$. This indicates that in this case the instructor could best satisfy the class by teaching Datalog only. □

Note that in the case of weighted model-fitting the instructor tries to satisfy the majority of the class, instead of trying to satisfy each member to the best degree possible. The outcome changes from Example 3.1 due to the large number of students who want to learn Datalog only.

Example 4.2 As another example, suppose that each sentence of a weighted knowledge base $\tilde{K}$ describes a different person. In particular the first person is a male married lawyer who likes to play golf, the second is a female non-smoking married lawyer who likes to play golf, the third is a non-smoking single real estate agent who does not like golf, the fourth is a male non-smoking married lawyer, and the fifth is a female smoking married real estate agent. Using the variables F for female, G for golfer, L for lawyer, M for married, R for real estate agent and S for smoker, we can describe the persons by the following propositional formulas.

$$\begin{aligned}
\psi_1 &= M \wedge L \wedge \neg R \wedge G \wedge \neg F \\
\psi_2 &= M \wedge L \wedge \neg R \wedge \neg S \wedge G \wedge F \\
\psi_3 &= \neg M \wedge \neg L \wedge R \neg S \wedge \neg G \\
\psi_4 &= M \wedge L \wedge \neg R \wedge \neg S \wedge \neg F \\
\psi_5 &= M \wedge \neg L \wedge R \wedge S \wedge F
\end{aligned}$$

Suppose now that somebody asks the following: what is a typical married lawyer like? Let $\mu = M \wedge L$. To answer the question, we need to consider only the sentences that describe married lawyers, that is those ψ_i for which $\psi_i \rightarrow \mu$ holds. It is easy to see that only sentences ψ_1, ψ_2 and ψ_4 are interesting. Hence we temporarily reset the weights of these sentences to 1 and the weights of the other two sentences to 0. Then intuitively typical married lawyers can be found by $\tilde{K} \triangleright \mu$, because this chooses among all possible models or descriptions of married lawyers those that are overall closest to the description of married lawyers in the knowledge base.

Using the sample weighted model fitting operator we find that $Mod(\tilde{K} \triangleright \mu) = \{M, L, G\}$. Hence in words the typical married lawyer according to the weighted knowledge base is a not a real estate agent and is a male non-smoker who likes to play golf. □

5 Arbitration

In this section we describe arbitration as a generalization of weighted model-fitting. An arbitration operator takes as input a weighted and a regular knowledge base. If $\tilde{K_1}$ and K_2 are such knowledge bases then the arbitration of $\tilde{K_1}$ by K_2, denoted as $\tilde{K_1} \,\triangle\, K_2$, will return a knowledge base containing the formulas within K_2 that are closest to $\tilde{K_1}$. Note that we are now interested in closest formulas instead of closest models. To simplify the exposition, in this section we will assume that each formula in a regular knowledge base is satisfiable.

In addition in this section if K_1 and K_2 are knowledge bases we define $Modset(K_1) = \{Mod(\mu) : \mu \in K_1\}$, and we take $K_1 \subseteq K_2$ to be true if and only if $Modset(K_1) \subseteq Modset(K_2)$ in the regular sense, and we take $K_1 \cap K_2 = \{\mu_1 \in K_1 : \exists \mu_2 \in K_2 \text{ and } Mod(\mu_1) = Mod(\mu_2)\}$. If μ is a formula, we take $\mu \in K_1$ to be true if and only if $Mod(\mu) \in Modset(K_1)$ in the regular sense.

We say that a knowledge base operator is an arbitration operator if and only if it satisfies the following axioms for all propositional knowledge base K_2, weighted knowledge bases $\tilde{K_1}$ and $\tilde{K_3}$ and propositional formula μ:

(A1) $\tilde{K_1} \,\triangle\, K_2 \subseteq K_2$.

(A2) If $\exists \mu \in K_2$ such that $\mu \wedge Form(\tilde{K_1})$ is satisfiable then
$\tilde{K_1} \bigtriangleup K_2 = \{\mu \in K_2 : \mu \wedge Form(\tilde{K_1}) \text{ is satisfiable}\}$.

(A3) If K_2 is nonempty then $\tilde{K_1} \bigtriangleup K_2$ is nonempty.

(A4) If $K_2 \leftrightarrow K_3$ then $\tilde{K_1} \bigtriangleup K_2 \leftrightarrow \tilde{K_1} \bigtriangleup K_3$.

(A5) $(\tilde{K_1} \bigtriangleup K_2) \cap K_3 \subseteq \tilde{K_1} \bigtriangleup (K_2 \cap K_3)$.

(A6) If $(\tilde{K_1} \bigtriangleup K_2) \cap K_3$ is nonempty then $\tilde{K_1} \bigtriangleup (K_2 \cap K_3) \subseteq (\tilde{K_1} \bigtriangleup K_2) \cap K_3$.

(A7) $(\tilde{K_1} \bigtriangleup K_2) \cap (\tilde{K_3} \bigtriangleup K_2) \subseteq (\tilde{K_1} \uplus \tilde{K_3}) \bigtriangleup K_2$.

(A8) If $(\tilde{K_1} \bigtriangleup K_2) \cap (\tilde{K_3} \bigtriangleup K_2)$ is nonempty then $(\tilde{K_1} \uplus \tilde{K_3}) \bigtriangleup K_2 \subseteq (\tilde{K_1} \bigtriangleup K_2) \cap (\tilde{K_3} \bigtriangleup K_2)$.

A *generalized loyal* assignment is a function that assigns for each weighted knowledge base $\tilde{K}$ a pre-order $\leq_{\tilde{K}}$ such that the following five conditions hold. For all formulas $\mu_1, \mu_2, \phi_1, \phi_2$ and weighted knowledge bases $\tilde{K}, \tilde{K_1}, \tilde{K_2}$:

(1) If $\mu_1 \wedge Form(\tilde{K})$ and $\mu_2 \wedge Form(\tilde{K})$ are satisfiable, then $\mu_1 <_{\tilde{K}} \mu_2$ does not hold.
(2) If $\mu_1 \wedge Form(\tilde{K})$ is satisfiable and $\mu_2 \wedge Form(\tilde{K})$ is unsatisfiable, then $\mu_1 <_{\tilde{K}} \mu_2$.
(3) If $\mu_1 \leftrightarrow \mu_2$ and $\phi_1 \leftrightarrow \phi_2$, then $\mu_1 \leq_{\tilde{K}} \phi_1$ if and only if $\mu_2 \leq_{\tilde{K}} \phi_2$.
(4) If $\mu_1 \leq_{\tilde{K_1}} \mu_2$ and $\mu_1 \leq_{\tilde{K_2}} \mu_2$ then $\mu_1 \leq_{\tilde{K_1} \uplus \tilde{K_2}} \mu_2$.
(5) If $\mu_1 <_{\tilde{K_1}} \mu_2$ and $\mu_1 \leq_{\tilde{K_2}} \mu_2$ then $\mu_1 <_{\tilde{K_1} \uplus \tilde{K_2}} \mu_2$.

We can now give a characterization of arbitration operators in terms of closest formulas that is the analogue of the characterization of model-fitting operators in terms of closest models. The proof of the theorem is given in the appendix.

Theorem 5.1 A knowledge base operator $\bigtriangleup$ satisfies axioms (A1-A8) if and only if there exists a generalized loyal assignment that maps each weighted knowledge base $\tilde{K_1}$ to a total pre-order $\leq_{\tilde{K_1}}$ such that $\tilde{K_1} \bigtriangleup K_2 = Min(K_2, \leq_{\tilde{K_1}})$. $\Box$

We now define an example arbitration operator and then apply it in some examples. We define the distance between a weighted knowledge base and a formula to be:

$$gdist(\tilde{K}, \mu) = \sum_{M \in 2^{\mathcal{M}}} \tilde{K}(M) * gdist(M, \mu)$$

where the distance between a model set and a formula is:

$$gdist(M, \mu) = \min_{I \in M,\, J \in Mod(\mu)} dist(I, J)$$

Next we define for each weighted knowledge base $\tilde{K}$ the total pre-order $\leq_{\tilde{K}}$ such that $\mu_1 \leq_{\tilde{K}} \mu_2$ if and only if $gdist(\tilde{K}, \mu_1) \leq gdist(\tilde{K}, \mu_2)$. Clearly this is a generalized loyal assignment. We also define the result of $Mod(\tilde{K_1} \bigtriangleup K_2)$ to be the minimal models of K_2 according to the ordering $\leq_{\tilde{K_1}}$. Then by Theorem 5.1 the operator is an arbitration operator.

