Notre Dame Journal of Formal Logic
Volume 36, Number 1, Winter 1995

Constructive Modelings for Theory Change

PAVLOS PEPPAS and MARY-ANNE WILLIAMS

Abstract Alchourron, Gärdenfors and Makinson have developed and investigated a set of rationality postulates which appear to capture much of what is required of any rational system of theory revision. This set of postulates describes a class of revision functions, however it does not provide a constructive way of defining such a function. There are two principal constructions of revision functions, namely an epistemic entrenchment and a system of spheres. We refer to their approach as the AGM paradigm. We provide a new constructive modeling for a revision function based on a *nice* preorder on models, and furthermore we give explicit conditions under which a nice preorder on models, an epistemic entrenchment, and a system of spheres yield the *same* revision function. Moreover, we provide an identity which captures the relationship between revision functions and *update operators* (as defined by Katsuno and Mendelzon).

1 Introduction Theory revision models the way we change our beliefs in response to the intrusion of various forms of new information, for instance the way we might revise our beliefs in the light of information which contradicts previously accepted beliefs.

Alchourron, Gärdenfors and Makinson [1], [2], [3] have developed and investigated a set of rationality postulates which appear to capture much of what is required of any rational system of theory revision. We refer to their approach as the AGM paradigm. This set of postulates embodies the *principle of minimal change*, and describes a class of revision functions, although it does not provide a constructive way of defining such a function. Within the AGM paradigm there are two principal constructions of revision functions, namely an epistemic entrenchment as in [2], [3] and a system of spheres as in [5].

Katsuno and Mendelzon [6] provide a model-theoretic characterization of revision functions for finitary propositional languages. Their representation result relies on the finiteness property which allows an interpretation to be construed as a formula.

Grove [5] used a syntactic representation based on maximal consistent extensions, or equivalently consistent complete theories, without the restrictions of [6]. Katsuno and Mendelzon [6] note that due to the one-to-one correspondence between consistent complete theories and interpretations in the finitary propositional case, their representation result is derivable from the work of Grove [5]. Furthermore, the one-to-one correspondence between consistent complete theories and interpretations does not require the finiteness property, and therefore in the propositional case Grove's results have a semantic counterpart. However this one-to-one correspondence does not hold for the more general first order case, and a model-theoretic characterization for this case has not hitherto been established.

We provide a generalized model-theoretic construction of revision without propositional restrictions where a consistent complete theory may possess more than one model, and we call our semantic construction a nice preorder on models (in section 5 we show how our nice preorder on models is related to Gärdenfors and Makinson's [4] nice preferential model structures).

Katsuno and Mendelzon [7] formally describe the difference between revision and update. According to them an update is used to model epistemic changes due to changes in the world, while on the other hand, revision is used to model epistemic changes initiated by the acquisition of new information about a static world. They introduced a set of postulates for an update operator on finitary propositional theories. We extend their set of postulates so that an update operator may be used on arbitrary first order theories, and we provide an identity which captures the fundamental relationship between revision and update within the AGM paradigm.

The purpose of this paper is threefold; first to extend the postulates for update, and provide a connection between update and revision, thereby firmly incorporating update into the AGM paradigm. Second, to provide a new construction for revision, namely a nice preorder on models. Third, to provide explicit conditions under which a nice preorder on models, an epistemic entrenchment, and a system of spheres represent the *same* revision function.

It is well known that a contraction function can also be defined by a revision function using the Harper Identity, and we provide an identity that defines an update operator in terms of revision functions. Consequently, the framework we develop for revision also supports belief change based on contraction and update operations.

For completeness and to establish our notation we outline the AGM paradigm in Section 2, and describe how a revision function is related to an update operator. In Sections 3 and 4, we describe two well known constructions of a revision function, namely an epistemic entrenchment ordering and a system of spheres, respectively. In Section 5, we describe a nice preorder on models and show how such a structure can be used to construct a revision function. In Section 6, we provide explicit conditions under which an epistemic entrenchment ordering, a system of spheres, and a nice preorder on models render the same revision function. A discussion of our results and future work is given in Section 7.

