

On the Semantics of Theory Change: Arbitration between Old and New Information*

Peter Z. Revesz

Department of Computer Science and Engineering
University of Nebraska–Lincoln

Abstract: Katsuno and Mendelzon divide theory change, the problem of adding new information to a logical theory, into two types: revision and update. We propose a third type of theory change: arbitration. The key idea is the following: the new information is considered neither better nor worse than the old information represented by the logical theory. The new information is simply one voice against a set of others already incorporated into the logical theory. From this follows that arbitration should be commutative. First we define arbitration by a set of postulates and then describe a model-theoretic characterization of arbitration for the case of propositional logical theories. We also study weighted arbitration where different models of a theory can have different weights.

1 Introduction

The problem of updating logical theories is a common fundamental concern to databases, to Artificial Intelligence [McC68, Rei92], and to belief revision [Mak85, Gär88]. It is well-

*Part of the research presented here was done while the author was a postdoctoral fellow at the University of Toronto and was supported by a grant from the Canadian Institute for Robotics and Intelligent Systems. For correspondence write to revesz@cse.unl.edu.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

ACM-PODS-5/93/Washington, D.C.

© 1993 ACM 0-89791-593-3/93/0005/0071...\$1.50

known that giving semantics to updates is difficult once the logical theory is complex enough to express views or integrity constraints [BS81, FUV83].

For a simple example from [GMR92] consider the propositional database represented by the theory $\{A, B, A \wedge B \rightarrow C\}$. If we want to add the information that proposition C is false, then the resulting theory could be $\{A, A \wedge B \rightarrow C, \neg C\}$, $\{B, A \wedge B \rightarrow C, \neg C\}$, or $\{A, B, \neg C\}$ or several others that leave the theory consistent.

In general we may treat each database as some logical theory T . Then the database update problem translates into the following: What should be the result of changing theory T with some sentence ϕ ?

We may try to answer this question directly by proposing a concrete operator for theory change. Or we may answer indirectly by giving a set of axioms or postulates that every adequate theory change operator should be expected to satisfy. Taking the first direction concrete operators were proposed by Borgida [Bor85], Dalal [Dal88], Fagin, Ullman and Vardi [FUV83], Grahne, Mendelzon and Revesz [GMR92], Satoh [Sat88], Weber [Web86], and Winslett [Win88]. Taking the second direction, sets of postulates were proposed by Alchourrón, Gärdenfors and Makinson [AGM85] and by Katsuno and Mendelzon [KM91].

Katsuno and Mendelzon point out that there are no universally desirable set of postulates. In particular they argue that the AGM postulates describe one type of theory change –*revision*– while the KM postulates describe another –*update*. An informal distinction between these two types was also made before by Abiteboul and Grahne [AG85] and by Keller and Winslett [KW85]. This paper introduces a third type of theory change –*arbitration*– and shows it to be distinct from both revision and update.

The three types of theory change assume quite different relations between the present theory and the new information. As a simple illustration of these relations consider a jury during a trial. As each witness tells his/her story, the jury has to change its theory of the crime.

Revision assumes that the new information is always more important and reliable than the present theory. This may happen if a prosecutor organizes a set of witnesses from the least reliable to the most reliable. For example, the first witness may be a distant relative, while the second a close relative of the defendant. The distant relative may say that the defendant was a social drinker. The close relative may say that the defendant was an alcoholic. Then the jury needs revision.

Update assumes that the new information is always more recent than the present theory. This may happen if the prosecutor organizes the witnesses chronologically. For example, the first witness may claim that the defendant bought a gun in January. The second witness may say that the defendant sold the gun in February. Then the jury needs update.

Arbitration assumes no asymmetry between the present theory and the new information. The witnesses may be all contemporary and equally important. For example, members of a crowd of witnesses of a brawl may

be all equally important and contemporary. However, their version of the events may conflict in some details. For example, nine witnesses may claim that A started the fight, but two witnesses may claim that it was B. In this case the jury needs arbitration. Arbitration alone is interested in the question: how can the jury reach a consensus?

The intuitive idea for any theory change is to modify the present theory *as little as possible* to accommodate the new information. Due to the different assumptions described above, the three types of theory change operators define least change in fundamentally different ways.

Katsuno and Mendelzon 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.

We will define arbitration in terms of an auxiliary operation called *model-fitting*. Model-fitting operators can be characterized as those operators that select from the models of the new information the overall closest models to the whole set of models of T .

We are also going to consider a generalization of arbitration when the models of the logical theory describing the database can have different weights. These weights are distinct from the priority values of Fagin et al. [FUV83], which are assigned to formulas, not to models. They are also different from the weights of Dalal [Dal88] which are assigned to propositions. They also have only vague connection with probabilities and possibility values [Zad78] which both have a constant one as an upper bound, unlike weights.