Example 5.1 Consider a hypothetical election where there are two candidates and three voters. The following statements are made by each.

First candidate: We should balance the budget and cut taxes.

Second candidate: We need a national health-care and higher taxes on the rich.

First voter: Balancing the budget should be first priority. If that is done, I'm willing to pay higher taxes, but I'm opposed to national health-care because it reduces patient choice.

Second voter: I want both a national health-care and lower taxes.

Third voter: I can't pay health insurance because of the high taxes. We should either cut taxes or introduce a national health-care. Balancing the budget during a recession is a bad idea.

We can represent this election situation by two propositional knowledge bases, a weighted knowledge base $\tilde{K_1}$ describing the voters and a regular knowledge base K_2 describing the candidates. We will use the propositional letters B for balanced budget, H for national health-care and T for higher taxes. We assume that taxes never stay the same, they either rise or fall. The only model sets in $\tilde{K_1}$ with nonzero weight will be $Mod(\psi_1)$, $Mod(\psi_2)$ and $Mod(\psi_3)$, each with weight 1, and we also will have $K_2 = \{\mu_1, \mu_2\}$, where the models of the sentences are the following.

$Mod(\psi_1) = \{\{B\}, \{B, T\}\}$
$Mod(\psi_2) = \{\{B, H\}, \{H\}\}$
$Mod(\psi_3) = \{\{\}, \{H\}\{H, T\}\}$
$Mod(\mu_1) = \{\{B\}, \{B, H\}\}$
$Mod(\mu_2) = \{\{H, T\}, \{B, H, T\}\}$

Intuitively, we need to find the candidate whose platform appeals more to the voters. That is, we need to find out which platform is closest to the voters' desires.

We assume that the voters tend to view candidates optimistically. They are likely to ask themselves: "What would be the best scenario for me if this candidate is elected?". In this case, the arbitration operator defined above can be applied, and we calculate that:

$gdist(Mod(\psi_1), \mu_1) = 0, \ \ gdist(Mod(\psi_1), \mu_2) = 1$
$gdist(Mod(\psi_2), \mu_1) = 0, \ \ gdist(Mod(\psi_2), \mu_2) = 1$
$gdist(Mod(\psi_3), \mu_1) = 1, \ \ gdist(Mod(\psi_3), \mu_2) = 0$

Hence $gdist(\tilde{K_1}, \mu_1) = 1$ and $gdist(\tilde{K_1}, \mu_2) = 2$. Therefore $\tilde{K_1} \triangle K_2 = \{\mu_1\}$. This indicates that the candidate advocating μ_1 appeals better to the voters in general. □

Note that both the first and the second voter will vote for the first candidate μ_1, even though it is impossible that μ_1 will satisfy both completely (as they optimistically expect). That is, when they vote, the first voter expects $\{B\}$ while the second voter expects $\{B, H\}$ to happen. Clearly the two cannot happen at the same time. This may seem like an error in our definition of arbitration, but it is not really. Perhaps it simply shows that life is not as logical as some may wish it to be. The example illustrates a potential problem during elections. Voters and news reporters wisely ask candidates to be specific on issues, otherwise a candidate who does not commit to anything specific could win.

Example 5.2 Suppose instead of three voters, we have three groups of voters with $20, 15$ and 65 members respectively. Assume that members of group one, two and three feel the same way as the first, second, and third voter in Example 5.1, respectively. This requires that the weights of $Mod(\psi_1)$, $Mod(\psi_2)$ and $Mod(\psi_3)$ be increased to $20, 15$ and 65 respectively. We calculate that $gdist(\tilde{K_1}, \mu_1) = 65$ and $gdist(\tilde{K_1}, \mu_2) = 35$. Therefore $K_1 \triangle K_2 = \{\mu_2\}$. In this case the second candidate could win the election. □

6 Arbitration of First-Order Knowledge Bases

In this section we describe arbitration of knowledge bases that contain first-order sentences instead of propositional formulas. Let the symbol ω denote the set of all natural numbers, and the symbol $\oplus$ denote the symmetric set difference operation[1].

Consider a first-order function free language $\mathcal{L}$ built from the following primitive components: A set $A = \{a_i : i \in \omega\}$ of *domain elements*, a set $X = \{x_i : i \in \omega\}$ of *variables*, a set $R = \{R_i : i \in \omega\}$ of *relation symbols*, $\wedge$ (and), $\neg$ (negation), $\exists$ (existential quantifier), $=$ (equality), and the parenthesis symbols. As a notational convenience, in the examples below we will also use the symbols $\vee$ (or), $\rightarrow$ (implication), $\leftrightarrow$ (mutual implication), $\neq$ (inequality) and $\forall$ (universal quantifier) defined in terms of the primitive components of $\mathcal{L}$ in the usual way.

With each relation symbol $R_i \in R$ we associate the *arity* $\alpha(i)$. A *k-ary term* is a tuple with k components, each in $A \cup X$. A *literal* is an expression of the form $R_i(x)$ where R_i is in R and x is an $\alpha(i)$-ary term. A *ground literal* is a literal that has no variables in it. An *atomic formula* is a literal, or an expression of the form $x_i = x_j$, or an expression of the form $x_i = a_j$, where $\{x_i, x_j\} \subseteq X$, and $a_j \in A$.

The set of all *well formed formulas* of $\mathcal{L}$ is defined in the usual way, and it is denoted Φ'. The subset of *sentences* in Φ' is denoted Φ. If ϕ is a formula where variable x_i occurs free, then $\phi(x_i/a_j)$ denotes the formula ϕ with each free occurrence of x_i substituted by a_j. For $\phi \in \Phi$, define the schema $\sigma(\phi)$ to be the set of all relation symbols appearing in ϕ.

The set of all ground literals is denoted by $\mathcal{G}$. A database db is a set of ground literals. The schema of db, denoted $\sigma(db)$ is the set of all relation symbols appearing in db. Let db_1 and db_2 be databases. Then we say that $\sigma(db_2)$ *dominates* $\sigma(db_1)$, if $\sigma(db_1)$ is a subset of $\sigma(db_2)$.

The set of all databases is denoted $\mathcal{DB}$. By $\mathcal{DB}_{\mathbf{s}}$ we mean the set of all databases on schema $\mathbf{s}$. Furthermore, if B is a subset of the domain A, then $\mathcal{DB}_{\mathbf{s}}^B$ denotes the set of all databases on schema $\mathbf{s}$ containing only values in B.

The *interpretation* of a sentence $\phi \in \Phi$ *w.r.t.* a database db is a relation $\models$ on $\mathcal{DB} \times \Phi$ defined for db and ϕ if and only if $\sigma(db)$ dominates $\sigma(\phi)$, in which case the recursive definition is:

$$\begin{array}{rll} db \models (a_i = a_j) & \text{iff} & i = j \\ db \models R_i(x) & \text{iff} & R_i(x) \in db \\ db \models (\phi \wedge \psi) & \text{iff} & db \models \phi \text{ and } db \models \psi \\ db \models (\neg\phi) & \text{iff} & (\mathcal{G} \setminus db) \models \phi \\ db \models (\exists x_i \phi) & \text{iff} & db \models \phi(x_i/a_j) \text{ for some } a_j \in A \end{array}$$

By $Mod(\phi)$ we mean the set of all databases that are models of ϕ, i.e., $Mod(\phi) = \{db \in \mathcal{DB} : db \models \phi\}$. We say that ϕ *finitely implies* ψ, if $Mod(\phi) \subseteq Mod(\psi)$. If $\mathbf{s}$ is a scheme and B is a subset of the domain A, then $Mod(\phi)_{\mathbf{s}}^B$ denotes the set $Mod(\phi) \cap \mathcal{DB}_{\mathbf{s}}^B$.

A *sentence base* is a finite set of databases. The scheme of a sentence base is the union of the schemes of the databases in it. The set of all sentence bases is denoted $\mathcal{SB}$. If B is a subset of the domain A, then $\mathcal{SB}_{\mathbf{s}}^B$ denotes the set of all sentence bases on scheme $\mathbf{s}$ containing only values in B.

