On the Difference between Updating a Knowledge Base and Revising it[1]

HIROFUMI KATSUNO

NTT Basic Research Laboratories
Musashino-shi, Tokyo 180, Japan

ALBERTO O. MENDELZON

Computer Systems Research Institute
University of Toronto
Toronto, Canada M5S 1A4

1. INTRODUCTION

Consider a knowledge base represented by a theory ψ of some logic, say propositional logic. We want to incorporate into ψ a new fact, represented by a sentence μ of the same language. What should the resulting theory be? A growing body of work (Dalal 1988, Katsuno and Mendelzon 1989, Nebel 1989, Rao and Foo 1989) takes as a departure point the *rationality postulates* proposed by Alchourrón, Gärdenfors and Makinson (1985). These are rules that every adequate revision operator should be expected to satisfy. For example: the new fact μ must be a consequence of the revised knowledge base.

In this paper, we argue that no such set of postulates will be adequate for every application. In particular, we make a fundamental distinction between two kinds of modifications to a knowledge base. The first one, *update*, consists of bringing the knowledge base up to date when the world described by it changes. For example, most database updates are of this variety, e.g. "increase Joe's salary by 5%". Another example is the incorporation into the knowledge base of changes caused in the world by the actions of a robot (Ginsberg and Smith 1987, Winslett 1988, Winslett 1990). We show that the AGM postulates must be drastically modified to describe update.

The second type of modification, *revision*, is used when we are obtaining new information about a static world. For example, we may be trying to diagnose a faulty circuit and want to incorporate into the knowledge base the results of successive tests, where newer results may contradict old ones. We claim the AGM postulates describe only revision.

---

[1]A preliminary version of this paper was presented at the Second International Conference on Principles of Knowledge Representation and Reasoning, Cambridge, Mass., 1991.

The distinction between update and revision was made by Keller and Winslett (1985) in the context of extended relational databases. They distinguished *change-recording updates* (which we call updates) and *knowledge-adding updates* (which we call revisions). Our work extends theirs in several ways. We formalize the distinction, which they made informally. We provide an axiomatization for update obtained from the AGM. Keller and Winslett's work does not treat inconsistent knowledge-bases or addition of inconsistent knowledge, while ours does. And we treat arbitrary propositional knowledge-bases, while their setting is relational databases extended with null values and disjunction.

Gärdenfors (1988) considers two types of revision functions in the context of probabilistic reasoning: *imaging* and *conditionalization*. We can regard imaging as a probabilistic version of update, and conditionalization as a probabilistic version of revision.

Morreau in this volume also recognizes the distinction between update and revision, and shows how to use update in planning.

Rao and Foo (1989) extend the AGM postulates in order to apply them to reasoning about action. They introduce the notion of time and consider a modal logic. However, they do not identify the difference between revision and update. In this paper we clarify exactly why the postulates apply to revision but not to update. We give a new set of postulates that apply to update operators, and characterize the set of operators that satisfy the postulates in terms of a set of partial orders defined among possible worlds.

The difference between the postulates for revision and for update can be explained intuitively as follows. Suppose knowledge base ψ is to be revised with sentence μ. Revision methods that satisfy the AGM postulates are exactly those that select from the models of μ those that are "closest" to models of ψ, where the notion of closeness is defined by an ordering relationship among models that satisfies certain conditions (Katsuno and Mendelzon, 1989). The models selected determine the new theory, which we denote by $\psi \circ \mu$. On the other hand, update methods select, for each model M of the knowledge base ψ, the set of models of μ that are closest to M. The new theory describes the union of all such models. Suppose that ψ has exactly two models, I and J; that is, there are two possible worlds described by the knowledge base. Suppose that μ describes exactly two worlds, K and L, and that K is "closer" to I than L is, and K is also closer to I than L is to J. Then K is selected for the new knowledge base, but L is not. Note the knowledge base has effectively forgotten that J used to be a possible world; the new fact μ has been used as evidence for the retroactive impossibility of J. That is, not only do we refuse to have J as a model of the new knowledge base, but we also conclude that J should not have been in the old knowledge base to begin with.

If we are doing revisions, this behaviour is rational. Since the real world has not changed, and μ has to be true in all the new possible worlds, we can forget about some of the old possible worlds on the grounds that they are too different from what we now know to be the case. On the other hand, suppose we are doing updates. The models of ψ are possible worlds; we think one of them is the real world, but we do not know which one. Now the real world has changed; we examine each of the old possible worlds and find the minimal way of changing each one of them so that it becomes a model of μ. The fact that the real world has changed gives us no grounds to conclude that some of the old worlds were actually not possible.

To illustrate this distinction between update and revision, let us consider two examples which are formally identical to the one above but have different intuitively desirable results. First, in the spirit of Ginsberg and Smith (1987) and Winslett (1988), suppose our knowledge base describes five objects A,B,C,D,E inside a room. There is a table in the room, and objects may be on or off the table. The sentence a means "object A is on the table," and similarly for sentences b,c,d, and e. The knowledge base ψ is the sentence

$$(a \wedge \neg b \wedge \neg c \wedge \neg d \wedge \neg e) \vee (\neg a \wedge \neg b \wedge c \wedge d \wedge e).$$

That is, either object A is on the table by itself, or objects C,D and E are. This knowledge base has exactly two models I and J. We send a robot into the room, instructing it to achieve a situation in which all or none of the objects are on the table. This change can be modelled by incorporating the following sentence μ:

$$(a \wedge b \wedge c \wedge d \wedge e) \vee (\neg a \wedge \neg b \wedge \neg c \wedge \neg d \wedge \neg e).$$

Let us take Dalal's notion of "closeness" and the revision operator that results Dalal (1988). According to this measure, the distance between two models is simply the number of propositional letters on which they differ. The models selected for the new KB will be those models of μ which are at minimal distance from models of ψ. Now K, the model where nothing is on the table, is at distance 1 from I (the model where A is on the table) and at distance 3 from J (the model where C,D and E are). On the other hand L, the model where every object is on the table, is at distance 4 from I and 2 from J. Dalal's revision operator will therefore select K as the only model of the new knowledge base. But intuitively, it seems clear that this is incorrect. After the robot is done, all we know is that either all objects are on the table or all are off; there is no reason to conclude that they are all off, which is what revision does.

Consider now an example that is formally identical, but where the desired result is given by revision, not by update. Suppose the knowledge base describes the state of a five bit register which we read through noisy communication lines. Each of the propositional letters a,b,c,d,e now represents one bit. The state of the register is unchanging. Two different readings have been obtained: 10000 and 00111. By an independent analysis of the circuits that control the register, we learn that all bits must have the same value. That is, only 11111 and 00000 are possible patterns.