The research described here has much practical relevance to databases. We believe that the three types of theory change operations complement each other. Not all database updates are alike. Some are true updates (in the sense of [KM91]) some are not. Especially promising as an application area for arbitration are large heterogeneous databases, which often require merging of large equally important sets of information to answer queries. It certainly seems beneficial if database users are provided with more options to choose from to suit their needs.

The outline of the paper is the following. Section 2 lists some basic definitions. Section 3 defines by postulates the operations of arbitration and model-fitting. This section also shows a model-theoretic characterization of these operators for the case when the database is a propositional theory. Section 4 considers the case when the models of the database can have different weights. Section 5 concludes with 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 the formula describing our present information about some problem the *knowledge base*. If our information is composed of a set of formulas, then we take their conjunction as the knowledge base.

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:

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

$$\begin{aligned} Mod(\psi \vee \phi) &= Mod(\psi) \cup Mod(\phi) \\ Mod(\psi \wedge \phi) &= Mod(\psi) \cap Mod(\phi) \end{aligned}$$

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 S of \mathcal{M} with respect to a pre-order \leq_ψ is defined as:

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

Katsuno and Mendelzon gave the following model-theoretic characterization of revision and update when the knowledge base is a propositional theory. 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_ψ of interpretations for closeness to ψ , where the pre-order \leq_ψ 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 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_ψ 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_ψ minimal models of μ . Hence by Katsuno and Mendelzon's characterization above, Dalal's operator is a true revision operator.

3 Arbitration and Model-Fitting

This section gives first a formal definition of the set of model-fitting operations and a model-theoretic characterization of it. A definition of arbitration is given in terms of model-fitting. It is then shown that the three types of theory change operators are disjoint.

We say that a theory change operator \triangleright is a model-fitting operator if and only if \triangleright satisfies the following axioms:

- (A1) $\psi \triangleright \mu$ implies μ .
- (A2) If ψ is unsatisfiable then $\psi \triangleright \mu$ is unsatisfiable.
- (A3) If both ψ and μ are satisfiable then $\psi \triangleright \mu$ is also satisfiable.
- (A4) If $\psi_1 \leftrightarrow \psi_2$ and $\mu_1 \leftrightarrow \mu_2$ then $\psi_1 \triangleright \mu_1 \leftrightarrow \psi_2 \triangleright \mu_2$.
- (A5) $(\psi \triangleright \mu) \wedge \phi$ implies $\psi \triangleright (\mu \wedge \phi)$.
- (A6) If $(\psi \triangleright \mu) \wedge \phi$ is satisfiable then $\psi \triangleright (\mu \wedge \phi)$ implies $(\psi \triangleright \mu) \wedge \phi$.
- (A7) $(\psi_1 \triangleright \mu) \wedge (\psi_2 \triangleright \mu)$ implies $(\psi_1 \vee \psi_2) \triangleright \mu$.

- (A8) If $(\psi_1 \triangleright \mu) \wedge (\psi_2 \triangleright \mu)$ is satisfiable then $(\psi_1 \vee \psi_2) \triangleright \mu$ implies $(\psi_1 \triangleright \mu) \wedge (\psi_2 \triangleright \mu)$.

Here axioms (A1) and (A3-A5) are the same as the axioms (U1) and (U3-U5). (See Appendix A for the list of revision and update postulates.) Axiom (A6) is the same as axiom (R6). Axioms (A2), (A7) and (A8) are new. Axiom (A2) says that no model can be fitted to an unsatisfiable knowledge base. Axioms (A7) and (A8) guarantee that the models in the knowledge-base are considered together for overall closeness. They express the following property: the overall closest models to $\psi_1 \vee \psi_2$ in μ are exactly the intersection of the overall closest models to ψ_1 in μ and the overall closest models to ψ_2 in μ if the intersection is nonempty.

The next theorem presents a model-theoretic characterization of model-fitting operators that satisfy axioms (A1-A8). The main idea is to define for each knowledge-base ψ a relation that orders interpretations in \mathcal{M} with respect to their closeness to ψ . This can be done as follows.

A *loyal* assignment is a function that assigns for each knowledge-base ψ a pre-order \leq_ψ such that the following three conditions hold. For each $I, J \in \mathcal{M}$ and knowledge bases ψ_1, ψ_2 :

- (1) If $\psi_1 \leftrightarrow \psi_2$ then $\leq_{\psi_1} = \leq_{\psi_2}$.
- (2) If $I <_{\psi_1} J$ and $I \leq_{\psi_2} J$ then $I <_{\psi_1 \vee \psi_2} J$.
- (3) If $I \leq_{\psi_1} J$ and $I \leq_{\psi_2} J$ then $I \leq_{\psi_1 \vee \psi_2} J$.

Using these definitions, the characterization theorem can now be stated as follows.