A *first-order knowledge base* K is a finite set of first-order sentences. A *weighted first-order knowledge base* $\tilde{K}$ is a function from sentence bases to nonnegative real numbers. The scheme of K is the union of the schemes of the sentences in it. The scheme of $\tilde{K}$ is the union of the schemes of the sentence bases with nonzero weights.

We now define an example arbitration operator for first-order knowledge bases. Let K be a first-order knowledge base and $\tilde{K}$ be a weighted first-order knowledge base. Let $\mathbf{s}$ be the scheme of $\tilde{K}$ and let B be the set of domain elements occurring in the sentence bases with nonzero weights in $\tilde{K}$. Then we define the distance between $\tilde{K}$ and any $\mu \in K$ as:

[1] $A \oplus B = (A \setminus B) \cup (B \setminus A)$

$$d(\tilde{K}, \mu) = \sum_{sb \in \mathcal{SB}_s^B} \tilde{K}(sb) * d(sb, \mu)$$

where the distance between a sentence base and a formula is:

$$d(sb, \mu) = \min_{db \in sb,\, db_1 \in Mod(\mu)_s^B} d(db, db_1)$$

where the distance between two databases is defined as the cardinality of their set differences. That is:

$$d(db, db_1) = card(db \oplus db_1)$$

Next we define for each weighted first-order knowledge base $\tilde{K}$ the total pre-order $\leq_{\tilde{K}}$ such that $\mu_1 \leq_{\tilde{K}} \mu_2$ if and only if $d(\tilde{K}, \mu_1) \leq d(\tilde{K}, \mu_2)$.

Then we say that μ_1 is $\leq_{\tilde{K}}$*-minimal in* K, if $\mu_1 \in K$ and there is no $\mu_2 \in K$ such that $\mu_2 \leq_{\tilde{K}} \mu_1$ and $\mu_1 \not\leq_{\tilde{K}} \mu_2$.

It is easy to verify that this is a generalized loyal assignment, that is, it satisfies all the five conditions of the previous section, with the knowledge bases taken to be first-order instead of propositional.

We then define an arbitration operator that always returns the set of minimal formulas according to the above assignment. It can be shown similarly to Theorem 5.1 that this first-order arbitration operator satisfies axioms (A1-A8).

Example 6.1 Suppose we have a database class with three students who make the following requests to the instructor:

Alice: I want to use the same languages as Brian.

Brian: I want to use SQL only.

Carl: I want to use any language if and only if it is used by at least one other student in the class.

It is clear to all students that every student has to use some programming language, that the only students in the class are Alice, Brian and Carl, and that the only programming languages that they can choose from are SQL and Datalog.

The instructor is trying to decide whether to require everyone to use both SQL and Datalog or to require everyone to use Datalog only. Which is better for this class?

At first we represent the student's requests using first-order sentences as follows. We take the scope of the quantifiers in the first-order sentences to be the set of integer numbers and strings of the English alphabet. In the sentences we will use only the binary relation symbol u, whose first argument will describe the name of a person and the second argument will describe the name of a database programming language. We can represent:

Alice's request by sentence ψ_1:

$$\forall_x \; u(Alice, x) \leftrightarrow u(Brian, x)$$

Brian's request by sentence ψ_2:

$$u(Brian, S) \wedge \forall_x \ u(Brian, x) \rightarrow x = S$$

Carl's request by sentence ψ_3:

$$\forall_x \ u(Carl, x) \leftrightarrow (\exists_y \ u(y, x) \wedge y \neq Carl)$$

The assumption that every student in the class uses some language by ϕ_1:

$$\exists_{x,y,z} \ u(Alice, x) \wedge u(Brian, y) \wedge u(Carl, z)$$

The assumption that the three students choose from the two languages by ϕ_2:

$$\forall_{x,y} \ u(x, y) \rightarrow ((x = Alice \vee x = Brian \vee x = Carl) \wedge (y = S \vee y = D))$$

The teacher's first option by μ_1:

$$\forall_x (x = Alice \vee x = Brian \vee x = Carl) \rightarrow (u(x, S) \wedge u(x, D) \wedge \neg \exists_y (y \neq S \wedge y \neq D \wedge u(x, y)))$$

And the teacher's second option by μ_2:

$$\forall_x (x = Alice \vee x = Brian \vee x = Carl) \rightarrow (u(x, D) \wedge \neg \exists_y (y \neq D \wedge u(x, y)))$$

Let $\phi = \phi_1 \wedge \phi_2$ and let $\psi'_i = \psi_i \wedge \phi$ and $\mu'_j = \mu_j \wedge \phi$ for all $1 \leq i \leq 3$ and $1 \leq j \leq 2$. It is implicit in this problem that each of the students and the teacher agrees to ϕ also.

Note that we can restrict consideration to databases with scheme $\mathbf{s} = \{u\}$ and domain $B = \{Alice, Brian, Carl, D, S\}$. We can represent the students' requests by a weighted first-order knowledge base $\tilde{K_1}$ which has in it only the sentence bases $Mod(\psi'_1)^B_{\mathbf{s}}$, $Mod(\psi'_2)^B_{\mathbf{s}}$ and $Mod(\psi'_3)^B_{\mathbf{s}}$ with nonzero weights. Moreover, since each student should be considered equally, each of the three sentence bases should have a weight of 1.

We can also represent the teacher's options by a first-order knowledge base $K_2 = \{\mu'_1, \mu'_2\}$. Then the problem can be solved by arbitrating $\tilde{K_1}$ by K_2. For this we can use the arbitration operator defined in this section.

As the table below shows, there are nine databases within each of the sentence bases $Mod(\psi'_1)^B_{\mathbf{s}}$, $Mod(\psi'_2)^B_{\mathbf{s}}$ and $Mod(\psi'_3)^B_{\mathbf{s}}$. To save some space in the table we use the constant symbols A, B and C instead of the names Alice, Brian and Carl. Also, only databases that are models of ϕ are shown in the table. The presence of a database within a sentence base is indicated by a "y" in the appropriate column.

Note that only the database $db_1 = \{u(A, S), u(A, D), u(B, S), u(B, D), u(C, S), u(C, D)\}$ is in $Mod(\mu'_1)^B_{\mathbf{s}}$ and that only the database $db_2 = \{u(A, D), u(B, D), u(C, D)\}$ belongs to $Mod(\mu'_2)^B_{\mathbf{s}}$. Therefore the distance between any database and μ_1 will be just the distance between that database and db_1. Similarly, the distance between any database and μ_2 will be the distance between that database and db_2. This is how the values in the last two columns can be calculated for each given row. Now we can calculate by taking the minimum of the appropriate values:

$d(Mod(\psi'_1)^B_{\mathbf{s}}, \mu'_1) = 0 \quad d(Mod(\psi'_1)^B_{\mathbf{s}}, \mu'_2) = 0$
$d(Mod(\psi'_2)^B_{\mathbf{s}}, \mu'_1) = 1 \quad d(Mod(\psi'_2)^B_{\mathbf{s}}, \mu'_2) = 2$

$d(Mod(\psi_3')^B_{\mathbf{s}}, \mu_1') = 0 \quad d(Mod(\psi_3')^B_{\mathbf{s}}, \mu_2') = 0$

By taking sums we have that $d(\tilde{K_1}, \mu_1) = 1$ and $d(\tilde{K_1}, \mu_2) = 2$. Therefore we can conclude that the first option is overall closer to the requests of the students. Hence it is better if the teacher asks everyone to use both SQL and Datalog. □