Dalal's revision method tells us to keep 00000 as the new knowledge base; that is, we conclude that 00111 is relatively too far from the possible patterns to be an acceptable result. It might be argued that it is better to forget the two readings in the KB and just keep both 00000 and 11111 as possible worlds. However, consider an example in which the register is thousands of bits long, the two readings agree on every bit except the first five, and the new fact only says that the first five bits must be all 0's or all 1's. It is clearly a waste of information now to discard the old KB and just keep the new fact.

With this motivation, let us postulate that an update method should give each of the old possible worlds equal consideration. One way of capturing this condition syntactically is to require that the result of updating $\psi \vee \phi$ with μ be equivalent to the disjunction of ψ updated with μ and ϕ updated with μ. Let us call this the *disjunction rule*. This rule turns out to have far-reaching consequences. In particular, consider the case where μ is consistent with ψ, that is, no conflict exists. The AGM postulates require the result of the revision to be simply the conjunction of ψ and μ. As we will see, this apparently obvious requirement is inconsistent with the disjunction rule.

The outline of the paper is as follows. In Section 2 we give preliminaries. We review the AGM postulates and our characterization from Katsuno and Mendelzon (1989) of all revision methods that satisfy the postulates in terms of a pre-order among models. In Section 3 we define the update operation and give a set of rationality postulates for it. We show that these postulates characterize all update operators that select for each model I of ψ those models of μ that are "closest" to I in a certain sense. In Section 4, we discuss briefly how update and revision could be combined for reasoning about action. In Section 5, we propose a new operation called *erasure*. Erasure is the analogue of contraction (Alchourrón, Gärdenfors and Makinson 1985, Makinson 1985) for update operators. We show that Winslett's *Forget* operator is a special case of *symmetric erasure*, an operator defined in terms of erasure. Finally, in Section 6 we sketch a way to unify update and revision by using a single theory change operator parameterized by time.

2. PRELIMINARIES

Throughout this paper, we consider a finitary propositional language L, and we denote the finite set consisting of all the propositional letters in L by Ξ. We represent a knowledge base by a propositional formula ψ, since we need a finite fixed representation of a KB to store it in a computer. An *interpretation* of L is a function from Ξ to $\{\mathtt{T}, \mathtt{F}\}$. A *model* of a propositional formula ψ is an interpretation that makes ψ true in the usual sense. $Mod(\psi)$ denotes the set of all the models of ψ. The knowledge base ψ may be *inconsistent*, in which case $Mod(\psi) = \emptyset$. A propositional formula ϕ is *complete* if for any propositional formula, μ, ϕ implies μ or ϕ implies $\neg\mu$.

2.1. Revision and the AGM Postulates

Given a knowledge base ψ and a sentence μ, $\psi \circ \mu$ denotes the *revision* of ψ by μ; that is, the new knowledge base obtained by adding new knowledge μ to the old knowledge base ψ. Note that other papers in this volume use the symbol $\dot{+}$ to denote revision. We, however, use $\circ$ instead of $\dot{+}$ to clarify that we represent a KB by a propositional formula ψ, while other papers use a (possibly infinite) set K of formulas.

Alchourrón, Gärdenfors and Makinson propose eight postulates, (G*1)$\sim$(G*8), which they argue must be satisfied by any reasonable revision function. By specializing to the case of propositional logic and rephrasing them in terms of finite covers for infinite "knowledge sets," the postulates become the six rules below. See (Katsuno and Mendelzon 1989, 1990) for a discussion of the intuitive meaning and formal properties of these rules.

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

2.2. Orders between Interpretations

The postulates (R5) and (R6) represent the condition that revision be accomplished with minimal change. In Katsuno and Mendelzon (1989), we gave a model theoretic characterization of minimal change.

Let $\mathcal{I}$ be the set of all the interpretations of L. A *pre-order* $\leq$ over $\mathcal{I}$ is a reflexive and transitive relation on $\mathcal{I}$. We define $<$ as $I < I'$ if and only if $I \leq I'$ and $I' \not\leq I$. A pre-order is *total* if for every $I, J \in \mathcal{I}$, either $I \leq J$ or $J \leq I$. Consider a function that assigns to each propositional formula ψ a pre-order $\leq_\psi$ over $\mathcal{I}$. We say this assignment is *faithful*[2] if the following three conditions hold:

1. If $I, I' \in Mod(\psi)$ then $I <_\psi I'$ does not hold.
2. If $I \in Mod(\psi)$ and $I' \notin Mod(\psi)$ then $I <_\psi I'$ holds.

[2]The term *persistent* was used instead of "faithful" in Katsuno and Mendelzon (1989).

3. If $\psi \leftrightarrow \phi$, then $\leq_\psi = \leq_\phi$.

That is, a model of ψ cannot be strictly less than any other model of ψ and must be strictly less than any non-model of ψ.

Let $\mathcal{M}$ be a subset of $\mathcal{I}$. An interpretation I is minimal in $\mathcal{M}$ with respect to $\leq_\psi$ if $I \in \mathcal{M}$ and there is no $I' \in \mathcal{M}$ such that $I' <_\psi I$. Let $Min(\mathcal{M}, \leq_\psi)$ be the set of all $I \in \mathcal{M}$ such that I is minimal in $\mathcal{M}$ with respect to $\leq_\psi$. The following characterization of all revision operators that satisfy the postulates was established in Katsuno and Mendelzon (1989).

Theorem 2.1 *Revision operator $\circ$ satisfies Conditions (R1)$\sim$(R6) if and only if there exists a faithful assignment that maps each KB ψ to a total pre-order $\leq_\psi$ such that $Mod(\psi \circ \mu) = Min(Mod(\mu), \leq_\psi)$.*

3. UPDATE

In this section we axiomatize all update operators that can be defined by partial orders or partial pre-orders over interpretations. The class of operators defined generalizes Winslett's *Possible Models Approach* (PMA) (Winslett 1988, 1989). Winslett argues that the PMA is suitable for reasoning about action in certain applications. According to our classification, the PMA is an update operator, because it changes each possible world independently. For background, we review this approach first.

3.1. Possible Models Approach

Let ψ be a KB and μ a new sentence. We denote the PMA operator by $\diamond_{pma}$. For each model I of ψ, the PMA selects from the models of μ those which are "closest" to I. The models of the new KB $(\psi \diamond_{pma} \mu)$ are the union of these selected models. Formally, the PMA is defined by

$$Mod(\psi \diamond_{pma} \mu) = \bigcup_{I \in Mod(\psi)} \text{Incorporate}(Mod(\mu), I),$$

where Incorporate$(Mod(\mu), I)$ is the set of models that are "closest" to I in $Mod(\mu)$.