Theorem 3.1 A theory change operator satisfies conditions (A1-A8) if and only if there exists a loyal assignment that maps each knowledge-base ψ to a total pre-order \leq_ψ such that $Mod(\psi \triangleright \mu) = Min(Mod(\mu), \leq_\psi)$. \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 ψ and an interpretation I as follows:

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

Then we assign to each knowledge base ψ the total pre-order \leq_ψ defined by $I \leq_\psi J$ if and only if $odist(\psi, I) \leq odist(\psi, J)$. Clearly this is a loyal assignment. Hence by Theorem 3.1 this operator satisfies axioms (A1-A8) and is a proper model-fitting operator.

Example 3.1 As an application of model-fitting consider a database class with three students. The instructor considers teaching either Datalog only or both SQL and Datalog. This can be represented as $\mu = (\neg S \wedge D) \vee (S \wedge D)$. The three students in order would like to learn SQL only, would like Datalog only, and would like SQL, Datalog and Query-by-Example. That is the students suggest to the instructor to teach $\psi = (S \wedge \neg D \wedge \neg Q) \vee (\neg S \wedge D \wedge \neg Q) \vee (S \wedge D \wedge Q)$.

Considering only the propositional terms S , D , and Q , $Mod(\mu)$ contains only $\{D\}$ and $\{S, D\}$, while $Mod(\psi)$ contains only $\{S\}$, $\{D\}$, and $\{S, D, Q\}$. We calculate that $odist(\psi, \{D\}) = 2$ and $odist(\psi, \{S, D\}) = 1$. Hence we find that $Mod(\psi \triangleright \mu) = \{S, D\}$. This indicates that the instructor could best satisfy the class by teaching both SQL and Datalog. \square

Example 3.1 is a situation which calls for arbitration instead of revision. Note that if the instructor decides to teach Datalog only—which would be suggested by a revision operator like Dalal's—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.

Arbitration can be defined as a special case of model-fitting:

$$\psi \triangle \phi = (\psi \vee \phi) \triangleright \mathcal{M}$$

where \mathcal{M} is the set of all interpretations. Therefore arbitration means finding the best fit interpretations to both the information ψ and the information ϕ . For example, if the instructor in Example 3.1 were willing to teach any combination of SQL, Datalog, and Query-by-Example, then he/she would be doing arbitration.

Corollary 3.1 A theory change operator is an arbitration operator if and only if there exists a loyal assignment that maps each knowledge-base ψ to a total pre-order \leq_ψ such that $Mod(\psi \triangle \phi) = Min(\mathcal{M}, \leq_{\psi \vee \phi})$. \square

An interesting question is whether the three types of theory change operators are disjoint. As Katsuno and Mendelzon [KM92] point out all update operators are *monotone*. That is, for any update operator if ϕ implies ψ , then $\phi \diamond \mu$ implies $\psi \diamond \mu$. However, Gärdenfors' impossibility theorem shows that no non-trivial revision operator can be monotone (see [Gär88] for the result and the logical definition of non-trivial). From these it follows that the set of non-trivial revision and the set of update operators are disjoint. This result can be strengthened by dropping the provision of non-triviality. In general, we have that:

Theorem 3.2 The set of revision, update, and model-fitting operators are pairwise disjoint. In particular, there is no theory change operator that satisfies both (R2) and (A8), or all of (U2), (U8) and (A8), or all of (R1), (R2), (R3) and (U8). \square

Theorem 3.2 has several implications. As Katsuno and Mendelzon [KM91, KM92] showed, the operators of Borgida [Bor85], Dalal [Dal88], Fagin et al. [FUV83], Satoh [Sat88], and Weber [Web86] each satisfy axiom (R2), and Winslett's operator [Win88], simplified for the propositional case, satisfies (U2) and (U8). Hence none of these operators can be model-fitting operators.

4 Weighted Arbitration

In this section we generalize the results of the previous section by considering weighted knowledge bases. A weighted knowledge base is a function $\tilde{\psi}$ from interpretations to nonnegative real numbers. The real numbers are intended to describe the relative degree of importance of interpretations 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 ψ into a weighted knowledge base $\tilde{\psi}$ having for all interpretations I $\tilde{\psi}(I) = 0$ if $I \notin \text{Mod}(\psi)$ and $\tilde{\psi}(I) = 1$ if $I \in \text{Mod}(\psi)$.

The model of a weighted knowledge base $\tilde{\psi}$ is a function that can be defined as follows.

$$\begin{aligned} \text{Mod}(\tilde{\psi}) &= \tilde{\psi} \\ \text{Mod}(\tilde{\psi} \vee \tilde{\phi}) &= \text{Mod}(\tilde{\psi}) \uplus \text{Mod}(\tilde{\phi}) \\ \text{Mod}(\tilde{\psi} \wedge \tilde{\phi}) &= \text{Mod}(\tilde{\psi}) \sqcap \text{Mod}(\tilde{\phi}) \end{aligned}$$

where \uplus takes for each interpretation I the sum, and \sqcap takes the minimum of the weights of I in $\text{Mod}(\tilde{\psi})$ and $\text{Mod}(\tilde{\phi})$ and assigns it as the weight of I in $\text{Mod}(\tilde{\psi} \vee \tilde{\phi})$ and $\text{Mod}(\tilde{\psi} \wedge \tilde{\phi})$ respectively.