Database *db*	$\models \psi_1'$	$\models \psi_2'$	$\models \psi_3'$	$d(db, \mu_1')$	$d(db, \mu_2')$
u(A,S),u(B,S),u(C,S),	y	y	y	3	6
u(A,S),u(B,S),u(C,D)	y	y		3	4
u(A,S),u(B,S),u(C,S),u(C,D)	y	y		2	5
u(A,S),u(B,D),u(C,S)				3	4
u(A,S),u(B,D),u(C,D)				3	2
u(A,S),u(B,D),u(C,S),u(C,D)			y	2	3
u(A,S),u(B,S),u(B,D),u(C,S)				2	5
u(A,S),u(B,S),u(B,D),u(C,D)				2	3
u(A,S),u(B,S),u(B,D),u(C,S),u(C,D)			y	1	4
u(A,D),u(B,S),u(C,S)		y		3	4
u(A,D),u(B,S),u(C,D)		y		3	2
u(A,D),u(B,S),u(C,S),u(C,D)		y	y	2	3
u(A,D),u(B,D),u(C,S)	y			3	2
u(A,D),u(B,D),u(C,D)	y		y	3	0
u(A,D),u(B,D),u(C,S),u(C,D)	y			2	1
u(A,D),u(B,S),u(B,D),u(C,S)				2	3
u(A,D),u(B,S),u(B,D),u(C,D)				2	1
u(A,D),u(B,S),u(B,D),u(C,S),u(C,D)			y	1	2
u(A,S),u(A,D),u(B,S),u(C,S)		y		2	5
u(A,S),u(A,D),u(B,S),u(C,D)		y		2	3
u(A,S),u(A,D),u(B,S),u(C,S),u(C,D)		y	y	1	4
u(A,S),u(A,D),u(B,D),u(C,S)				2	3
u(A,S),u(A,D),u(B,D),u(C,D)				2	1
u(A,S),u(A,D),u(B,D),u(C,S),u(C,D)			y	1	2
u(A,S),u(A,D),u(B,S),u(B,D),u(C,S)	y			1	4
u(A,S),u(A,D),u(B,S),u(B,D),u(C,D)	y			1	2
u(A,S),u(A,D),u(B,S),u(B,D),u(C,S),u(C,D)	y		y	0	3

7 Relationship with other Operators

Model-fitting operators are related to the various decision making protocols of Borgida and Imielinski [BI84], which is the first syntax-independent proposal for handling inconsistencies that arise from more than two sources. In [BI84] decision making in committees is used as a model for resolving inconsistencies. Unfortunately, the decision making protocols described in the paper do not go far enough in resolving inconsistencies. In particular some protocols allow committees to vacillate, that is, to support a decision p and at the same time also support its negation $\neg p$. This is different from model-fitting operators, which by definition only allow consistent answers.

Arbitration operators and the combination operator of Baral et al. [BKM91, BKMS92] also have a strong similarity in their aims. The combination operator was developed for helping building expert systems, where the knowledge of several experts has to be combined. Arbitration can be also applied in this case as mentioned in the introduction. However, there are significant differences in the way arbitration and combination work.

One important difference is that combination operators do not use weights. This makes it difficult to emphasize properly two experts' opinions who happen to agree. More importantly, combination operators are syntax-sensitive, that is, they violate axiom (M4). Beyond these differences, there are a number of others which can be best illustrated by giving an example. Before that let us first consider the original definitions.

Definition 7.1 **[BKMS92]** Let P be a knowledge base and IC be a set of integrity constraints. A subset $Q \subseteq P$ is said to be maximally consistent with priority to IC if and only if $Q \cup IC$ is consistent and for every $Q \subset Q' \subseteq P$, it is the case that $Q' \cup IC$ is inconsistent. MAXCONS(P,IC) is the set of maximally consistent subsets of P with priority to IC. □

Definition 7.2 **[BKMS92]** Let $K_1, \ldots, K_k$ and IC be knowledge bases. Then the combination of $K_1, \ldots, K_k$ subject to the integrity constraints IC is $Comb(\{K_1, \ldots, K_k\}, IC) =_{\mathtt{def}} MAXCONS(K_1 \cup \ldots \cup K_k,\, IC)$. □

Let us consider again Example 6.1. We can represent this problem in the framework of [BKMS92] as having three knowledge bases $K_1 = \{\psi_1, \phi\}$, $K_2 = \{\psi_2, \phi\}$ and $K_3 = \{\psi_3, \phi\}$. Then for any integrity constraint IC the combination of the three knowledge bases will be:

$$Comb(\{K_1, K_2, K_3\}, IC) = MAXCONS(K_1 \cup K_2 \cup K_3, IC) = MAXCONS(\{\psi_1, \psi_2, \psi_3, \phi\}, IC)$$

If we take IC to be the theory $\{\mu_1, \mu_2, \phi\}$, then the result of the combination is undefined, because μ_1 and μ_2 cannot be both true, hence no subset of the union of the three knowledge bases can be consistent with IC.

If we take $IC = \{\mu_1 \vee \mu_2, \phi\}$ as the set of integrity constraints, then the only maximal consistent subset of the union of the three knowledge bases with priority to IC will be $\{\psi_1, \psi_3, \phi\}$.

This example illustrates that while there are some similarities between combination and arbitration, there are subtle differences as well. While the combination operator finds the largest set of students whose wishes can be simultaneously satisfied, arbitration finds a solution that satisfies to the best degree possible each student. Hence combination does not solve the teacher's problem. Inasmuch as combination suggests a solution it must be that Brian should be kicked out. However, a moment reflection tells us that even that would not work, because if Brian is not taking the course, then Alice's wish also will not be fulfilled.

Since this paper was submitted for review Jinxin Lin and Alberto Mendelzon [JM94] and Paolo Liberatore and Marco Schaerf [LSch95] have proposed other operators for merging information within knowledge bases. Benczur, Novak and Revesz [BNR95] have also considered a different way of adding weights to arbitration in the propositional case. A detailed comparison of these operators can be found in [BNR96].

8 Open Problems

An open problem is to consider instead of total orders partial orders among models and formulas, similarly to [Win88]. Another open problem is to further analyze and compare the computational complexity of various cases of revision, update, and arbitration with each other [ASV90, EG92, GMR92]. A third interesting open problem is to consider the interaction of the three operators in a system that alternately uses all of them. As one facet of this interaction, the weights of the formulas could now be considered to be dynamic. For example, the weight of a witness' testimony declines each time it is revised during cross examination. This would effect subtly the result of any later arbitration operations that are done in the system.

Acknowledgements: I thank Gosta Grahne and Alberto Mendelzon for helpful comments on previous drafts of this paper.

References

[AG85] S. Abiteboul & G. Grahne. Update semantics for incomplete databases. *Proceedings of the Eleventh International Conference on Very Large Databases*, pages 1–12, 1985.

[ASV90] S. Abiteboul, E. Simon & V. Vianu. Non-deterministic languages to express deterministic transformations. *Proceedings of the Ninth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 218–229, 1990.

[AGM85] C. E. Alchourrón, P. Gärdenfors & D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, (50), pages 510–530, 1985.

[BS81] F. Bancilhon & N. Spyratos. Update semantics of relational views. *ACM Transactions on Database Systems*, (4), pages 557–575, 1981.

[BKM91] C. Baral, S. Kraus & J. Minker. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering*, vol. 3, no. 2, pages 208–220, 1991.

[BKMS92] C. Baral, S. Kraus, J. Minker & V.S. Subrahmanian. Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, vol. 8, no. 1, pages 45–71, 1992.

[BNR95] A. Benczur, A. Novak, & P. Z. Revesz. On weighted knowledgebase transformation. In: *Proc. Fourth Symposium on Programming Languages and Software Tools*, p. 404-415, Visegrad, Hungary, 1995.

[BNR96] A. Benczur, A. Novak, & P. Z. Revesz. On classical and weighted knowledgebase transformations. *Computers and Mathematics*, in print.

[Bon90] A. J. Bonner. Hypothetical Datalog: Complexity and expressibility. *Theoretical Computer Science*, vol. 76, pages 3–51, 1990.

[Bor85] A. Borgida. Language features for flexible handling of exceptions in information systems. *ACM Transaction on Database Systems*, 10, pages 563–603, 1985.

[BI84] A. Borgida & T. Imielinski. Decision making in committees: A framework for dealing with inconsistency and non-monotonicity. *Proceedings Workshop on Nonmonotonic Reasoning*, pages 21–32, 1984.

[Dal88] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. *Proceedings of the AAAI*, pages 475–479, 1988.

[EG92] T. Eiter & G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Proceedings of the Eleventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 261-273, 1992.

[FUV83] R. Fagin, J. D. Ullman & M. Y. Vardi. On the semantics of updates in databases. *Proceedings of the Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 352–365, 1983.

[FKUV86] R. Fagin, G. Kuper, J. D. Ullman & M. Y. Vardi. Updating logical databases. In: *Advances in Computing Research*, Vol. 3, P. C. Kanellakis and F. Preparata eds., pages 1–18, JAI Press, 1986.