The closeness between two interpretations, I and J is measured by the set $Diff(I, J)$ of propositional letters that have different truth values under I and J. For two interpretations, J_1 and J_2, J_1 is closer to I than J_2 (denoted by $J_1 \leq_{I,pma} J_2$) if and only if $Diff(I, J_1)$ is a subset of $Diff(I, J_2)$. Then, Incorporate$(Mod(\mu), I)$ is the set of all the minimal elements with respect to $\leq_{I,pma}$ in the set $Mod(\mu)$, that is, $Min(Mod(\mu), \leq_{I,pma})$.

Example 3.1 Let L have only two propositional letters, b and m. Let $\psi \leftrightarrow (b \wedge \neg m) \vee (\neg b \wedge m)$ and $\mu \leftrightarrow b$. Then, $I = \langle F, T \rangle$ is a model of ψ. $J_1 = \langle T, T \rangle$ and $J_2 = \langle T, F \rangle$ are two models of μ. $J_1 \leq_{I,pma} J_2$ follows from the fact $\mathit{Diff}(I, J_1) = \{b\}$ is a subset of $\mathit{Diff}(I, J_2) = \{b, m\}$. Similarly, by considering the case where J_2 is a model of ψ, we obtain $\psi \diamond_{pma} \mu \leftrightarrow b$.

To interpret this example in the context of (Winslett 1988, 1989), let us go back to a room with two objects in it, a book and a magazine. Suppose b means the book is on the floor, and m means the magazine is on the floor. Then, ψ states that either the book is on the floor or the magazine is, but not both. Now, we order a robot to put the book on the floor. The result of this action should be represented by the update of ψ with b. After the robot puts the book on the floor, all we know is b, and this is in fact the result of appying the PMA. Note that ψ is consistent with μ. According to revision postulate (R2), the result of $\psi \circ \mu$ should therefore be $\psi \wedge \mu$, that is, $b \wedge \neg m$. But why should we conclude that the magazine is not on the floor?

3.2. Postulates for Update

The PMA is defined in terms of a certain partial order over interpretations. This subsection generalizes the PMA by axiomatizing all update operators that can be defined by partial orders or partial pre-orders over interpretations.

We use $\psi \diamond \mu$ to denote the result of updating KB ψ with sentence μ. Our postulates for update are:

(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 complete then $(\psi \diamond \mu_1) \wedge (\psi \diamond \mu_2)$ implies $\psi \diamond (\mu_1 \vee \mu_2)$.

(U8) $(\psi_1 \vee \psi_2) \diamond \mu \leftrightarrow (\psi_1 \diamond \mu) \vee (\psi_2 \diamond \mu)$.

Postulates (U1)$\sim$(U5) correspond directly to the the corresponding postulates for revision given in Section 2. Note that postulate (U2) says that if a new sentence μ is derivable from KB ψ, then updating by μ does not influence the KB. In the case where ψ is consistent, postulate (U2) is strictly weaker than (R2). An immediate consequence of (U2) is the following.

Lemma 3.1 *If an update operator $\diamond$ satisfies (U2), and ψ is inconsistent, then $\psi \diamond \mu$ is inconsistent for any μ.*

The property above might appear undesirable: once an inconsistency is introduced in the knowledge base, there is no way to eliminate it. However, all we are saying is there is no way to eliminate it *by using update*. For example, revision does not have this behaviour; in fact, (R3) guarantees that the result of a revision is consistent provided that the new sentence introduced is itself consistent. This is another manifestation of the difference between update and revision. An inconsistent knowledge base is the result of an inadequate theory, and can be remedied with revision (or contraction) by adding new knowledge that supersedes the inconsistency (or removing contradictory knowledge using contraction). We can never repair an inconsistent theory using update, because update specifies a change in the world. If there is no set of worlds that fits our current description, we have no way of recording the change in the real world.

We drop rule (R6), and add instead three new postulates, (U6)~(U8). (U6) says that if updating a knowledge base with μ_1 guarantees μ_2, and updating the same knowledge base with μ_2 guarantees μ_1, then the two updates have the same effect. This is similar to condition (C7) in Gardenfors's analysis of minimal changes of belief Gärdenfors (1978) and to conditional logic axiom *CSO* in Nute (1984). (U7) applies only to complete KB's, in which there is no uncertainty over what are the possible worlds. If some possible world results from updating a complete KB with μ_1 and it also results from updating it with μ_2, then this possible world must also result from updating the KB with $\mu_1 \vee \mu_2$. Finally, (U8) is what we called the "disjunction rule" in the Introduction. It guarantees that each possible world of the KB is given independent consideration. (U8) can be regarded as a nonprobabilistic version of the *homomorphic* condition about probabilistic revision functions in Gärdenfors (1988).

The following lemma shows that we can obtain one direction of (R2) by using (U2) and (U8)

Lemma 3.2 *If an update operator $\diamond$ satisfies (U2) and (U8), then $\psi \wedge \mu$ implies $\psi \diamond \mu$.*

Proof. Since ψ is equivalent to $(\psi \wedge \mu) \vee (\psi \wedge \neg\mu)$, it follows from (U8) that $\psi \diamond \mu$ is equivalent to $((\psi \wedge \mu) \diamond \mu) \vee ((\psi \wedge \neg\mu) \diamond \mu)$. By (U2), $(\psi \wedge \mu) \diamond \mu$ is equivalent to $\psi \wedge \mu$. Hence, $\psi \wedge \mu$ implies $\psi \diamond \mu$.

However, as Example 3.1 showed, update operators do not necessarily satisfy that $\psi \diamond \mu$ implies $\psi \wedge \mu$ when ψ is consistent with μ.

An interesting consequence of the postulates is *monotonicity.*

Lemma 3.3 *If an update operator $\diamond$ satisfies (U8), and ϕ implies ψ, then $\phi \diamond \mu$ implies $\psi \diamond \mu$.*

Monotonicity has been deemed undesirable by the philosophers of theory revision. The reason is a result called "Gärdenfors's impossibility theorem" (Arló Costa 1989, Gärdenfors 1988, Makinson 1989), which shows that monotonicity is incompatible with postulates (R1)$\sim$(R4). More precisely, Theorem 7.10 of Gärdenfors (1988) implies that there is no non-trivial revision operator that satisfies monotonicity and (R1)–(R4). Since update operators do not satisfy (R2), this result does not apply to update.

Gärdenfors's motivation in studying this problem is to use theory revision to define the conditional connective used in counterfactual reasoning. The idea is to use the *Ramsey Test*: interpret the conditional statement "given the state of the world described by ψ, if μ were true, then η would also be true" as $\psi \circ \mu$ implies η. Intuitively, it would seem that this kind of statement is better modelled by using update instead of revision in the Ramsey Test. This intuition, together with the immunity of updates to Gärdenfors's result, suggest further study of the connection between updates and conditional reasoning may be fruitful. Preliminary results are reported by Katsuno and Saoth (1991) and Grahne (1991).