This generalizes the model for knowledge bases and leaves a clear distinction between syntax and semantics. We could have $\psi \neq \phi$ and $\text{Mod}(\psi) = \text{Mod}(\phi)$. Similarly, we now can have $\tilde{\psi} \neq \tilde{\phi}$ and $\text{Mod}(\tilde{\psi}) = \text{Mod}(\tilde{\phi})$.

We say that a weighted knowledge base $\tilde{\psi}$ is unsatisfiable if and only if for all I $\text{Mod}(\tilde{\psi})(I) = 0$. A weighted knowledge base is satisfiable if and only if it is not unsatisfiable. We say that a weighted knowledge base $\tilde{\psi}$ implies another weighted knowledge base $\tilde{\phi}$, written as $\tilde{\psi} \rightarrow \tilde{\phi}$, if and only if for all I $\text{Mod}(\tilde{\psi})(I) \leq \text{Mod}(\tilde{\phi})(I)$.

A *weighted loyal assignment* is a function that assigns for each weighted knowledge base $\tilde{\psi}$ a total pre-order $\leq_{\tilde{\psi}}$ such that the following conditions hold. For each interpretation I and J and weighted knowledge base $\tilde{\psi}$ and $\tilde{\phi}$:

- (1) If $\tilde{\psi} \leftrightarrow \tilde{\phi}$ then $\leq_{\tilde{\psi}} = \leq_{\tilde{\phi}}$.
- (2) If $I <_{\tilde{\psi}} J$ and $I \leq_{\tilde{\phi}} J$ then $I <_{\tilde{\psi} \vee \tilde{\phi}} J$.
- (3) If $I \leq_{\tilde{\psi}} J$ and $I \leq_{\tilde{\phi}} J$ then $I \leq_{\tilde{\psi} \wedge \tilde{\phi}} J$.

Let $\tilde{\mu}$ be a weighted knowledge base and let $S = \{I : \tilde{\mu}(I) > 0\}$. The set of *minimal models* of $\tilde{\mu}$ with respect to a pre-order $\leq_{\tilde{\psi}}$ is defined as:

$$\text{Min}(\text{Mod}(\tilde{\mu}), \leq_{\tilde{\psi}}) =$$

$$\begin{aligned} \{\tilde{\phi} : \text{if } I \in \text{Min}(S, \leq_{\tilde{\psi}}) \text{ then } \tilde{\phi}(I) = \tilde{\mu}(I) \\ \text{else } \tilde{\phi}(I) = 0\} \end{aligned}$$

We rewrite axioms (A1-A8) into axioms (F1-F8) by simply replacing regular knowledge bases by weighted knowledge bases. We say that a theory change operator is a weighted model-fitting operator if and only if it satisfies axioms (F1-F8). The following theorem is a model-theoretic characterization of weighted model-fitting.

Theorem 4.1 A theory change operator satisfies conditions (F1-F8) if and only if there exists a weighted loyal assignment that maps each weighted knowledge-base $\tilde{\psi}$ to a total pre-order $\leq_{\tilde{\psi}}$ such that $\text{Mod}(\tilde{\psi} \triangleright \tilde{\mu}) = \text{Min}(\text{Mod}(\tilde{\mu}), \leq_{\tilde{\psi}})$. \square

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

$$wdist(\tilde{\psi}, I) = \sum_{J \in \mathcal{M}} dist(I, J) * \tilde{\psi}(J)$$

Next we define for each weighted knowledge base $\tilde{\psi}$ the total pre-order $\leq_{\tilde{\psi}}$ such that $I \leq_{\tilde{\psi}} J$ if and only if $wdist(\tilde{\psi}, I) \leq wdist(\tilde{\psi}, J)$. Clearly this is a weighted loyal assignment. Hence 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 weighted knowledge base $\tilde{\mu}$ with $\tilde{\mu}(\{D\}) = \tilde{\mu}(\{S, D\}) = 1$ and $\tilde{\mu}(I) = 0$ for any other interpretation I .

Suppose that 10 students would like to learn SQL only, 20 would like Datalog only, and 5 would like SQL, Datalog, and Query-by-Example. The students' requests can be represented as $\tilde{\psi}(\{S\}) = 10$, $\tilde{\psi}(\{D\}) = 20$, and $\tilde{\psi}(\{S, D, Q\}) = 5$. Now we calculate that $wdist(\tilde{\psi}(\{D\}), I) = 30$ and $wdist(\tilde{\psi}(\{S, D\}), I) = 35$. Hence $Mod(\tilde{\psi} \triangleright \tilde{\mu})$ will be a knowledge base $\tilde{\phi}$ with $\tilde{\phi}(\{D\}) = 1$ and $\tilde{\phi}(I) = 0$ for any other interpretation I . This indicates that in this case the instructor could best satisfy the class by teaching Datalog only. \square

Note that in the case of weighted arbitration 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.