[Gab85] D. M. Gabbay. N-Prolog: an extension of Prolog with hypothetical implications. II. Logical foundations and negation as failure. *Journal of Logic Programming*, (2) 251–283, 1985.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* Bradford Books, MIT Press, Cambridge, MA, 1988.

[Gin86] M. L. Ginsberg. Counterfactuals. *Artificial Intelligence*, vol. 30, no. 1, 1986.

[Gra91] G. Grahne. Updates and counterfactuals. *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 269–276, 1991.

[GM91] G. Grahne & A. O. Mendelzon. Updates and subjunctive queries. R. Demolombe, T. Imielinski & L. Fariñãs del Cerro (Eds.), *Workshop on Nonstandard Queries and Answers*, pages 33–52, 1991.

[GMR92] G. Grahne, A. O. Mendelzon, & P. Z. Revesz. Knowledgebase transformations. *Proceedings of the Eleventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 246–260, 1992.

[Gro88] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, Vol. 17, pages 157–170, 1988.

[Han91] S. O. Hansson. *Belief Base Dynamics.* Ph.D. thesis, Department of Philosophy, Uppsala University, 1991.

[JM94] J. Lin & A. O. Mendelzon. Private communication.

[KW85] A. M. Keller & M. Winslett Wilkins. On the use of an extended relational model to handle changing incomplete information. *IEEE Trans. on Software Engineering*, 11:7, pages 620–633, 1985.

[KM89] H. Katsuno & A. O. Mendelzon. A unified view of propositional knowledge base updates. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1413–1419, 1989.

[KM91] H. Katsuno & A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52, pages 263–294, 1991.

[KM92] H. Katsuno & A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In: *Belief Revision*, Gärdenfors, ed., pages 183–203, Cambridge University Press, 1992. (Preliminary version in: *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 387–394, 1991.)

[LSch95] P. Liberatore & M. Schaerf. Arbitration: A commutative operator for belief revision. *Proc. World Conference on Artificial Intelligence*, 1995.

[McC68] J. McCarthy. Programs with common sense. In: M. Minsky, (Ed.), *Semantic Information Processing*, MIT Press, Cambridge, MA, pages 403–418, 1968.

[Mak85] D. Makinson. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthèse*, (62), pages 347–363, 1985

[Neb90] B. Nebel. *Reasoning and Revision in Hybrid Representation Systems.* Springer-Verlag Lecture Notes in Artificial Intelligence, 422, 1990.

[Mey90] R. van der Meyden. Recursively indefinite databases *Proceedings of the Third International Conference on Database Theory*, pages 364–378, 1990.

[Rei92] R. Reiter. On specifying database updates. *Proceedings of the Third International Conference on Extending Database Technology*, 1992.

[Rev93] P. Z. Revesz. On the semantics of theory change: Arbitration between old and new information. *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 71–82, 1993.

[Ry91] M. D. Ryan. Belief revision and ordered theory presentations. *Proceedings of the Eighth Amsterdam Colloquium on Logic*, 1991.

[Sat88] K. Satoh. Nonmonotonic reasoning by minimal belief revision. *Proceedings International Conference on Fifth Generation Systems*, pages 455–462, 1988.

[Web86] A. Weber. Updating propositional formulas. *Proceedings First Conference on Expert Database Systems*, pages 487–500, 1986.

[Win88] M. Winslett. Reasoning about action using a possible models approach. *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 89–93, 1988.

[Zad78] L.A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, vol. 1, no. 1, pages 3–28, 1978.

A Proofs

Proof of Theorem 3.1: (Only-if) Assume that axioms (M1-M7) hold for a model-fitting operator $\triangleright$. We define a loyal assignment as follows. For each knowledge base K we define a total pre-order $\leq_K$ in terms of the $\triangleright$ operator as follows. For each (not necessarily distinct) pair I, J of models, let $I \leq_K J$ if and only if $I \in Mod(K \triangleright form(I, J))$.

We have to show three things: (1) that for each knowledge base K the assignment $\leq_K$ is a total pre-order, (2) that the function from knowledge bases to assignments is loyal, and (3) that $Mod(K \triangleright \mu) = Min(Mod(\mu), \leq_K)$.

(1) We need to show that $\leq_K$ is total, reflexive, and transitive when K is satisfiable.

total By axioms (M1) and (M3) the $Mod(K \triangleright form(I, J))$ is a nonempty subset of $\{I, J\}$. Hence any pair of models are comparable, making $\leq_K$ a total relation.

reflexive By axioms (M1) and (M3) the $Mod(K \triangleright form(I))$ is a nonempty subset of $\{I\}$. Hence $\leq_K$ is also reflexive.

transitive Assume that the relation $\leq_K$ is not transitive, that is, for some I, J_1, and J_2 models $I \leq_K J_1$, $J_1 \leq_K J_2$, and $I \not\leq_K J_2$.

Then by the definition of $\leq_K$, $I \notin Mod(K \triangleright form(I, J_2))$. By axiom (M5), $I \notin Mod(K \triangleright form(I, J_1, J_2)) \cap \{I, J_2\}$. Hence $I \notin Mod(K \triangleright form(I, J_1, J_2))$. There are two possible cases. Either (i) $J_1 \in Mod(K \triangleright form(I, J_1, J_2))$ or (ii) $J_1 \notin Mod(K \triangleright form(I, J_1, J_2))$.

In case (i), we know that I is not in $Mod(K \triangleright form(I, J_1, J_2)) \cap \{I, J_1\}$ and that $Mod(K \triangleright form(I, J_1, J_2)) \wedge form(I, J_1)$ is satisfiable. Then by (M6) also $I \notin Mod(K \triangleright form(I, J_1))$. This contradicts the assumption that $I \leq_K J_1$.

In case (ii), by (M1) and (M3) we know that $J_2 = Mod(K \triangleright form(I, J_1, J_2))$. Hence $Mod(K \triangleright form(I, J_1, J_2)) \cap \{J_1, J_2\}$ is satisfiable but does not contain J_1. Hence by (M6) also $J_1 \notin Mod(K \triangleright form(J_1, J_2))$. This contradicts the assumption that $J_1 \leq_K J_2$.

(2) The first condition of loyalness follows easily from the definition of $\leq_K$. To see the second condition, assume that $I \in Mod(K)$ and $J \notin Mod(K)$. Then by axiom (M2), $Mod(K \triangleright form(I, J)) = \{I\}$. Therefore $I <_K J$ holds. The third condition of loyalness follows from axiom (M4). To see the fourth condition assume that $I \leq_{K_1} J$ and $I \leq_{K_2} J$ both hold for knowledge bases K_1 and K_2. Then $I \in Mod(K_1 \triangleright form(I, J))$ and $I \in Mod(K_2 \triangleright form(I, J))$. Hence by (M7) $I \in Mod((K_1 \cup K_2) \triangleright form(I, J))$ and therefore by the assignment we chose $I \leq_{K_1 \cup K_2} J$ must also hold.

(3) We need to show both the $\subseteq$ and the $\supseteq$ directions. If μ is unsatisfiable, then $Mod(K \triangleright \mu) = \emptyset = Min(Mod(\mu), \leq_K)$. Hence assume that μ is satisfiable.

($\subseteq$) Assume that $I \in Mod(K \triangleright \mu)$ and $I \notin Min(Mod(\mu), \leq_K)$. By (M1) $I \in Mod(\mu)$. Since I is not a minimal model, according to the definition of minimal there must be another model J in $Mod(\mu)$ such that $J <_K I$ (i.e., such that $J \leq_K I$ and $I \not\leq_K J$). By the definition of $\leq_K$ then $J \in Mod(K \triangleright form(I, J))$ and $I \notin Mod(K \triangleright form(I, J))$.
Since both I and J are in $Mod(\mu)$, $\mu \wedge form(I, J) = form(I, J)$. Hence I is also not in $Mod(K \triangleright (\mu \wedge form(I, J)))$. By (M5) and using $\phi = form(I, J)$ we know that $Mod((K \triangleright \mu) \wedge form(I, J))$ implies $Mod(K \triangleright (\mu \wedge form(I, J)))$. Hence also $I \notin Mod((K \triangleright \mu) \wedge form(I, J))$. Therefore, I cannot be in $Mod(K \triangleright \mu)$, which is a contradiction.