We can now formalize a notion of closeness between models that generalizes the particular measure used in the PMA. Instead of associating each KB with an ordering, let us consider a function that maps each interpretation I to a partial pre-order $\leq_I$. We say that this assignment is *faithful* if the following condition holds:

- For any $J \in \mathcal{I}$, if $I \neq J$ then $I <_I J$.

The following theorem shows that the postulates exactly capture all update operators defined by a partial pre-order. It turns out that the classes of operators defined by partial orders and partial pre-orders are the same.

Theorem 3.4 *Let $\diamond$ be an update operator. The following conditions are equivalent:*

1. *The update operator $\diamond$ satisfies Conditions (U1)$\sim$(U8).*
2. *There exists a faithful assignment that maps each interpretation I to a partial pre-order $\leq_I$ such that*

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

3. *There exists a persistent assignment that maps each interpretation I to a partial order $\leq_I$ such that*

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

We give a proof sketch here. A detailed proof can be found in the Appendix.

Proof Sketch. ($1 \Rightarrow 2$) We assign to each interpretation I a relation $\leq_I$ defined as follows. For any interpretations J and J', $J \leq_I J'$ if and only if either $J = I$ or $Mod(form(I) \diamond form(J, J')) = \{J\}$. We verify that Conditions (U1)$\sim$(U8) imply that this mapping is a faithful assignment such that

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

($2 \Rightarrow 3$) For a pre-order $\leq_I$, we define a relation $\leq'_I$ as $J \leq'_I J'$ if and only if $J = J'$ or $J <_I J'$. It is easy to show that $\leq'_I$ is a partial order and that $J <_I J'$ if and only if $J <'_I J'$. Hence, Statement 3 follows from Statement 2 by changing $\leq_I$ to $\leq'_I$.

($3 \Rightarrow 1$) Assume that there is a faithful assignment mapping each interpretation I to a partial order $\leq_I$. We define an update operator $\diamond$ by

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

We show that the update operator $\diamond$ satisfies (U1)$\sim$(U8).

Comparing this result with Theorem 2.1, we see two differences between revision and update from a model-theoretic point of view. First, Theorem 3.1 refers to *partial* preorders while Theorem 2.1 uses *total* preorders. It turns out that a version of the revision postulates that accommodates partial preorders can be given, and we show this in Katsuno and Mendelzon (1990). It is also possible to design a class of update operators based on total pre-orders. If we replace (U6) and (U7) by postulate (U9) below, then we can prove the total pre-order analogue of Theorem 3.1. The proof is similar to that of Theorem 3.1, by defining, for any two interpretations J and J', $J \leq_I J'$ if and only if either $J = I$ or $J \in Mod(form(I) \diamond form(J, J'))$.

(U9) If ψ is complete and $(\psi \diamond \mu) \wedge \phi$ is satisfiable then $\psi \diamond (\mu \wedge \phi)$ implies $(\psi \diamond \mu) \wedge \phi$.

It is worth pointing out that a total preorder associated with interpretation I is what Lewis (1973) calls *a system of spheres centered at I*. Systems of spheres play a central role in the semantics of Lewis's conditional logic; this brings up again the suggested connection between updates and conditional logic, which is explored further by Grahne (1991).

The second and more important difference between revision and update is that, in the case of update, a different ordering is induced by each model of ψ, while for revision, only one ordering is induced by the whole of ψ. This "local" behaviour of update, contrasted with the "global" behaviour of revision, is essential to the difference between the two operators.

4. REASONING ABOUT ACTION

For the purposes of reasoning about action, the usual approach is to represent a particular action as a pair of a precondition and a postcondition. The precondition for the action encodes what the world must be like in order for the action to be executable. The postcondition describes the immediate consequences resulting from the action. Any update operator that satisfies our postulates can be used for reasoning about action by regarding postconditions for an action as new knowledge and by assuming that preconditions for the action are satisfied by the current KB. That is, the effect on KB ψ of performing action with precondition α and postcondition β will be ψ if ψ does not imply α, and $\psi \diamond \beta$ otherwise. Winslett (1989) discusses how the frame, qualification and ramification problems are handled by this approach.[3]

Let us extend this idea by examining more closely what happens when ψ does not satisfy the precondition α. Presumably, the robot will return and report one of two outcomes: either α was true, and the action was carried out, or α failed and the action was not carried out. If we want a more elaborate model, we can also allow other outcomes, such as: α was true but the action could not be carried out for other reasons, or α could not be either verified or falsified. In each case, we can take advantage of the distinction between revision and update to incorporate into ψ all the information gained by the robot. For example, if the action was carried out, we can change the KB to $(\psi \circ \alpha) \diamond \beta$. If the precondition was found false, we use $\psi \circ \neg\alpha$. If the truth value of the precondition could not be determined, we use $\psi \bullet \alpha$ (contraction is discussed in the next section).

5. CONTRACTION AND ERASURE

Contraction is a change of belief or knowledge state induced by the loss of confidence in some sentence. For example, if we believed that a paper was written by Turing, but new evidence has cast doubt on this belief, we contract the corresponding sentence from our knowledge base.

Alchourrón *et al.* (1985) proposed rationality postulates for contraction. We denote by $\psi \bullet \mu$ a new knowledge base obtained from an old knowledge base ψ by contracting μ. The postulates for contraction, rephrased in our terms, are as follows.

(C1) ψ implies $\psi \bullet \mu$.

(C2) If ψ does not imply μ then $\psi \bullet \mu$ is equivalent to ψ.

(C3) If μ is not a tautology then $\psi \bullet \mu$ does not imply μ.

[3]Actually, Winslett uses for this purpose a variant of the PMA that orders interpretations in a way similar to the partial pre-order used in prioritized circumscription. Such variants are included in our class of update operators.

(C4) If $\psi_1 \leftrightarrow \psi_2$ and $\mu_1 \leftrightarrow \mu_2$ then $\psi_1 \bullet \mu_1 \leftrightarrow \psi_2 \bullet \mu_2$.

(C5) $(\psi \bullet \mu) \wedge \mu$ implies ψ.

Alchourrón *et al.* (1985) showed that contraction and revision are closely related: they proved that, given a revision operator $\circ$ that satisfies (R1)$\sim$(R4), if we define a contraction operator $\bullet$ by

$$\psi \bullet \mu \leftrightarrow \psi \vee (\psi \circ \neg\mu)$$

then the operator $\bullet$ satisfies (C1)$\sim$(C5). Conversely, given a contraction operator $\bullet$ that satisfies (C1)$\sim$(C4), if we define a revision operator $\circ$ by

$$\psi \circ \mu \leftrightarrow (\psi \bullet \neg\mu) \wedge \mu$$

then the operator $\circ$ satisfies (R1)$\sim$(R4).