2 The AGM paradigm Let $\mathcal{L}$ be a language which contains a complete set of Boolean connectives. We will denote sentences in $\mathcal{L}$ by lower case Greek letters. We assume $\mathcal{L}$ is governed by a logic that is identified with its consequence relation $\vdash$. The relation $\vdash$ is assumed to satisfy the following conditions (see [2]):

(a) If φ is a truth-functional tautology, then $\vdash \varphi$.
(b) If $\vdash \varphi \rightarrow \psi$ and $\vdash \varphi$, then $\vdash \psi$ (*modus ponens*).
(c) $\vdash$ is consistent, that is, $\nvdash \perp$, where $\perp$ denotes the inconsistent theory.
(d) $\vdash$ satisfies the deduction theorem.

(e) $\vdash$ is compact.

Our results can be applied to any logic satisfying these properties, however we are principally interested in first order logic. The set of all logical consequences of a set $T \subseteq \mathcal{L}$, that is $\{\varphi: T \vdash \varphi\}$, is denoted by $Cn(T)$. A *theory* of $\mathcal{L}$ is any subset of $\mathcal{L}$ closed under $\vdash$. A *consistent* theory of $\mathcal{L}$ is any theory of $\mathcal{L}$ that does not contain both φ and $\neg\varphi$, for any sentence φ of $\mathcal{L}$. A *complete* theory of $\mathcal{L}$ is any theory of $\mathcal{L}$ such that for any sentence φ of $\mathcal{L}$, the theory contains φ or $\neg\varphi$. We shall denote by $\Theta_{\mathcal{L}}$ the set of all consistent complete theories of $\mathcal{L}$, and by $\mathcal{K}_{\mathcal{L}}$ the set of all theories of $\mathcal{L}$. Finally for a set of sentence Δ, we define $[\Delta]$ to be the set of all consistent complete theories of $\mathcal{L}$ containing Δ. If Δ is inconsistent then we define $[\Delta] = \varnothing$, while if $\Delta = \varnothing$ then $[\Delta] = \Theta_{\mathcal{L}}$. For a sentence $\varphi \in \mathcal{L}$, we shall use $[\varphi]$ as an abbreviation of $[\{\varphi\}]$.

Epistemic states are belief sets which are usually partial or incomplete descriptions of the world. In the AGM paradigm of [1], [2], [3] belief sets are taken to be theories, and changes of belief are regarded as transformations on theories. There are three types of AGM transformations: expansion, contraction and revision. These transformations allow us to model changes of belief based on the *principle of minimal change*. Expansion is the simplest change, and it is most effectively employed in modeling the incorporation of beliefs that are consistent with the current set of beliefs. The *expansion* of a theory T with respect to a sentence φ, denoted as T^+_φ, is defined to be the logical closure of T and φ, that is $T^+_\varphi = Cn(T \cup \{\varphi\})$.

In contradistinction, contraction and revision are nonunique operations and cannot be represented using logical or set theoretical notions alone, but rather are constrained by a set of rationality postulates. It is these rationality postulates that attempt to embody the principle of minimal change.

A *contraction* of T with respect to φ, denoted by T^-_φ, involves the removal of a set of sentences from T so that φ is no longer implied. Formally, a contraction operator $^-$ is any function from $\mathcal{K}_{\mathcal{L}} \times \mathcal{L}$ to $\mathcal{K}_{\mathcal{L}}$, mapping $\langle T, \varphi\rangle$ to T^-_φ which satisfies the following postulates, for any $\varphi, \psi \in \mathcal{L}$ and any $T \in \mathcal{K}_{\mathcal{L}}$:

($^-$1) $T^-_\varphi \in \mathcal{K}_{\mathcal{L}}$.
($^-$2) $T^-_\varphi \subseteq T$.
($^-$3) If $\varphi \notin T$ then $T^-_\varphi = T$.
($^-$4) If $\nvdash \varphi$ then $\varphi \notin T^-_\varphi$.
($^-$5) If $\varphi \in T$, then $T \subseteq (T^-_\varphi)^+_\varphi$.
($^-$6) If $\vdash \varphi \equiv \psi$ then $T^-_\varphi = T^-_\psi$.
($^-$7) $T^-_\varphi \cap T^-_\psi \subseteq T^-_{\varphi\wedge\psi}$.
($^-$8) If $\varphi \notin T^-_{\varphi\wedge\psi}$ then $T^-_{\varphi\wedge\psi} \subseteq T^-_\varphi$.