($\supseteq$) Assume now that $I \notin Mod(K \triangleright \mu)$ and $I \in Min(Mod(\mu), \leq_K)$. By the definition of minimal, $I \in Mod(\mu)$. Since μ is satisfiable, by (M3) there is some model J in $Mod(K \triangleright \mu)$, and by (M1) also $J \in Mod(\mu)$. Since both I and J are in $Mod(\mu)$, $\mu \wedge form(I, J) = form(I, J)$. Hence by (M5) and (M6) and letting ϕ be $form(I, J)$ we derive that $Mod((K \triangleright \mu) \wedge form(I, J)) = Mod(K \triangleright \mu) \cap \{I, J\} = Mod(K \triangleright form(I, J))$.

Since ϕ is satisfiable, by (M1) and (M3), $Mod(K \triangleright form(I,J))$ is a nonempty subset of $\{I,J\}$. But the identity above and $I \notin Mod(K \triangleright \mu)$ implies that also $I \notin Mod(K \triangleright form(I,J))$. Hence $J = Mod(K \triangleright form(I,J))$. Therefore $J <_K I$. Hence I cannot be a minimal model according to $\leq_K$, i.e., $I \notin Min(Mod(\mu), \leq_K)$. This is again a contradiction.

(If) Assume that for a knowledge base operator $\triangleright$ there is a loyal function that assigns to each satisfiable knowledge base K a total pre-order $\leq_K$ such that $Mod(K \triangleright \mu) = Min(Mod(\mu), \leq_K)$. We need to show that $\triangleright$ satisfies axioms (M1-M7).

(M1) Axiom (M1) follows because the minimal model of μ with respect to any total pre-order is always by definition some subset of $Mod(\mu)$.

(M2) Suppose $K \wedge \mu$ is satisfiable. Then it suffices to show that $Mod(K \wedge \mu) = Min(Mod(\mu), \leq_K)$. Suppose $I \in Mod(K \wedge \mu)$. Then $I \in Mod(\mu)$ and $I \in Mod(K)$. Then by conditions (1&2) of loyalness, $J <_K I$ does not hold for any $J \in Mod(K)$ and $I <_K J$ holds for any $J \notin Mod(K)$. Therefore I must be a minimal model among $Mod(\mu)$ according to $\leq_K$. Hence $I \in Min(Mod(\mu), \leq_K)$ also holds. This shows that $Mod(K \wedge \mu) \subseteq Min(Mod(\mu), \leq_K)$

Now to show the other direction that $Mod(K \wedge \mu) \supseteq Min(Mod(\mu), \leq_K)$ assume that $I \in Min(Mod(\mu), \leq_K)$ but $I \notin Mod(K \wedge \mu)$. By axiom (M1) we know that $I \in Mod(\mu)$, hence also $I \notin Mod(K)$. Since $K \wedge \mu$ is satisfiable, there is an interpretation J such that $J \in Mod(K \wedge \mu)$, that is $J \in Mod(K)$ and $J \in Mod(\mu)$. By the second condition of loyalness, $J <_K I$ holds. Hence I is not a minimal model of μ with respect to $\leq_K$, which is a contradiction.

(M3) Axiom (M3) follows because as long as μ is satisfiable there is some minimal model in $Mod(\mu)$ with respect to K.

(M4) Axiom (M4) follows from the third condition of loyalness.

(M5) Assume that (M5) is false. Then for some I model (1) $I \in Mod((K \triangleright \mu) \wedge \phi)) = Min(Mod(\mu), \leq_K) \cap Mod(\phi)$ and (2) $I \notin Mod(K \triangleright (\mu \wedge \phi)) = Min(Mod(\mu \wedge \phi), \leq_K)$.

Then $I \in Mod(\phi)$ and $I \in Min(Mod(\mu), \leq_K)$ must hold. By the definition of minimal models $I \in Mod(\mu)$ must also hold. Therefore, $I \in Mod(\mu \wedge \phi)$ must hold. That and assumption (2) imply there must be another model J such that $J \in Mod(\mu \wedge \phi)$ and $J <_K I$. Note that $J \in Mod(\mu)$ must also hold and this and $J <_K I$ imply that $I \notin Min(Mod(\mu), \leq_K)$, which is a contradiction.

(M6) Assume that (M6) is false. Then $(K \triangleright \mu) \wedge \phi$ must be satisfiable and for some I model (1) $I \notin Min(Mod(\mu), \leq_K) \cap Mod(\phi)$ and (2) $I \in Min(Mod(\mu \wedge \phi), \leq_K)$.

Since we assume that $(K \triangleright \mu) \wedge \phi$ is satisfiable, there must be a model J such that $J \in Min(Mod(\mu), \leq_K) \cap Mod(\phi)$. Then $J \in Mod(\mu \wedge \phi)$ must hold. By assumption (2) and because $\leq_K$ is total, $I \leq_K J$ holds. Since $J \in Min(Mod(\mu), \leq_K) \cap Mod(\phi)$, also $J \in Min(Mod(\mu), \leq_K)$ must hold. We claim that $I \in Min(Mod(\mu), \leq_K)$ must also hold. Suppose that it does not hold. Then there must be a model $H \in Mod(\mu)$ such that $H <_K I$, i.e., $H \leq_K I$ holds and $I \leq_K H$ does not hold. Since J is a minimal model in μ, $J \leq_K H$ holds. Then by transitivity of $I \leq_K J$ and $J \leq_K H$ we have $I \leq_K H$, which is a contradiction. Hence $I \in Min(Mod(\mu), \leq_K)$ must hold.

From assumption (2) and the definition of minimal models, it is clear that $I \in Mod(\mu)$ and $I \in Mod(\phi)$. The latter and assumption (1) imply that $I \notin Min(Mod(\mu), \leq_K)$. This is a contradiction to the claim that we proved in the previous paragraph. Hence assumptions (1-2) were wrong.

(M7) If I is minimal in $Mod(\mu)$ according to both K_1 and K_2, then for any J in $Mod(\mu)$ both $I \leq_{K_1} J$ and $I \leq_{K_2} J$ must be true. Then by the fourth condition of loyalness I is also minimal in $Mod(K_1 \cup K_2)$. This implies that axiom (M7) also holds.

□

Proof of Theorem 4.1: The proof of this theorem will be similar to the proof of Theorem 3.1. Let $\triangleright$ be any weighted model-fitting operator. Then take for each weighted knowledge base $\tilde{K}$ the total pre-order $\leq_{\tilde{K}}$ to be the relation that assigns for each (not necessarily distinct) pair I, J of models $I \leq_{\tilde{K}} J$ if and only if $I \in Mod(\tilde{K} \triangleright form(I,J))$.

We can argue similarly as in Theorem 3.1, but we also have to show the mutual dependence of the fourth condition of weighted loyal assignments and axiom (W8).

To show the fourth condition of weighted loyal assignment when (W8) holds, assume that $I <_{K_1} J$ and $I \leq_{K_2} J$. Then I is and J is not in $Mod(K_1 \triangleright form(I,J))$, and I is also in $Mod(K_2 \triangleright form(I,J))$. Hence $I = Mod((K_1 \triangleright form(I,J)) \wedge (K_2 \triangleright form(I,J)))$. Then by (W7) and (W8) also $I = Mod((K_1 \uplus K_2) \triangleright form(I,J))$. Then by the definition of assignments $I <_{K_1 \uplus K_2} J$.

For the reverse, to show (W8) when the forth condition of loyal assignments holds, suppose that I is both $\leq_{K_1}$ and $\leq_{K_2}$ minimal in $Mod(\mu)$ and that axiom (W8) does not hold. Then there is some model J that is $\leq_{K_1 \uplus K_2}$ minimal in $Mod(\mu)$ but w.l.g. not $\leq_{K_1}$ minimal in $Mod(\mu)$. Then $I <_{K_1} J$ and $I \leq_{K_2} J$. Then by the fourth condition of loyalness $I <_{K_1 \uplus K_2} J$. Hence J cannot be $\leq_{K_1 \uplus K_2}$ minimal in $Mod(\mu)$, which is a contradiction. Hence (W8) holds. □

Proof of Theorem 5.1: (Only-if) Assume that axioms (A1-A8) hold for an arbitration operator $\triangle$. We define a generalized loyal assignment as follows. For each weighted knowledge base $\tilde{K_1}$ we define a total pre-order $\leq_{\tilde{K_1}}$ in terms of the $\triangle$ operator as follows. For each (not necessarily distinct) pair μ_1, μ_2 of formulas, let $\mu_1 \leq_{\tilde{K_1}} \mu_2$ if and only if $\mu_1 \in \tilde{K_1} \triangle \{\mu_1, \mu_2\}$.