We propose a new operator, *erasure*, which is to contraction as update is to revision. Erasing sentence μ from ψ means adding models to ψ; for each model I, we add all those models closest to I in which μ is false. Intuitively, erasing μ means the world may have changed in such a way that μ is not true. In contrast, contracting μ means our description of the set of possible worlds must be adjusted to the possibility of μ being false.

The erasure operator $\blacklozenge$ for a given update operator $\diamond$ is defined by

$$\psi \blacklozenge \mu \leftrightarrow \psi \vee (\psi \diamond \neg\mu) \qquad (U \rightarrow E).$$

This erasure operator satisfies the following postulates (E1)$\sim$(E5) and (E8) if the update operator satisfies (U1)$\sim$(U4) and (U8).

(E1) ψ implies $\psi \blacklozenge \mu$.

(E2) If ψ implies $\neg\mu$ then $\psi \blacklozenge \mu$ is equivalent to ψ.

(E3) If ψ is satisfiable and μ is not a tautology then $\psi \blacklozenge \mu$ does not imply μ.

(E4) If $\psi_1 \leftrightarrow \psi_2$ and $\mu_1 \leftrightarrow \mu_2$ then $\psi_1 \blacklozenge \mu_1 \leftrightarrow \psi_2 \blacklozenge \mu_2$.

(E5) $(\psi \blacklozenge \mu) \wedge \mu$ implies ψ.

(E8) $(\psi_1 \vee \psi_2) \blacklozenge \mu$ is equivalent to $(\psi_1 \blacklozenge \mu) \vee (\psi_2 \blacklozenge \mu)$.

There are two differences between contraction and erasure in terms of postulates. One is that (E2) is weaker than (C2); since contraction of a sentence μ does not influence a KB ψ if ψ does not imply μ, but erasure of μ might modify ψ if ψ does not imply $\neg\mu$. The other one is that erasure needs the disjunctive rule (E8), but contraction does not.

Example 5.1 Consider Example 3.1 again. Recall we have a room with two objects in it, a book and a magazine, b means the book is on the floor, and m means the magazine is on the floor. The knowledge base ψ states that either the book is on the floor or the magazine is, but not both. Suppose that a contraction operator $\bullet$ satisfies (C2). If we contract ψ by b then $\psi \bullet b$ is equivalent to ψ, since ψ does not imply b. This means that since the sentence that the book is on the floor is already questionable under ψ, contraction does not change ψ.

On the other hand, let an erasure operator $\blacklozenge$ be defined based on the PMA $\diamond_{pma}$. If we erase b from ψ then $\psi \blacklozenge b$ is equivalent to $(b \wedge \neg m) \vee \neg b$. This can be interpreted as follows. ψ represents two possible worlds, M_1 and M_2. In world M_1, the book is on the floor but the magazine is not. Since b holds in M_1, M_1 is altered to two worlds, M_1 itself and the world M_3 represented by $\neg b \wedge \neg m$, that is, neither the book nor the magazine is on the floor. In world M_2, the magazine is on the floor but the book is not. Since b does not hold in M_2, M_2 is retained as itself. Hence, $\psi \blacklozenge \mu$ represents the three worlds, M_1, M_2 and M_3.

The intuitive difference between contraction and erasure can be explained in this example as follows. Contracting b means nothing has changed in the room, but if the KB believes that the book is on the floor, make sure this belief is retracted. Since the KB has no such belief, the contraction has no effect. Erasing b means the state of the room has changed in such a way that, if the book was on the floor before, it has now been moved in an unpredictable way. This affects only those possible worlds in which the book was on the floor. The result is that we can no longer deduce anything about the location of the magazine from the fact that the book is not on the floor.

There is another operation which appears perhaps more natural than erasure. Suppose the state of the room has changed in such a way that the location of the book is now unpredictable, and we want to reflect this change in the knowledge base. We formalize this operation, called *symmetric erasure*, after the Theorem below.

The following theorem, proved in the Appendix, gives a correspondence between update and erasure similar to the correspondence between revision and contraction.

Theorem 5.1

1. *If an update operator $\diamond$ satisfies (U1)$\sim$(U4) and (U8), then the erasure operator $\blacklozenge$ defined by $(U \rightarrow E)$ satisfies (E1)$\sim$(E5) and (E8).*

2. *If an erasure operator $\blacklozenge$ satisfies (E1)$\sim$(E4) and (E8), then the update operator $\diamond$ defined by*

$$\psi \diamond \mu \leftrightarrow (\psi \blacklozenge \neg\mu) \wedge \mu \qquad (E \rightarrow U)$$

satisfies (U1)$\sim$(U4) and (U8).

3. *Suppose that an update operator $\diamond$ satisfies (U1)$\sim$(U4) and (U8). Then, we can define an erasure operator by $(U \rightarrow E)$. The update operator obtained from the erasure operator by $(E \rightarrow U)$ is equal to the original update operator $\diamond$.*

4. *Suppose that an erasure operator $\diamond$ satisfies (E1)$\sim$(E5) and (E8). Then, we can define an update operator by $(E \rightarrow U)$. The erasure operator obtained from the update operator by $(U \rightarrow E)$ is equal to the original erasure operator $\blacklozenge$.*

Winslett (1989) discusses an operator called *Forget*, which she compares with contraction. It turns out that *Forget*, given an update operator $\diamond$, is equivalent to

$$(\psi \diamond \mu) \vee (\psi \diamond \neg\mu).$$

We call this operator *symmetric erasure* because μ and its negation play the same role in its definition. The main difference between erasure and symmetric erasure is that erasure does not affect the possible worlds in which $\neg\mu$ holds, but symmetric erasure does. Going back to Example 5.1, the symmetric erasure of b from ψ reflects the fact that someone has picked up the book and unpredictably decided to place it on the floor or on the table. The result of this symmetric erasure is the knowledge base with no information, since there is nothing we can say about either the book or the magazine after this change.

We can show similar postulates for symmetric erasure to those for erasure, and prove a similar theorem to Theorem 5.1. A natural definition of *symmetric contraction* follows from the above discussions, and similar results can be shown for it. Gärdenfors (1981) defines an operator similar to symmetric contraction, which he calls *complete contraction*, and proposes to use it to model "even if" conditionals.

6. TIME, REVISION AND UPDATE

So far in this paper we have devoted our efforts to distinguishing update from revision. We would like now to suggest how they can be unified. The essential difference between revision and update is a temporal one: revision is a change to our description of a world that has not itself changed, while update is the incorporation into our world description of the fact that the world has changed. Suppose now that we make this hidden temporal parameter explicit in the knowledge base. That is, instead of just a theory, a knowledge base is now a pair $\langle \psi, t \rangle$ where ψ is a theory and t denotes a time instant. This is in the spirit of the situation calculus MaCarthy and Hayes (1969) and other temporal formalisms. It is not important for our purposes what exactly is the ontology of time, whether it is discrete or continuous, etc. For example, returning to our familiar book and magazine example, the knowledge base that says exactly one of them is on the table at 10am is $\langle (b \wedge \neg m) \vee (\neg b \wedge m), 10am \rangle$.