A *revision* attempts to transform a theory "as little as possible" in order to incorporate a new sentence, possibly inconsistent with the theory. Formally, a revision operator * is any function from $\mathcal{K}_{\mathcal{L}} \times \mathcal{L}$ to $\mathcal{K}_{\mathcal{L}}$, mapping $\langle T, \varphi\rangle$ to T^*_φ which satisfies the following postulates, for any $\varphi, \psi \in \mathcal{L}$ and any $T \in \mathcal{K}_{\mathcal{L}}$:

(*1) $T^*_\varphi \in \mathcal{K}_{\mathcal{L}}$.
(*2) $\varphi \in T^*_\varphi$.
(*3) $T^*_\varphi \subseteq T^+_\varphi$.

(*4) If $\neg\varphi \notin T$ then $T^+_\varphi \subseteq T^*_\varphi$.
(*5) $T^*_\varphi = \bot$ if and only if $\vdash \neg\varphi$.
(*6) If $\vdash \varphi \equiv \psi$ then $T^*_\varphi = T^*_\psi$.
(*7) $T^*_{\varphi\wedge\psi} \subseteq (T^*_\varphi)^+_\psi$.
(*8) If $\neg\psi \notin T^*_\varphi$ then $(T^*_\varphi)^+_\psi \subseteq T^*_{\varphi\wedge\psi}$.

We relate contraction and revision by means of the *Levi Identity*, $T^*_\varphi = (T^-_{\neg\varphi})^+_\varphi$, which defines a revision in terms of a contraction, and conversely by means of the *Harper Identity*, $T^-_\varphi = T \cap T^*_{\neg\varphi}$, which defines a contraction in terms of a revision. Consequently, contraction and revision are interdefinable.

Katsuno and Mendelzon [7] introduced an update operator on finitary propositional theories as a mechanism for modifying a theory in response to changes in the world. We extend their set of postulates so that an update operator may be used on arbitrary first order theories (another difference is discussed at the end of Section 5).

Formally, an update operator $^\diamond$ is any function from $\mathcal{K}_\mathcal{L} \times \mathcal{L}$ to $\mathcal{K}_\mathcal{L}$, mapping $\langle T, \varphi\rangle$ to $T^\diamond_\varphi$ which satisfies the following postulates, for any $\varphi, \psi \in \mathcal{L}$ and any $T \in \mathcal{K}_\mathcal{L}$:

($^\diamond$1) $T^\diamond_\varphi \in \mathcal{K}_\mathcal{L}$.
($^\diamond$2) $\varphi \in T^\diamond_\varphi$.
($^\diamond$3) If $\varphi \in T$ then $T^\diamond_\varphi = T$.
($^\diamond$4) $T^\diamond_\varphi = \bot$ iff T or φ is inconsistent.
($^\diamond$5) If $\vdash \varphi \equiv \psi$ then $T^\diamond_\varphi = T^\diamond_\psi$.
($^\diamond$6) $T^\diamond_{\varphi\wedge\psi} \subseteq (T^\diamond_\varphi)^+_\psi$.
($^\diamond$7) If T is complete and $\neg\psi \notin T^\diamond_\varphi$ then $(T^\diamond_\varphi)^+_\psi \subseteq T^\diamond_{\varphi\wedge\psi}$.
($^\diamond$8) If T is consistent then $T^\diamond_\varphi = \bigcap_{K\in[T]} K^\diamond_\varphi$.

The identity below captures the association between updates and revisions, and was christened the Winslett Identity due to its close association with an identity introduced by Winslett [12].

$$\textit{Winslett Identity} \qquad T^\diamond_\varphi = \begin{cases} \bigcap_{K\in[T]} K^*_\varphi & \text{if } T \neq \bot \\ \bot & \text{otherwise} \end{cases}$$

Theorem 2.1 and Theorem 2.2 below show that for every revision function *, the function $^\diamond$ defined from * by the Winslett Identity is an update operator, and conversely, for every update operator $^\diamond$ there exists a revision function * satisfying the Winslett Identity. Consequently the update operator dwells within the realm of the AGM paradigm.

Theorem 2.1 *Let * be a revision function. Then the function $^\diamond$ defined from * and the Winslett Identity is an update operator.*

Proof: We show that $^\diamond$ satisfies the postulates ($^\diamond$1)–($^\diamond$8). Postulates ($^\diamond$1)–($^\diamond$5) follow directly from the postulates (*1)–(*8) for revision.

For ($^\diamond$6), if T is inconsistent, then ($^\diamond$6) trivially holds. Assume therefore that T is consistent. Then by the Winslett Identity we have that $T^\diamond_{\varphi\wedge\psi} = \bigcap_{K\in[T]} K^*_{\varphi\wedge\psi}$.