Weighted arbitration can be defined as a special case of weighted model-fitting:

$$\tilde{\psi} \Delta \tilde{\phi} = (\tilde{\psi} \vee \tilde{\phi}) \triangleright \tilde{\mathcal{M}}$$

where $\tilde{\mathcal{M}}$ is the weighted knowledge base that has $\tilde{\mathcal{M}}(I) = 1$ for all $I \in \mathcal{M}$. Weighted arbitration can be characterized by the following.

Corollary 4.1 A theory change operator is a weighted arbitration operator if and only if there exists a weighted loyal assignment that maps each weighted knowledge base $\tilde{\psi}$ to a total pre-order $\leq_{\tilde{\psi}}$ such that $Mod(\tilde{\psi} \Delta \tilde{\phi}) = Min(\tilde{\mathcal{M}}, \leq_{\tilde{\psi} \vee \tilde{\phi}})$. \square

5 Open Problems

An open problem is to extend arbitration from propositional to first-order, similarly perhaps to the first order update language in [GMR92]. 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].

Acknowledgements: I thank Alex Borgida, Mariano Consens, 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*, 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) 510–530, 1985.
- [BS81] F. Bancilhon & N. Spyratos. Update semantics of relational views. *ACM Transactions on Database Systems*, (4) 557–575, 1981.
- [Bor85] A. Borgida. Language features for flexible handling of exceptions in information systems. *ACM Transaction on Database Systems*, 10, pages 563–603, 1985.
- [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.
- [Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge, MA, 1988.
- [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.
- [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. On the difference between updating a knowledge base and revising it. *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 387–394, 1991.
- [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. manuscript.
- [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. *Synthese*, (62) 347–363, 1985
- [Rei92] R. Reiter. On specifying database updates. *Proceedings of the Third International Conference on Extending Database Technology*, 1992.
- [Rei78] R. Reiter. On closed world databases. In: *Logic and Databases*, H. Gallaire & J. Minker, editors. Plenum Press, New York, pages 55–76, 1978.
- [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*, pages 3–28, 1978.

A Revision and Update Axioms

A set of axioms, called the AGM-postulates, for the class of revision operators was presented by Alchourrón, Gärdenfors and Makinson [AGM85]. These are axioms that every revision operator has to satisfy. As Katsuno and Mendelzon [KM91] showed, the AGM postulates can be rewritten into the following equivalent set of axioms on propositional knowledge-bases.

- (R1) $\psi \circ \mu$ implies μ .
- (R2) If $\psi \wedge \mu$ is satisfiable then $\psi \circ \mu \leftrightarrow \psi \wedge \mu$.
- (R3) If μ is satisfiable then $\psi \circ \mu$ is also satisfiable.
- (R4) If $\psi_1 \leftrightarrow \psi_2$ and $\mu_1 \leftrightarrow \mu_2$ then $\psi_1 \circ \mu_1 \leftrightarrow \psi_2 \circ \mu_2$.
- (R5) $(\psi \circ \mu) \wedge \phi$ implies $\psi \circ (\mu \wedge \phi)$.
- (R6) If $(\psi \circ \mu) \wedge \phi$ is satisfiable then $\psi \circ (\mu \wedge \phi)$ implies $(\psi \circ \mu) \wedge \phi$.

Axiom (R1) assures that the new knowledge will hold in the revised knowledge-base. Axiom (R2) 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 (R3) assures that no unwarranted inconsistency will be introduced. Axiom (R4) 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 form of those sentences. This rule is called Dalal's Principle of Irrelevance of Syntax. Axioms (R5) and (R6) assure that the set of the models of 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.

For the class of update operators Katsuno and Mendelzon [KM92] present the following

postulates:

- (U1) $\psi \diamond \mu$ implies μ .
- (U2) If ψ implies μ then $\psi \diamond \mu$ is equivalent to ψ .
- (U3) If both ψ and μ are satisfiable then $\psi \diamond \mu$ is also satisfiable.
- (U4) If $\psi_1 \leftrightarrow \psi_2$ and $\mu_1 \leftrightarrow \mu_2$ then $\psi_1 \diamond \mu_1 \leftrightarrow \psi_2 \diamond \mu_2$.
- (U5) $(\psi \diamond \mu) \wedge \phi$ implies $\psi \diamond (\mu \wedge \phi)$.
- (U6) If $\psi \diamond \mu_1$ implies μ_2 and $\psi \diamond \mu_2$ implies μ_1 then $\psi \diamond \mu_1 \leftrightarrow \psi \diamond \mu_2$.
- (U7) If ψ is a singleton¹ then $(\psi \diamond \mu_1) \wedge (\psi \diamond \mu_2)$ implies $\psi \diamond (\mu_1 \vee \mu_2)$.
- (U8) $(\psi_1 \vee \psi_2) \diamond \mu \leftrightarrow (\psi_1 \diamond \mu) \vee (\psi_2 \diamond \mu)$.