Instead of two distinct change operations, update and revision, let us introduce a single one called Tell(μ, t) where μ is the new formula to be incorporated and t is a time instant. The effect of applying Tell(μ, t) to a knowledge base is to *replace* the knowledge base with a new one that incorporates the sentence μ and has time parameter t, unless t is earlier than the KB's time. More precisely, we define the result of applying Tell(μ, t') to $\langle \psi, t \rangle$ as $\langle \psi \circ \mu, t \rangle$ if $t = t'$, and $\langle \psi \diamond \mu, t' \rangle$ if $t' > t$. For now, the result will be left undefined when $t' < t$. So, when we send the robot into the room to put the book on the table, and the robot returns at 10:05 reporting mission accomplished, we apply Tell$(b, 10{:}05am)$ to the KB. This behaves as an update, yielding $\langle b, 10{:}05am \rangle$ as a result. On the other hand, suppose the reason we knew there was exactly one object on the table was because of an aerial photograph taken at 10am from a high altitude. Further analysis of the photograph reveals that the object on the table was actually the book. We then apply the change Tell$(b, 10am)$, which behaves as revision, and obtain $\langle b \wedge \neg m, 10am \rangle$. Intuitively, it is now correct to conclude that the magazine is not on the table at time 10am.

This proposal relieves the user from the burden of deciding whether each change is a revision or an update, which become special cases of a more general operator parameterized by time. It also raises interesting questions that we cannot answer in this paper, but leave as topics for further research. For example, we did not define the meaning of Tell(μ, t') when t' is earlier than the KB time. An obvious generalization is to have not one pair of theory and instant, but a whole sequence of theories, one for each instant, and to allow changes to any past, present, or future KB. The next step would be to introduce *persistence*: if we know something is true at time t, and have no reason to believe it has changed, we assume it is still true at time $t+1$. We can then distinguish at each instant t between knowledge, that is, those sentences we have been told are true at time t, and defeasible knowledge, those that have been inferred by persistence from the past (or from the future). An appealing way of doing this is to define the set of worlds described by the KB at time $t+1$ as the result of *updating* all knowledge, defeasible or not, about instant t, with the non-defeasible knowledge at $t+1$. A symmetric construction can be used for supporting persistence from the future into the past. This approach will be elaborated in the future.

7. CONCLUSION

The distinction between update and revision is an important one, and it has been overlooked in the literature since it was pointed out by Keller and Winslett (1985). We have formalized this distinction and given a model-theoretic characterization of updates in terms of orderings among interpretations. We have defined and characterized *erasure*, which is to update as contraction is to revision.

Many problems remain to explore. The connection between updates and conditional logic is one being pursued by several researchers (Katsuno and Satoh 1991, Grahne 1991). Another is computational tractability of updates and erasures. For example, Grahne and Mendelzon (1991) show that by restricting the form of the knowledge base, PMA updates can be computed in time polynomial in the size of the knowledge base. A third is the combined use of different theory change operators–revision, contraction, update, erasure– in specific applications, as suggested in Section 4. A temporal framework that unifies these operators, as sketched in Section 5, may be the best way to do this.

Acknowledgements

We thank Ken Satoh for suggesting rule (U6). The second author thanks the Natural Sciences and Engineering Research Council of Canada and the Institute for Robotics and Intelligent Systems for their support. A preliminary version of this paper appeared in Katsuno and Mendelzon (1991)

Appendix

Theorem 3.1 *Let $\diamond$ be an update operator. The following conditions are equivalent:*

1. *The update operator $\diamond$ satisfies Conditions (U1)$\sim$(U8).*
2. *There exists a faithful assignment that maps each interpretation I to a partial pre-order $\leq_I$ such that*

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

3. *There exists a faithful assignment that maps each interpretation I to a partial order $\leq_I$ such that*

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

Proof. (1 $\Rightarrow$ 2) For any interpretations J and J' ($J = J'$ is permitted), we define a relation $\leq_I$ as $J \leq_I J'$ if and only if either $J = I$ or $Mod(form(I) \diamond form(J, J')) = \{J\}$.

We first show that $\leq_I$ is a pre-order. In order to show that $\leq_I$ is reflexive, we show $Mod(form(I) \diamond form(J)) = \{J\}$. It is obvious that the equation follows from (U1) and (U3).

We show that $\leq_I$ is transitive. Assume $J_1 \leq_I J_2$ and $J_2 \leq_I J_3$. Then, we obtain that $form(I) \diamond form(J_1, J_2) \leftrightarrow form(J_1)$ and $form(I) \diamond form(J_2, J_3) \leftrightarrow form(J_2)$. Let $\mu \leftrightarrow form(J_1, J_2, J_3)$. By (U5), $(form(I) \diamond \mu) \wedge form(J_2, J_3)$ implies $form(I) \diamond form(J_2, J_3)$. Since $form(I) \diamond form(J_2, J_3) \leftrightarrow form(J_2)$, J_3 is not a model of $form(I) \diamond \mu$. We can also obtain that J_2 is not a model of $form(I) \diamond \mu$ in a similar way by using (U5) and $form(I) \diamond form(J_1, J_2) \leftrightarrow form(J_1)$. Therefore, it follows from (U3) that $form(I) \diamond \mu$ is logically equivalent to $form(J_1)$. Thus, $form(I) \diamond \mu$ implies $form(J_1, J_3)$. On the other hand, it follows from (U1) that $\mathit{form}(I) \diamond \mathit{form}(J_1, J_3)$ implies μ. By (U6), we obtain that $\mathit{form}(I) \diamond \mu$ is logically equivalent to $\mathit{form}(I) \diamond \mathit{form}(J_1, J_3)$. Thus, $\mathit{form}(I) \diamond \mathit{form}(J_1, J_3)$ is logically equivalent to $\mathit{form}(J_1)$. Therefore, $I_1 \leq_I I_3$ holds.

It follows from (U2) that the assignment mapping each interpretation I to $\leq_I$ is faithful.