Moreover by (*7) we have that for every $K \in [T]$, $K^*_{\varphi\wedge\psi} \subseteq Cn(K^*_\varphi \cup \{\psi\})$, and therefore $\bigcap_{K\in[T]} K^*_{\varphi\wedge\psi} \subseteq \bigcap_{K\in[T]} Cn(K^*_\varphi \cup \{\psi\}) \subseteq Cn((\bigcap_{K\in[T]} K^*_\varphi) \cup \{\psi\})$. Combining the above we derive that $T^\diamond_{\varphi\wedge\psi} \subseteq Cn((\bigcap_{K\in[T]} K^*_\varphi) \cup \{\psi\})$, and since by the Winslett Identity, $\bigcap_{K\in[T]} K^*_\varphi = T^\diamond_\varphi$, we have that $T^\diamond_{\varphi\wedge\psi} \subseteq Cn(T^\diamond_\varphi \cup \{\psi\})$.

For ($^\diamond$7), assume that T is complete and $\neg\psi \notin T^\diamond_\varphi$. From the Winslett Identity we derive that T is consistent and $T^\diamond_\varphi = T^*_\varphi$. Then $\neg\psi \notin T^*_\varphi$ and consequently by (*8), $Cn(T^*_\varphi \cup \{\psi\}) \subseteq T^*_{\varphi\wedge\psi}$. Moreover, since T is consistent and complete, by the Winslett Identity again it follows that $T^\diamond_{\varphi\wedge\psi} = T^*_{\varphi\wedge\psi}$. Combining the above we derive that $Cn(T^\diamond_\varphi \cup \{\psi\}) \subseteq T^\diamond_{\varphi\wedge\psi}$.

Finally for ($^\diamond$8), let T be a consistent theory of $\mathcal{L}$. From the Winslett Identity it follows that for any consistent complete theory K of $\mathcal{L}$, $K^*_\varphi = K^\diamond_\varphi$. Therefore, again from the Winslett Identity we have, $T^\diamond_\varphi = \bigcap_{K\in[T]} K^*_\varphi = \bigcap_{K\in[T]} K^\diamond_\varphi$.

Theorem 2.2 *Let $^\diamond$ be an update operator. Then there exists a revision function * from which $^\diamond$ is derived via the Winslett Identity.*

Proof: We shall first prove that when focusing on consistent complete theories, $^\diamond$ satisfies the postulates (*1)–(*8) for revision.

Let K be an arbitrary consistent complete theory and let φ, ψ be any two sentences of $\mathcal{L}$. Postulates (*1), (*2), (*5), (*6), (*7), and (*8), follow directly from ($^\diamond$1), ($^\diamond$2), ($^\diamond$4), ($^\diamond$5), ($^\diamond$6), and ($^\diamond$7), respectively. Consider now (*3). If $\varphi \notin K$ then since K is complete, $\neg\varphi \in K$ and therefore $Cn(K \cup \{\varphi\}) = \mathcal{L} \supseteq K^\diamond_\varphi$. If on the other hand $\varphi \in K$ then by ($^\diamond$3), $K^\diamond_\varphi = K = Cn(K \cup \{\varphi\})$. Therefore in both cases (*3) is satisfied. Finally for (*4) assume that $\neg\varphi \notin K$. Then since K is complete, $\varphi \in K$ and therefore by ($^\diamond$3), $K^\diamond_\varphi = K = Cn(K \cup \{\varphi\})$, from which we derive that (*4) is satisfied.

From the above it follows that there exists a revision function * such that for every consistent complete theory K of $\mathcal{L}$ and every sentence $\varphi \in \mathcal{L}$, $K^\diamond_\varphi = K^*_\varphi$. Given such a revision function * it is not hard to see that $^\diamond$ is derived from * via the Winslett Identity.

Henceforth we consider revisions, but in view of the Harper Identity and the Winslett Identity our results are easily extended to both contractions and updates; see Williams [11] for details.

As noted earlier the postulates for revision merely describe the class of revision functions however they do not provide a constructive way of defining such a function. In Sections 3 and 4 we describe two well known constructions of a revision function, namely an epistemic entrenchment ordering as in [2], [3], and a system of spheres as in [5]. Then in Section 5 we introduce our own construction, a nice preorder on models which is closely related to a system of spheres.