Note that axioms (U1) and (U4-U5) are the same as axioms (R1) and (R4-R5). Axiom (U2) is a weakening of axiom (R2) in the case when ψ is satisfiable. Axiom (U3) is a weakening of axiom (R3) that is needed to avoid defining the update of an empty knowledge-base. Axioms (U6-U7) replace axiom (R6). They generalize (R6) slightly by admitting orderings where some pair of models of the new information are not comparable as to closeness to the knowledge-base. Axiom (U8) guarantees that each model in the knowledge-base is updated independently. Katsuno and Mendelson [KM92] prove that a simplified version of Winslett's operator satisfies all the KM-postulates, and Grahne et al. [GMR92] do the same for their update operator.

B Proofs

Proof of Theorem 3.1: (Only-if) In this proof let $form(I_1, \dots, I_k)$ denote the formula that has exactly the models I_1, \dots, I_k . We define a loyal assignment as follows. For each knowledge-base ψ we define a total pre-order \leq_ψ in terms of the \triangleright operator as follows.

¹A propositional sentence ψ is singleton if there is exactly one interpretation in $Mod(\psi)$.

For each (not necessary distinct) pair I, J of models, let $I \leq_\psi J$ if and only if $I \in Mod(\psi \triangleright form(I, J))$.

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

- (1) We need to show that \leq_ψ is total, reflexive, and transitive when ψ is satisfiable.

total When ψ is satisfiable, by axioms (A1) and (A3) the $Mod(\psi \triangleright form(I, J))$ is a nonempty subset of $\{I, J\}$. Hence any pair of models are comparable, making \leq_ψ a total relation.

reflexive When ψ is satisfiable, by axioms (A1) and (A3) the $Mod(\psi \triangleright form(I))$ is a nonempty subset of $\{I\}$. Hence \leq_ψ is also reflexive.

transitive Assume that ψ is satisfiable and that the relation \leq_ψ is not transitive, that is, for some I, J , and K models $I \leq_\psi J$, $J \leq_\psi K$, and $I \not\leq_\psi K$.

Then by the definition of \leq_ψ , $I \notin Mod(\psi \triangleright form(I, K))$. By axiom (A5), $I \notin Mod(\psi \triangleright form(I, J, K)) \wedge form(I, K)$. Hence $I \notin Mod(\psi \triangleright form(I, J, K))$. There are two possible cases. Either (i) $J \in Mod(\psi \triangleright form(I, J, K))$ or (ii) J is not in $Mod(\psi \triangleright form(I, J, K))$.

In case (i), we know that I is not in $Mod(\psi \triangleright form(I, J, K)) \wedge form(I, J)$ and that $Mod(\psi \triangleright form(I, J, K)) \wedge form(I, J)$ is satisfiable. Then by (A6) also $I \notin Mod(\psi \triangleright form(I, J))$. This contradicts the assumption that $I \leq_\psi J$.

In case (ii), by (A1) and (A3) we know that $K = Mod(\psi \triangleright form(I, J, K))$. Hence $Mod(\psi \triangleright form(I, J, K)) \wedge form(J, K)$ is satisfiable but does not contain J . Hence

by (A6) also $J \notin \text{Mod}(\psi \triangleright \text{form}(J, K))$. This contradicts the assumption that $J \leq_\psi K$.

(2) The first condition of loyalty follows from axiom (A4). To show the second condition, assume that $I <_{\psi_1} J$ and $I \leq_{\psi_2} J$. Then I is and J is not in $\text{Mod}(\psi_1 \triangleright \text{form}(I, J))$, and I is also in $\text{Mod}(\psi_2 \triangleright \text{form}(I, J))$. Hence $I = \text{Mod}((\psi_1 \triangleright \text{form}(I, J)) \wedge (\psi_2 \triangleright \text{form}(I, J)))$. Then by (A7) and (A8) also $I = \text{Mod}((\psi_1 \vee \psi_2) \triangleright \text{form}(I, J))$. Then by the definition of assignments $I <_{\psi_1 \vee \psi_2} J$. The third condition of loyalty can be shown similarly to the second condition.

(3) We need to show both the \subseteq and the \supseteq directions. If either ψ or μ are unsatisfiable, then $\text{Mod}(\psi \triangleright \mu) = \emptyset = \text{Min}(\text{Mod}(\mu), \leq_\psi)$. Hence assume that both are satisfiable.