We show $\mathit{Mod}(\mathit{form}(I) \diamond \mu) = \mathit{Min}(\mathit{Mod}(\mu), \leq_I)$. If μ is inconsistent then $\mathit{Mod}(\mu)$ is empty and it also follows from (U1) that $\mathit{Mod}(\mathit{form}(I) \diamond \mu)$ is empty. Hence, the equation holds. So, we assume in the following that μ is consistent. Suppose that J is a model of $\mathit{form}(I) \diamond \mu$ and J is not minimal in $\mathit{Mod}(\mu)$ with respect to $\leq_I$. There is a model J' of $\mathit{Mod}(\mu)$ such that $J' <_I J$. By (U5), $(\mathit{form}(I) \diamond \mu) \wedge \mathit{form}(J, J')$ implies $\mathit{form}(I) \diamond \mathit{form}(J, J')$. Since $J' <_I J$, $\mathit{form}(I) \diamond \mathit{form}(J, J')$ is equivalent to $\mathit{form}(J')$. Hence, J is not a model of $(\mathit{form}(I) \diamond \mu) \wedge \mathit{form}(J, J')$. This contradicts the assumption that J is a model of $\mathit{form}(I) \diamond \mu$. Therefore, $\mathit{Mod}(\mathit{form}(I) \diamond \mu) \subset \mathit{Min}(\mathit{Mod}(\mu), \leq_I)$ holds.

We show the converse inclusion. Assume that J is minimal in $\mathit{Mod}(\mu)$ with respect to $\leq_I$. Let $\mathit{Mod}(\mu) = \{J_1, \ldots, J_k\}$. Note that μ is logically equivalent to

$$\mathit{form}(J, J_1) \vee \mathit{form}(J, J_2) \vee \ldots \vee \mathit{form}(J, J_k).$$

Also, since there is no $J_j \in \mathit{Mod}(\mu)$ such that $J_j <_I J$, it follows that

$$J \in \mathit{Mod}(\mathit{form}(I) \diamond \mathit{form}(J, J_j))$$

for every $J_j \in \mathit{Mod}(\mu)$. Hence, J is a model of

$$(\mathit{form}(I) \diamond \mathit{form}(J, J_1)) \wedge \ldots \wedge (\mathit{form}(I) \diamond \mathit{form}(J, J_k)).$$

By repeated applications of (U7), this implies J is a model of

$$\mathit{form}(I) \diamond (\mathit{form}(J, J_1) \vee \ldots \vee \mathit{form}(J, J_k)),$$

that is, $J \in \mathit{Mod}(\mathit{form}(I) \diamond \mu)$.

If ψ is consistent then it follows from (U8) that

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

If ψ is inconsistent then both sides of the above equation are empty, that is, the equation holds.

($2 \Rightarrow 3$) The proof of this part is shown in the main text.

($3 \Rightarrow 1$) Assume that there is a faithful assignment mapping each interpretation I to a partial order $\leq_I$. We define an update operator $\diamond$ by

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

We show that the update operator $\diamond$ satisfies (U1)$\sim$(U8). (U1), (U3), (U4) and (U8) are obvious. If ψ is inconsistent then (U2), (U5), (U6) and (U7) trivially hold. We assume in the following that ψ is consistent.

We show (U2). It follows from the definition of faithfulness that if I is a model of μ then $form(I) \diamond \mu$ is equivalent to $form(I)$. Hence, we obtain (U2) by using (U8).

We show (U5). If $(\psi \diamond \mu) \wedge \phi$ is inconsistent then (U5) holds trivially. Let J be a model of $(\psi \diamond \mu) \wedge \phi$. There is some model I of ψ such that J is minimal in $Mod(\mu)$ with respect to $\leq_I$. Since $Mod(\mu \wedge \phi)$ is a subset of $Mod(\mu)$ and J is a model of ϕ, J is minimal in $Mod(\psi \diamond (\mu \wedge \phi))$ with respect to $\leq_I$. Hence, $(\psi \diamond \mu) \wedge \phi$ implies $\psi \diamond (\mu \wedge \phi)$.

We show (U6). Suppose that $\psi \diamond \mu_1$ implies μ_2 and that $\psi \diamond \mu_2$ implies μ_1. Assume that J is a model of $\psi \diamond \mu_1$, but J is not a model of $\psi \diamond \mu_2$. Since $\psi \diamond \mu_1$ implies μ_2, J is a model of μ_2. Since we assume that J is not a model of $\psi \diamond \mu_2$, for each model I of ψ, there exists a model J_I of μ_2 such that $J_I <_I J$ and J_I is minimal in $Mod(\mu_2)$ with respect to $\leq_I$. Then, each J_I is a model of $\psi \diamond \mu_2$. Since $\psi \diamond \mu_2$ implies μ_1, J_I is also a model of μ_1. Hence, for any model I of ψ, J is not minimal in $Mod(\mu_1)$ with respect to $\leq_I$. This contradicts that J is a model of $\psi \diamond \mu_1$. Therefore, $\psi \diamond \mu_1$ implies $\psi \diamond \mu_2$. Similarly, we can obtain $\psi \diamond \mu_2$ implies $\psi \diamond \mu_1$.

We show (U7). Let ψ be complete. Then, there exists a model I of ψ such that ψ is equivalent to $form(I)$. Let J be a model of $(\psi \diamond \mu_1) \wedge (\psi \diamond \mu_2)$. Assume that J is not a model of $\psi \diamond (\mu_1 \vee \mu_2)$. Then, there is a model J' of $\psi \diamond (\mu_1 \vee \mu_2)$ such that $J' <_I J$. If J' is a model of μ_1, this contradicts the minimality of J in $Mod(\mu_1)$ with respect to $\leq_I$. If J' is a model of μ_2, this also contradicts the minimality of J in $Mod(\mu_2)$ with respect to $\leq_I$.

Theorem 5.1

1. *If an update operator $\diamond$ satisfies (U1)$\sim$(U4) and (U8), then the erasure operator $\blacklozenge$ defined by $(U \rightarrow E)$ satisfies (E1)$\sim$(E5) and (E8).*

2. *If an erasure operator $\blacklozenge$ satisfies (E1)$\sim$(E4) and (E8), then the update operator $\diamond$ defined by*
$$\psi \diamond \mu \leftrightarrow (\psi \blacklozenge \neg\mu) \wedge \mu \qquad (E \rightarrow U)$$
satisfies (U1)$\sim$(U4) and (U8).

3. *Suppose that an update operator $\diamond$ satisfies (U1)$\sim$(U4) and (U8). Then, we can define an erasure operator by $(U \rightarrow E)$. The update operator obtained from the erasure operator by $(E \rightarrow U)$ is equal to the original update operator $\diamond$.*

4. *Suppose that an erasure operator $\diamond$ satisfies (E1)$\sim$(E5) and (E8). Then, we can define an update operator by $(E \rightarrow U)$. The erasure operator obtained from the update operator by $(U \rightarrow E)$ is equal to the original erasure operator $\blacklozenge$.*

Proof. 1. Assume that an update operator $\diamond$ satisfies (U1)$\sim$(U4) and (U8), and an erasure operator $\blacklozenge$ is defined by $(U \to E)$. (E1) follows from $(U \to E)$. We show (E2) If ψ implies $\neg\mu$ then it follows from (U2) that $\psi \diamond \neg\mu$ is equivalent to ψ. Therefore, $\psi \blacklozenge \mu$ is equivalent to ψ. (E3), (E4) and (E8) easily follow from (U3), (U4) and (U8), respectively. We show (E5). By (U1), $(\psi \diamond \neg\mu) \wedge \mu$ is inconsistent. Hence, $(\psi \blacklozenge \mu) \wedge \mu$ is equivalent to $\psi \wedge \mu$. Therefore, $(\psi \blacklozenge \mu) \wedge \mu$ implies ψ.