3 Epistemic entrenchment orderings An epistemic entrenchment is an ordering of the sentences in the language which attempts to capture the importance of a sentence in the face of change as in [2], [3]. Given a theory T of $\mathcal{L}$, an *epistemic entrenchment related to* T is any binary relation $\leq$ on $\mathcal{L}$ satisfying the following postulates:

(EE1) For every $\varphi, \psi, \xi \in \mathcal{L}$, if $\varphi \leq \psi$ and $\psi \leq \xi$ then $\varphi \leq \xi$.
(EE2) For any $\varphi, \psi \in \mathcal{L}$, if $\varphi \vdash \psi$ then $\varphi \leq \psi$.

(EE3) For all $\varphi, \psi \in T$, $\varphi \leq \varphi \wedge \psi$ or $\psi \leq \varphi \wedge \psi$.
(EE4) When $T \neq \bot$, $\varphi \notin T$ if and only if $\varphi \leq \psi$ for all $\psi \in \mathcal{L}$.
(EE5) If $\psi \leq \varphi$ for all $\psi \in \mathcal{L}$, then $\vdash \varphi$.

It can be shown from these postulates that an epistemic entrenchment is a total preorder of the sentences in $\mathcal{L}$.

An epistemic entrenchment $\leq$ related to a theory K represents the relative epistemic importance of the various beliefs in K. The epistemic importance of a belief in K determines its fate when K is revised. Loosely speaking, for any two sentences φ and ψ such that $\varphi \leq \psi$, whenever a choice exists between giving up φ and giving up ψ the former will be surrendered in order to minimize the epistemic loss. Formally the idea of epistemic entrenchment determining the result of belief revision is captured by the following condition:

(E*) $\psi \in T^*_\varphi$ if and only if either $\neg\varphi < \neg\varphi \vee \psi$ or $\vdash \neg\varphi$

Theorem 3.1 below shows that the family of functions over theories constructed from epistemic entrenchments by means of (E*) is precisely the class of functions satisfying the AGM postulates for revision. Theorem 3.1 follows directly from the work of Gärdenfors and Makinson [3], [2].

Theorem 3.1 (Gärdenfors and Makinson [3]) *Let T be a theory of $\mathcal{L}$. For every revision function * there exists an epistemic entrenchment $\leq$ related to T such that (E*) is true for every $\varphi, \psi \in \mathcal{L}$. Conversely, for every epistemic entrenchment $\leq$ related to T, there exists a revision function * such that (E*) is true for every $\varphi, \psi \in \mathcal{L}$.*

4 Systems of spheres The construction of a revision function using a system of a spheres is not based on the sentences in the language as in the case of an epistemic entrenchment but rather on the set of consistent complete theories, and is due to the work of Grove [5].

A *system of spheres* **S** *centered on* $[T]$, is any collection of subsets of $\Theta_\mathcal{L}$, the elements of which we call *spheres*, that satisfies the following conditions:

(**S**1) **S** is totally ordered by set inclusion.
(**S**2) $[T]$ is the $\subseteq$-minimum element of **S**.
(**S**3) $\Theta_\mathcal{L}$ is the $\subseteq$-maximum element of **S**.
(**S**4) For every sentence φ, if there is any sphere in **S** intersecting $[\varphi]$, then there is a smallest sphere in **S** intersecting $[\varphi]$.

For a system of spheres **S** and a consistent sentence $\varphi \in \mathcal{L}$, define $C_\mathbf{S}(\varphi)$ to be the smallest sphere in **S** intersecting $[\varphi]$, and define $f_\mathbf{S}(\varphi)$ to be the intersection of $[\varphi]$ with the smallest sphere in **S** having a common element with $[\varphi]$, i.e., $f_\mathbf{S}(\varphi) = [\varphi] \cap C_\mathbf{S}(\varphi)$.

A system of spheres **S** centered on $[T]$ is depicted in Figure 1, which has been adapted from [2]. Theorem 4.1 below, due to Grove [5], shows that the following condition (**S***) can be used to construct a revision function from a system of spheres.

$$(\mathbf{S}^*)\quad T^*_\varphi = \begin{cases} \cap f_\mathbf{S}(\varphi) & \text{if } \varphi \text{ is consistent} \\ \bot & \text{otherwise} \end{cases}$$