(\subseteq) Assume that $I \in \text{Mod}(\psi \triangleright \mu)$ and $I \notin \text{Min}(\text{Mod}(\mu), \leq_\psi)$. By (A1) $I \in \text{Mod}(\mu)$. Since I is not a minimal model, according to the definition of minimal there must be another model J in $\text{Mod}(\mu)$ such that $J <_\psi I$ (i.e., such that $J \leq_\psi I$ and $I \not\leq_\psi J$). By the definition of \leq_ψ then $J \in \text{Mod}(\psi \triangleright \text{form}(I, J))$ and $I \notin \text{Mod}(\psi \triangleright \text{form}(I, J))$.

Since both I and J are in $\text{Mod}(\mu)$, $\mu \wedge \text{form}(I, J) = \text{form}(I, J)$. Hence I is also not in $\text{Mod}(\psi \triangleright (\mu \wedge \text{form}(I, J)))$. By (A5) and using $\phi = \text{form}(I, J)$ we know that $\text{Mod}((\psi \triangleright \mu) \wedge \text{form}(I, J))$ implies $\text{Mod}(\psi \triangleright (\mu \wedge \text{form}(I, J)))$. Hence also $I \notin \text{Mod}((\psi \triangleright \mu) \wedge \text{form}(I, J))$. Therefore, I cannot be in $\text{Mod}(\psi \triangleright \mu)$, which is a contradiction.

(\supseteq) Assume now that $I \notin \text{Mod}(\psi \triangleright \mu)$ and $I \in \text{Min}(\text{Mod}(\mu), \leq_\psi)$. By the definition of minimal, $I \in \text{Mod}(\mu)$. Since both ψ and μ are satisfiable, by (A3) there is some model J in $\text{Mod}(\psi \triangleright \mu)$,

and by (A1) also $J \in \text{Mod}(\mu)$. Since both I and J are in $\text{Mod}(\mu)$, $\mu \wedge \text{form}(I, J) = \text{form}(I, J)$. Hence by (A5) and (A6) and letting ϕ be $\text{form}(I, J)$ we get that $\text{Mod}((\psi \triangleright \mu) \wedge \text{form}(I, J)) = \text{Mod}(\psi \triangleright \mu) \cap \{I, J\} = \text{Mod}(\psi \triangleright \text{form}(I, J))$. Since both ψ and ϕ are satisfiable, by (A1) and (A3), $\text{Mod}(\psi \triangleright \text{form}(I, J))$ is a nonempty subset of $\{I, J\}$. But the identity above and $I \notin \text{Mod}(\psi \triangleright \mu)$ implies that also $I \notin \text{Mod}(\psi \triangleright \text{form}(I, J))$. Hence $J = \text{Mod}(\psi \triangleright \text{form}(I, J))$. Therefore $J <_\psi I$. Hence I cannot be a minimal model according to \leq_ψ , i.e., $I \notin \text{Min}(\text{Mod}(\mu), \leq_\psi)$. This is again a contradiction.

(If) Assume that for a theory change operator \triangleright there is a loyal function that assigns to each satisfiable knowledge-base ψ a total pre-order \leq_ψ such that $\text{Mod}(\psi \triangleright \mu) = \text{Min}(\text{Mod}(\mu), \leq_\psi)$. We need to show that \triangleright satisfies axioms (A1-A8).

(A1) Axiom (A1) follows because the minimal model of μ with respect to any total pre-order is always by definition some subset of $\text{Mod}(\mu)$.

(A2) Axiom (A2) follows because if ψ is unsatisfiable, then the minimal model with respect to ψ is the emptyset. Hence $\psi \triangleright \mu$ is also unsatisfiable.

(A3) Axiom (A3) follows because as long as ψ and μ are satisfiable there is some minimal model in $\text{Mod}(\mu)$ with respect to ψ .

(A4) Axiom (A4) follows from the first condition of loyalty.

(A5-A6) Let ϕ be any formula. We have that $\text{Mod}((\psi \triangleright \mu) \wedge \phi) = \text{Min}(\text{Mod}(\mu), \leq_\psi) \cap \text{Mod}(\phi)$ and that $\text{Mod}(\psi \triangleright (\mu \wedge \phi)) = \text{Min}(\text{Mod}(\mu \wedge \phi), \leq_\psi)$.

Note that for any total pre-order \leq_ψ and any model I , if there is nothing closer than I within $\text{Mod}(\mu)$ to ψ then there is

nothing closer than I within $Mod(\mu \wedge \phi)$ to ψ . Hence if I is a minimal model of $Mod(\mu)$ then I is also a minimal model of $Mod(\mu \wedge \phi)$ with respect to ψ . Hence $Mod((\psi \triangleright \mu) \wedge \phi)$ implies $Mod(\psi \triangleright (\mu \wedge \phi))$ proving axiom (A5).

Suppose that $(\psi \triangleright \mu) \wedge \phi$ is satisfiable. We want to show the reverse of the above, that is, “if I is a minimal model of $Mod(\mu \wedge \phi)$ then I is also a minimal model of $Mod(\mu)$ ”. Since $(\psi \triangleright \mu) \wedge \phi$ is satisfiable, there is a minimal model J in $Mod(\mu)$ that is also in $Mod(\phi)$. Therefore every minimal model of $Mod(\mu \wedge \phi)$ must be a minimal model of $Mod(\mu)$. Hence if $(\psi \triangleright \mu) \wedge \phi$ is satisfiable then $Mod(\psi \triangleright (\mu \wedge \phi))$ implies $Mod((\psi \triangleright \mu) \wedge \phi)$ proving axiom (A6).