We have to show three things: (1) that for each weighted knowledge base $\tilde{K_1}$ the assignment $\leq_{\tilde{K_1}}$ is a total pre-order, (2) that the function from weighted knowledge bases to assignments is generalized loyal, and (3) that $\tilde{K_1} \triangle K_2 = Min(K_2, \leq_{\tilde{K_1}})$.

(1) We need to show that $\leq_{\tilde{K_1}}$ is total, reflexive, and transitive.

total By axioms (A1) and (A3) the result of $\tilde{K_1} \triangle \{\mu_1, \mu_2\}$ is a nonempty subset of $\{\mu_1, \mu_2\}$. Hence any pair of formulas are comparable, making $\leq_{\tilde{K_1}}$ a total relation.

reflexive By axioms (A1) and (A3) the result of $\tilde{K_1} \triangle \{\mu_1\}$ is a nonempty subset of $\{\mu_1\}$. Hence $\leq_{\tilde{K_1}}$ is also reflexive.

transitive Assume that the relation $\leq_{\tilde{K_1}}$ is not transitive, that is, for some μ_1, μ_2, and μ_3 formulas $\mu_1 \leq_{\tilde{K_1}} \mu_2$, $\mu_2 \leq_{\tilde{K_1}} \mu_3$, and $\mu_1 \not\leq_{\tilde{K_1}} \mu_3$.

Then by the definition of $\leq_{\tilde{K_1}}$, $\mu_1 \notin \tilde{K_1} \triangle \{\mu_1, \mu_3\}$. By axiom (A5), $\mu_1 \notin (\tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}) \cap \{\mu_1, \mu_3\}$. Hence $\mu_1 \notin \tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}$. There are two possible cases. Either (i) $\mu_2 \in \tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}$ or (ii) $\mu_2 \notin \tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}$.

In case (i), we know that μ_1 is not in $(\tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}) \cap \{\mu_1, \mu_2\}$ and that $(\tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}) \cap \{\mu_1, \mu_2\}$ is satisfiable. Then by (A6) also $\mu_1 \notin \tilde{K_1} \triangle \{\mu_1, \mu_2\}$. This contradicts the assumption that $\mu_1 \leq_{\tilde{K_1}} \mu_2$.

In case (ii), by (A1) and (A3) we know that $\{\mu_3\} = \tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}$. Hence $(\tilde{K_1} \triangle \{\mu_1, \mu_2, \mu_3\}) \cap \{\mu_2, \mu_3\}$ is satisfiable but does not contain μ_2. Hence by (A6) also $\mu_2 \notin \tilde{K_1} \triangle \{\mu_2, \mu_3\}$. This contradicts the assumption that $\mu_2 \leq_{\tilde{K_1}} \mu_3$.

(2) The first condition of generalized loyalness follows easily from the definition of $\leq_{\tilde{K_1}}$. To see the second condition, assume that $\mu_1 \wedge Form(\tilde{K_1})$ is satisfiable and $\mu_1 \wedge Form(\tilde{K_1})$ is unsatisfiable. Then by axiom (A2), $\tilde{K_1} \triangle \{\mu_1, \mu_2\} = \{\mu_1\}$. Therefore $\mu_1 <_{\tilde{K_1}} \mu_2$ holds. The third condition of generalized loyalness follows from axiom (A4). The see the fourth condition assume that $\mu_1 \leq_{\tilde{K_1}} \mu_2$ and $\mu_1 \leq_{\tilde{K_2}} \mu_2$ both hold for weighted knowledge bases $\tilde{K_1}$ and $\tilde{K_2}$. Then $\mu_1 \in (\tilde{K_1} \triangle \{\mu_1, \mu_2\})$ and

$\mu_1 \in (\tilde{K_2} \triangle \{\mu_1, \mu_2\})$. Hence by (A7) $\mu_1 \in ((\tilde{K_1} \uplus \tilde{K_2}) \triangle \{\mu_1, \mu_2\})$ and therefore by the assignment we chose $\mu_1 \leq_{\tilde{K_1} \uplus \tilde{K_2}} \mu_2$ must also hold.

To show the fifth condition of generalized loyal assignment, assume that $\mu_1 <_{\tilde{K_1}} \mu_2$ and $\mu_1 \leq_{\tilde{K_2}} \mu_2$. Then μ_1 is and μ_2 is not in $\tilde{K_1} \triangle \{\mu_1, \mu_2\}$, and μ_1 is also in $\tilde{K_2} \triangle \{\mu_1, \mu_2\}$. Hence $\{\mu_1\} = (\tilde{K_1} \triangle \{\mu_1, \mu_2\}) \cap (\tilde{K_2} \triangle \{\mu_1, \mu_2\})$. Then by (A7) and (A8) also $\{\mu_1\} = (\tilde{K_1} \uplus \tilde{K_2}) \triangle \{\mu_1, \mu_2\}$. Then by the definition of assignments $\mu_1 <_{\tilde{K_1} \uplus \tilde{K_2}} \mu_2$.

(3) We need to show both the $\subseteq$ and the $\supseteq$ directions. If K_2 is empty, then $\tilde{K_1} \triangle K_2 = \emptyset = Min(K_2, \leq_{\tilde{K_1}})$. Hence assume that K_2 is nonempty.

($\subseteq$) Assume that $\mu_1 \in \tilde{K_1} \triangle K_2$ and $\mu_1 \notin Min(K_2, \leq_{\tilde{K_1}})$. By (A1) $\mu_1 \in K_2$. Since μ_1 is not a minimal formula, according to the definition of minimal there must be another formula μ_2 in K_2 such that $\mu_2 <_{\tilde{K_1}} \mu_1$ (i.e., such that $\mu_2 \leq_{\tilde{K_1}} \mu_1$ and $\mu_1 \not\leq_{\tilde{K_1}} \mu_2$). By the definition of $\leq_{\tilde{K_1}}$ then $\mu_2 \in (\tilde{K_1} \triangle \{\mu_1, \mu_2\})$ and $\mu_1 \notin (\tilde{K_1} \triangle \{\mu_1, \mu_2\})$.

Since both μ_1 and μ_2 are in K_2, $K_2 \cap \{\mu_1, \mu_2\} = \{\mu_1, \mu_2\}$. Hence μ_1 is also not in $\tilde{K_1} \triangle (K_2 \cap \{\mu_1, \mu_2\})$. By (A5) and using $K_3 = \{\mu_1, \mu_2\}$ we know that $(\tilde{K_1} \triangle K_2) \cap \{\mu_1, \mu_2\} \subseteq \tilde{K_1} \triangle (K_2 \cap \{\mu_1, \mu_2\})$. Hence also $\mu_1 \notin ((\tilde{K_1} \triangle K_2) \cap \{\mu_1, \mu_2\})$. Therefore, μ_1 cannot be in $\tilde{K_1} \triangle K_2$, which is a contradiction.

($\supseteq$) Assume now that $\mu_1 \notin (\tilde{K_1} \triangle K_2)$ and $\mu_1 \in Min(K_2, \leq_{\tilde{K_1}})$. By the definition of minimal, $\mu_1 \in K_2$. Since K_2 is nonempty, by (A3) there is some formula μ_2 in $\tilde{K_1} \triangle K_2$, and by (A1) also $\mu_2 \in K_2$. Since both μ_1 and μ_2 are in K_2, $K_2 \cap \{\mu_1, \mu_2\} = \{\mu_1, \mu_2\}$. Hence by (A5) and (A6) and letting K_3 be $\{\mu_1, \mu_2\}$ we get that $(\tilde{K_1} \triangle K_2) \cap \{\mu_1, \mu_2\} = \tilde{K_1} \triangle (K_2 \cap \{\mu_1, \mu_2\}) = \tilde{K_1} \triangle \{\mu_1, \mu_2\}$.