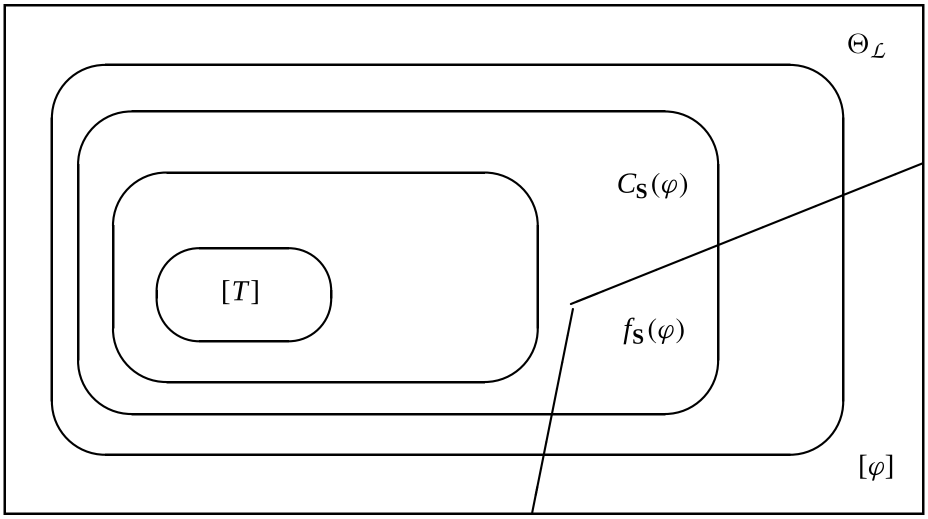

Figure 1: A System of Spheres $\mathbf{S}$ centered on $[T]$

Theorem 4.1 (Grove [5]) *Let T be a theory of $\mathcal{L}$. For every revision function * there exists a system of spheres $\mathbf{S}$ centered on $[T]$ such that $(\mathbf{S}^*)$ is true for every $\varphi \in \mathcal{L}$. Conversely, for every system of spheres $\mathbf{S}$ centered on $[T]$, there exists a revision function * such that $(\mathbf{S}^*)$ is true for every $\varphi \in \mathcal{L}$.*

We can associate an epistemic entrenchment ordering on sentences, and a system of spheres based on consistent complete theories, with a revision function and vice versa. It is upon these results that our work builds, and we give a new construction for revision, based on certain orderings on models, and provide explicit translations among these constructions such that the *same* revision is obtained.

5 Nice preorders on models So far we have described and recast well known constructions. In this section we extend the construction of Katsuno and Mendelzon [6], which is based on finite propositional interpretations, to a more general case which includes first order. First however we need to introduce some notation.

Let $\mathcal{M}_{\mathcal{L}}$ be any set of models of $\mathcal{L}$ such that every consistent complete theory has a model in $\mathcal{M}_{\mathcal{L}}$. For every nonempty subset S of $\mathcal{M}_{\mathcal{L}}$ define as $\mathrm{Th}_{\mathcal{L}}(S)$ the set of sentences that are valid in every model in S, i.e., $\mathrm{Th}_{\mathcal{L}}(S) = \{\varphi \in \mathcal{L}: M \models \varphi$ for every $M \in S\}$. If $S = \varnothing$, define $\mathrm{Th}_{\mathcal{L}}(S) = \bot$, the inconsistent theory. Clearly, for every nonempty collection S of models, $\mathrm{Th}_{\mathcal{L}}(S)$ is a consistent theory. For a set of sentences Δ, we define $\mathrm{Mod}(\Delta)$ to be the set of models in $\mathcal{M}_{\mathcal{L}}$ for which every sentence in Δ is valid, i.e., $\mathrm{Mod}(\Delta) = \{M \in \mathcal{M}_{\mathcal{L}}: M \models \varphi$ for each $\varphi \in \Delta\}$. If $\Delta = \varnothing$ we define $\mathrm{Mod}(\Delta) = \mathcal{M}_{\mathcal{L}}$, while if Δ is inconsistent $\mathrm{Mod}(\Delta) = \varnothing$.

For a theory T of $\mathcal{L}$, we define a *nice preorder on $\mathcal{M}_{\mathcal{L}}$ starting from* $\mathrm{Mod}(T)$ to be any preorder $\preceq$ on $\mathcal{M}_{\mathcal{L}}$ satisfying the following conditions.

(M1) For all $M, M' \in \mathcal{M}_{\mathcal{L}}$, either $M \preceq M'$ or $M' \preceq M$.

(M2) For all $M, M', M'' \in \mathcal{M}_{\mathcal{L}}$, if $M \preceq M'$ and $M' \preceq M''$ then $M \preceq M''$.