2. Assume that an erasure operator $\blacklozenge$ satisfies (E1)$\sim$(E4) and (E8), and an update operator $\diamond$ is defined by $(E \to U)$. Then, (U1) follows from $(E \to U)$. We show (U2). If ψ implies μ then it follows from (E2) that $\psi \blacklozenge \neg\mu$ is equivalent to ψ. Hence, we obtain (U2). (U3), (U4) and (U8) easily follow from (E3), (E4), and (E8).

3. Assume that an update operator $\diamond$ satisfies (U1)$\sim$(U4) and (U8). We show that $(\psi \vee (\psi \diamond \mu)) \wedge \mu$ is equivalent to $\psi \diamond \mu$. By (U1), $(\psi \diamond \mu) \wedge \mu$ is equivalent to $\psi \diamond \mu$. By Lemma 3.2, $\psi \wedge \mu$ implies $\psi \diamond \mu$. Hence, $(\psi \vee (\psi \diamond \mu)) \wedge \mu$ is equivalent to $\psi \diamond \mu$.

4. Assume that an erasure operator $\blacklozenge$ satisfies (E1)$\sim$(E5) and (E8). Let $\psi \blacklozenge' \mu$ be $\psi \vee ((\psi \blacklozenge \mu) \wedge \neg\mu)$. We show that $\psi \blacklozenge' \mu$ is equivalent to $\psi \blacklozenge \mu$. First, we show that $(\psi \blacklozenge' \mu) \wedge \mu$ is equivalent to $(\psi \blacklozenge \mu) \wedge \mu$. We know $(\psi \blacklozenge' \mu) \wedge \mu$ is equivalent to $\psi \wedge \mu$. By (E5), $(\psi \blacklozenge \mu) \wedge \mu$ implies $\psi \wedge \mu$. By (E1), $\psi \wedge \mu$ implies $(\psi \blacklozenge \mu) \wedge \mu$. Hence, $(\psi \blacklozenge \mu) \wedge \mu$ is equivalent to $\psi \wedge \mu$. Therefore, $(\psi \blacklozenge' \mu) \wedge \mu$ is equivalent to $(\psi \blacklozenge \mu) \wedge \mu$.

Next, we show that $(\psi \blacklozenge' \mu) \wedge \neg\mu$ is equivalent to $(\psi \blacklozenge \mu) \wedge \neg\mu$. By (E1), $\psi \wedge \neg\mu$ implies $(\psi \blacklozenge \mu) \wedge \neg\mu$. Hence, $(\psi \blacklozenge' \mu) \wedge \neg\mu$ is equivalent to $(\psi \blacklozenge \mu) \wedge \neg\mu$.

REFERENCES

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

Arló Costa, H.L. (1989): "Conditionals and monotonic belief revisions: the success postulate", *VI Simposio Latinoamericano de Lógica Matemática.*

Dalal, M. (1988): "Investigations into a theory of knowledge base revision: Preliminary Report", *Proceedings of the Seventh National Conference on Artificial Intelligence*, 475–479

Gärdenfors, P. (1978): "Conditionals and changes of belief", *Acta Philosophica Fennica*, Vol. XXX, 381–404.

Gärdenfors, P. (1981): "An Epistemic Approach to Conditionals", *Americal Philosophical Quarterly 18*, 203–211.

Gärdenfors, P. (1988): *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* Bradford Books, Cambridge, MA: The MIT Press, Bradford Books.

Ginsberg, M.L., and D.E. Smith. (1987): "Reasoning about action I: a possible worlds approach", *Readings in Nonmonotonic Reasoning*, M.L. Ginsberg, ed., Los Altos: Morgan Kaufmann, 434–463

Grahne, G. (1991): "Updates and counterfactuals", *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, Allen, J.A., Fikes, R., and Sandewell, E. (eds.), San Mateo, California: Morgan Kaufmann, 269–276.

Grahne, G. and A.O. Mendelzon. (1991): "Updates and Subjunctive Queries", *Proceedings of the Workshop on Non-Standard Queries and Non-Standard Answers*, Toulouse, to appear in 1991.

Katsuno, H. and A.O. Mendelzon. (1989): "A unified view of propositional knowledge base updates", In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1413–1419, 1989.

Katsuno, H. and A.O. Mendelzon. (1990): "Propositional knowledge base revision and minimal change", Technical Report KRR-TR-90-3, Department of Computer Science, University of Toronto, March 1990. Revised version to appear in *Artificial Intelligence.*

Katsuno, H. and A.O. Mendelzon. (1991): "On the difference between updating a knowledge base and revising it", *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, Allen, J.A., Fikes, R., and Sandewell, E. (eds.) San Mateo, California: Morgan Kaufmann, 387–394.

Katsuno, H. and K. Satoh. (1991): "A unified view of consequence relations, belief revision, and conditional logic", *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, 406–412

Keller, A.M. and M. Winslett Wilkins. (1985): "On the use of an extended relational model to handle changing incomplete information", *IEEE Trans. on Software Engineering*, SE-11:7, 620–633.

Lewis, D.K. (1973): *Counterfactuals.* Blackwell, Oxford.

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

Makinson, D. (1989): "The Gärdenfors impossibility theorem in non-monotonic contexts", Unpublished. To appear in *Studia Logica.*

McCarthy, J.M. and P.J. Hayes (1969). "Some philosophical problems from the standpoint of artificial intelligence", In *Machine Intelligence 4,* B. Meltzer and D. Mitchie eds, Unversity of Edinburgh Press.

Nebel, B. (1989): "A knowledge level analysis of belief revision", *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, 301–311.

Nute, D. (1984): "Conditional logic", In Dov M. Gabbay and Franz Guenthner, editors, *Handbook of Philosophical Logic*, D. Reidel, Dordrecht, 1984.

Rao, A.S. and N.Y. Foo. (1989): "Minimal change and maximal coherence: A basis for belief revision and reasoning about action", *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 966–971.

Winslett, M. (1988): "Reasoning about action using a possible models approach", In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 89–93.

Winslett, M. (1989): "Theory revision semantics for use in reasoning about action", Unpublished manuscript.

Winslett, M. (1990): *Updating Logical Databases.* Cambridge, Cambridge University Press.