By (A1) and (A3), $\tilde{K_1} \triangle \{\mu_1, \mu_2\}$ is a nonempty subset of $\{\mu_1, \mu_2\}$. But the identity above and $\mu_1 \notin (\tilde{K_1} \triangle K_2)$ implies that also $\mu_1 \notin (\tilde{K_1} \triangle \{\mu_1, \mu_2\})$. Hence $\{\mu_2\} = \tilde{K_1} \triangle \{\mu_1, \mu_2\}$. Therefore $\mu_2 <_{\tilde{K_1}} \mu_1$. Hence μ_1 cannot be a minimal formula according to $\leq_{\tilde{K_1}}$, i.e., $\mu_1 \notin Min(K_2, \leq_{\tilde{K_1}})$. This is again a contradiction.

(If) Assume that for a knowledge base operator $\triangle$ there is a generalized loyal function that assigns to each satisfiable knowledge base $\tilde{K_1}$ a total pre-order $\leq_{\tilde{K_1}}$ such that $\tilde{K_1} \triangle K_2 = Min(K_2, \leq_{\tilde{K_1}})$. We need to show that $\triangle$ satisfies axioms (A1-A8).

(A1) Axiom (A1) follows because the minimal formula of K_2 with respect to any total pre-order is always by definition some subset of K_2.

(A2) Suppose that there is a formula μ in K_2 such that $\mu \wedge Form(\tilde{K_1})$ is satisfiable. Then it suffices to show that $\{\mu \in K_2 : \mu \wedge Form(\tilde{K_1}) \text{ is satisfiable}\} = Min(K_2, \leq_{\tilde{K_1}})$.

Suppose $\mu_1 \in \{\mu \in K_2 : \mu \wedge Form(\tilde{K_1}) \text{ is satisfiable}\}$. Then $\mu_1 \in K_2$ and $\mu_1 \wedge Form(\tilde{K_1})$ is satisfiable. Consider any other formula μ_2 in K_2. If $\mu_2 \wedge Form(\tilde{K_1})$ is satisfiable, then by the first condition of generalized loyal assignments $\mu_2 <_{\tilde{K_1}} \mu_1$ does not hold. If $\mu_2 \wedge Form(\tilde{K_1})$ is unsatisfiable, then by the second condition of generalized loyal assignments $\mu_1 <_{\tilde{K_1}} \mu_2$ holds. Therefore μ_1 must be a minimal formula in K_2 according to $\leq_{\tilde{K_1}}$. Hence $\mu_1 \in Min(K_2, \leq_{\tilde{K_1}})$ also holds. This shows one direction of the equality.

To show the other direction $\{\mu \in K_2 : \mu \wedge Form(\tilde{K_1}) \text{ is satisfiable}\} \supseteq Min(K_2, \leq_{\tilde{K_1}})$ assume $\mu_1 \notin \{\mu \in K_2 : \mu \wedge Form(\tilde{K_1}) \text{ is satisfiable}\}$ but $\mu_1 \in Min(K_2, \leq_{\tilde{K_1}})$. By axiom (A1) we know that $\mu_1 \in K_2$, hence it must be the case that $\mu_1 \wedge Form(\tilde{K_1})$ is unsatisfiable. By the second condition of generalized loyalness, $\mu <_{\tilde{K_1}} \mu_1$ holds. Hence μ_1 is not a minimal formula of K_2 with respect to $\leq_{\tilde{K_1}}$, which is a contradiction.

(A3) Axiom (A3) follows because as long as K_2 is nonempty there is some minimal formula in K_2 with respect to $\tilde{K_1}$.

(A4) Axiom (A4) follows from the third condition of generalized loyal assignments.

(A5) Assume that (A5) is false. Then for some μ_1 formula (1) $\mu_1 \in Min(K_2, \leq_{\tilde{K}_1}) \cap K_3$ and (2) $\mu_1 \notin Min(K_2 \cap K_3, \leq_{\tilde{K}_1})$.

Then $\mu_1 \in K_3$ and $\mu_1 \in Min(K_2, \leq_{\tilde{K}_1})$ must hold. By the definition of minimal models $\mu_1 \in K_2$ must also hold. Therefore, $\mu_1 \in K_2 \cap K_3$ must hold. That and assumption (2) imply there must be another formula μ_2 such that $\mu_2 \in K_2 \cap K_3$ and $\mu_2 <_{\tilde{K}_1} \mu_1$. Note that $\mu_2 \in K_2$ must also hold and this and $\mu_2 <_{\tilde{K}_1} \mu_1$ imply that $\mu_1 \notin Min(K_2, \leq_{\tilde{K}_1})$, which is a contradiction.

(A6) Assume that (A6) is false. Then $(\tilde{K_1} \triangle K_2) \cap K_3$ must be nonempty and for some μ_1 formula (1) $\mu_1 \notin Min(K_2, \leq_{\tilde{K}_1}) \cap K_3$ and (2) $\mu_1 \in Min(K_2 \cap K_3, \leq_{\tilde{K}_1})$.

Since we assume that $(\tilde{K_1} \triangle K_2) \cap K_3$ is nonempty, there must be a formula μ_2 such that $\mu_2 \in Min(K_2, \leq_{\tilde{K}_1}) \cap K_3$. Then $\mu_2 \in K_2 \cap K_3$ must hold. By assumption (2) and because $\leq_{\tilde{K}_1}$ is total, $\mu_1 \leq_{\tilde{K}_1} \mu_2$ holds. Since $\mu_2 \in Min(K_2, \leq_{\tilde{K}_1}) \cap K_3$, also $\mu_2 \in Min(K_2, \leq_{\tilde{K}_1})$ must hold. We claim that $\mu_1 \in Min(K_2, \leq_{\tilde{K}_1})$ must also hold. Suppose that it does not hold. Then there must be a formula $\mu_3 \in K_2$ such that $\mu_3 <_{\tilde{K}_1} \mu_1$, i.e., $\mu_3 \leq_{\tilde{K}_1} \mu_1$ holds and $\mu_1 \leq_{\tilde{K}_1} \mu_3$ does not hold. Since μ_2 is minimal model in K_2, $\mu_2 \leq_{\tilde{K}_1} \mu_3$ holds. Then by transitivity of $\mu_1 \leq_{\tilde{K}_1} \mu_2$ and $\mu_2 \leq_{\tilde{K}_1} \mu_3$ we have that $\mu_1 \leq_{\tilde{K}_1} \mu_3$, which is a contradiction. Hence $\mu_1 \in Min(K_2, \leq_{\tilde{K}_1})$ must hold.

From assumption (2) and the definition of minimal models, it is clear that $\mu_1 \in K_2$ and $\mu_1 \in K_3$. The latter and assumption (1) imply that $\mu_1 \notin Min(K_2, \leq_{\tilde{K}_1})$. This is a contradiction to the claim that we proved in the previous paragraph. Hence assumptions (1-2) were wrong.

(A7) If μ_1 is a minimal formula in K_2 according to both $\tilde{K_1}$ and $\tilde{K_3}$, then for any μ_2 in K_2 both $\mu_1 \leq_{\tilde{K}_1} \mu_2$ and $\mu_1 \leq_{\tilde{K}_3} \mu_2$ must be true. Then by the fourth condition of generalized loyal assignments μ_1 is also minimal in $\tilde{K_1} \uplus \tilde{K_3}$. This implies that axiom (A7) also holds.

(A8) Suppose that μ_1 is both $\leq_{\tilde{K}_1}$ and $\leq_{\tilde{K}_3}$ minimal in K_2 and that axiom (A8) does not hold. Then there is some formula μ_2 that is $\leq_{\tilde{K}_1 \uplus \tilde{K}_3}$ minimal in K_2 but w.l.g. not $\leq_{\tilde{K}_1}$ minimal in K_2. Then $\mu_1 <_{\tilde{K}_1} \mu_2$ and $\mu_1 \leq_{\tilde{K}_3} \mu_2$. Then by the fifth condition of generalized loyal assignments $\mu_1 <_{\tilde{K}_1 \uplus \tilde{K}_3} \mu_2$. Hence μ_2 cannot be $\leq_{\tilde{K}_1 \uplus \tilde{K}_3}$ minimal in K_2, which is a contradiction. Hence (A8) holds.

□