(A7) If I is minimal in $Mod(\mu)$ according to both ψ_1 and ψ_2 , then for any J in $Mod(\mu)$ it must be that $(I <_{\psi_1} J$ and $I \leq_{\psi_2} J)$ or $(I \leq_{\psi_1} J$ and $I <_{\psi_2} J)$ or $(I \leq_{\psi_1} J$ and $I \leq_{\psi_2} J)$ is true. Then by the second and third conditions of loyalty I is also minimal in $Mod(\psi_1 \vee \psi_2)$. This implies that axiom (A7) also holds.

(A8) Suppose that I is both \leq_{ψ_1} and \leq_{ψ_2} minimal in $Mod(\mu)$ and that axiom (A8) does not hold. Then there is some model J that is $\leq_{\psi_1 \vee \psi_2}$ minimal in $Mod(\mu)$ but w.l.g. not \leq_{ψ_1} minimal in $Mod(\mu)$. Then $I <_{\psi_1} J$ and $I \leq_{\psi_2} J$. Then by the second condition of loyalty $I <_{\psi_1 \vee \psi_2} J$. Hence J cannot be $\leq_{\psi_1 \vee \psi_2}$ minimal in $Mod(\mu)$, which is a contradiction. Hence (A8) holds.

□

Proof of Theorem 3.2: Assume that there is a theory change operator $*$ that satisfies both (R2) and (A8). Let m_1 and m_2 be any two singletons. Let $\psi_1 = m_1 \vee m_2$, $\psi_2 = m_2$, and $\mu = m_1 \vee m_2$. Then $(\psi_1 \vee \psi_2) \wedge \mu = m_1 \vee m_2$. Hence by (R2), $(\psi_1 \vee \psi_2) * \mu = m_1 \vee m_2$.

But $\psi_1 \wedge \mu = m_1 \vee m_2$. Hence by (R2),

$\psi_1 * \mu = m_1 \vee m_2$. Similarly, $\psi_2 \wedge \mu = m_2$ and by (R2), $\psi_2 * \mu = m_2$. Taking conjunctions we have that $(\psi_1 * \mu) \wedge (\psi_2 * \mu) = (m_1 \vee m_2) \wedge m_2 = m_2$. Hence by (A8), $(\psi_1 \vee \psi_2) * \mu$ implies m_2 . This is a contradiction.

Assume now that there is a theory change operator $*$ that satisfies all of (U2), (U8), and (A8). Let ψ_1, ψ_2 , and μ be as above. Then note that ψ_1 implies μ , and ψ_2 implies μ . Hence by (U2), $\psi_1 * \mu = \psi_1 = m_1 \vee m_2$, and $\psi_2 * \mu = \psi_2 = m_2$. Hence by (U8), $(\psi_1 \vee \psi_2) * \mu = (\psi_1 * \mu) \vee (\psi_2 * \mu) = (m_1 \vee m_2) \vee m_2 = m_1 \vee m_2$.

But $(\psi_1 * \mu) \vee (\psi_2 * \mu) = (m_1 \vee m_2) \wedge m_2 = m_2$. Hence by (A8) $(\psi_1 \vee \psi_2) * \mu$ implies m_2 . This is again a contradiction.

Assume now that there is a theory change operator $*$ that satisfies all of (R1), (R2), (R3), and (U8). Let m_1, m_2 , and m_3 be any three singletons. Let $\psi_1 = m_1$ and $\mu = m_2 \vee m_3$. By (R1) $\psi_1 * \mu$ implies μ , and by (R3) $\psi_1 * \mu$ is satisfiable. Hence without loss of generality either (i) $\psi_1 * \mu = m_2 \vee m_3$, or (ii) $\psi_1 * \mu = m_3$ must be true. Let $\psi_2 = m_2$. Since $\psi_2 \wedge \mu = m_2 \wedge (m_2 \vee m_3) = m_2$, by (R2) $\psi_2 * \mu = m_2$. Similarly, since $(\psi_1 \vee \psi_2) \wedge \mu = (m_1 \vee m_2) \wedge (m_2 \vee m_3) = m_2$, by (R2) $(\psi_1 \vee \psi_2) * \mu = m_2$.

But, by (U8), $(\psi_1 \vee \psi_2) * \mu = (\psi_1 * \mu) \vee (\psi_2 * \mu) = (\psi_1 * \mu) \vee m_2$. Hence, in both cases (i) and (ii) above, $(\psi_1 \vee \psi_2) * \mu = m_2 \vee m_3$. This is again a contradiction. □