(M3) For every consistent sentence φ, $\mathrm{Mod}(\{\varphi\})$ has a $\preceq$-minimal element.

(M4) If T is consistent, a model $M \in \mathcal{M}_{\mathcal{L}}$ is minimal in $\mathcal{M}_{\mathcal{L}}$ if and only if $M \in \mathrm{Mod}(T)$.

where for any subset B of $\mathcal{M}_{\mathcal{L}}$, an element M of B is minimal in B if and only if $M' \preceq M$ entails $M \preceq M'$ for every $M' \in B$. For any subset B of $\mathcal{M}_{\mathcal{L}}$ define as $\min(B)$ to be the set of minimal elements in B with respect to $\preceq$. If $B = \varnothing$, then $\min(B) = \varnothing$.

Gärdenfors and Makinson [4] use a "nice preferential model structure," and it can be seen that $\langle \mathcal{M}_{\mathcal{L}}, \models, \prec \rangle$ is a nice preferential model structure.

The condition (M*) below shows how a revision function can be constructed using a nice preorder on models, and the two theorems that follow prove that the class of revision functions so constructed correspond precisely to the class that satisfy the AGM postulates.

(M*) $T^*_\varphi = \mathrm{Th}_{\mathcal{L}}(\min(\mathrm{Mod}(\{\varphi\})))$.

Theorem 5.1 *Let T be a theory of $\mathcal{L}$. For every revision function * there exists a nice preorder $\preceq$ on $\mathcal{M}_{\mathcal{L}}$ starting from* $\mathrm{Mod}(T)$ *such that (*M**) is true for every $\varphi \in \mathcal{L}$.*

Proof: Let * be an arbitrary revision function and let $\mathbf{S}$ be a system of spheres centered on $[T]$ that is associated with * by means of $(\mathbf{S}^*)$. We will prove Theorem 5.1 by constructing a nice preorder $\preceq$ starting from $\mathrm{Mod}(T)$ such that for all consistent $\varphi \in \mathcal{L}$, $f_{\mathbf{S}}(\varphi) = \{K \in \Theta_{\mathcal{L}}\colon K = \mathrm{Th}_{\mathcal{L}}(\{M\})$ for some $M \in \min(\mathrm{Mod}(\{\varphi\}))\}$. Clearly such a nice preorder $\preceq$ would satisfy condition (M*).

We construct $\preceq$ from $\mathbf{S}$ as follows. For all $M, M' \in \mathcal{M}_{\mathcal{L}}$, $M \preceq M'$ iff every sphere in $\mathbf{S}$ that contains $\mathrm{Th}_{\mathcal{L}}(\{M'\})$, also contains $\mathrm{Th}_{\mathcal{L}}(\{M\})$. It is not hard to verify that $\preceq$ so constructed, has all the desired properties, i.e. it satisfies (M1)–(M4) and moreover for any consistent $\varphi \in \mathcal{L}$, $f_{\mathbf{S}}(\varphi) = \{K \in \Theta_{\mathcal{L}}\colon K = \mathrm{Th}_{\mathcal{L}}(\{M\})$ for some $M \in \min(\mathrm{Mod}(\{\varphi\}))\}$.

Theorem 5.2 *Let T be a theory of $\mathcal{L}$. For every nice preorder $\preceq$ on $\mathcal{M}_{\mathcal{L}}$ starting from* $\mathrm{Mod}(T)$, *there exists a revision function * such that (*M**) is true for every $\varphi \in \mathcal{L}$.*

Proof: Let $\preceq$ be a nice preorder starting from $\mathrm{Mod}(T)$ and let * be the function generated from $\preceq$ by means of (M*). We show that * satisfies the postulates (*1)–(*8) for revision (our proof is essentially a reconstruction in the present context of Grove's proof of Theorem 1 in [5], and it is included mainly for completeness).

Let φ be an arbitrary formula of $\mathcal{L}$. If φ is inconsistent then by (M*), T^*_φ is also inconsistent and all eight postulates for revision are trivially satisfied. Assume therefore that φ is consistent.

Postulates (*1), (*2) and (*5) trivially follow from (M*). For (*3), if $\neg\varphi \in T$ then $T^+_\varphi = \bot$ and (*3) is trivially true. Assume therefore that $\neg\varphi \notin T$. Then $\mathrm{Mod}(T) \cap \mathrm{Mod}(\{\varphi\}) \neq \varnothing$ and therefore by (M4), $\min(\mathrm{Mod}(\{\varphi\})) = \mathrm{Mod}(T) \cap \mathrm{Mod}(\{\varphi\})$, which again implies that $T^*_\varphi = \mathrm{Th}_{\mathcal{L}}(\min(\mathrm{Mod}(\{\varphi\}))) = \mathrm{Th}_{\mathcal{L}}(\mathrm{Mod}(T) \cap \mathrm{Mod}(\{\varphi\})) = T^+_\varphi$ as desired. The above argument also proves (*4).

For (*6), observe that if $\vdash \varphi \equiv \psi$ then $\mathrm{Mod}(\{\varphi\}) = \mathrm{Mod}(\{\psi\})$ and therefore $\min(\mathrm{Mod}(\{\varphi\})) = \min(\mathrm{Mod}(\{\psi\}))$.

For (*7), if $\neg\psi \in T^*_\varphi$ then (*7) trivially holds. Assume therefore that $\neg\psi \notin T^*_\varphi$. Then $\min(\mathrm{Mod}(\{\varphi\})) \cap \mathrm{Mod}(\{\psi\}) \neq \varnothing$, and hence $\min(\mathrm{Mod}(\{\varphi\})) \cap \mathrm{Mod}(\{\psi\}) =$

$\min(\mathrm{Mod}(\{\varphi\}) \cap \mathrm{Mod}(\{\psi\})) = \min(\mathrm{Mod}(\{\varphi \wedge \psi\}))$. This again by (M*) entails that $T^*_{\varphi\wedge\psi} = (T^*_\varphi)^+_\psi$ as desired. The above argument also proves (*8).

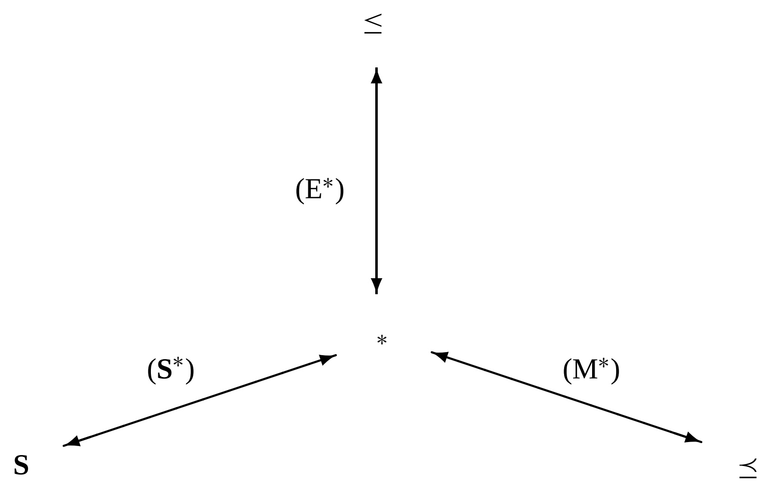

Figure 2: A revision function and its associated constructive modelings

Theorems 3.1, 4.1, 5.1, and 5.2 show that we can associate an epistemic entrenchment, a system of spheres, and a nice preorder on models with a revision function, and vice versa, as illustrated in Figure 2.

We say that $\mathcal{M}_\mathcal{L}$ is *injective*, as in Makinson [8], if and only if $\mathrm{Mod}(K)$ is a singleton for every $K \in \Theta_\mathcal{L}$.

As noted in the introduction, it is straightforward to extend Grove's results to a total preorder on models whenever $\mathcal{M}_\mathcal{L}$ is injective. However, Theorem 5.1 and Theorem 5.2 do not require injectiveness. Therefore the results hold for more general cases than those that can be immediately derived from Grove's systems of spheres, in particular they hold for first order logic where a consistent complete theory may possess more than one model.

Makinson [8] observes that injectiveness is required for update as defined by Katsuno and Mendelzon [7], and we note that the update operator defined in Section 2 is also subject to this requirement.

There are two major differences between the update operators defined in Section 2, and the ones introduced by Katsuno and Mendelson (henceforth KM updates). The first obvious difference is that the updates of section 2 apply to arbitrary theories, while the KM updates apply to sentences of a finitary propositional language. The second and perhaps more important difference between the two kinds of updates is related to the preorders on models that they induce. More precisely, a nice preorder is a certain *total* preorder on models, and inspection of the Winslett Identity reveals that an update operator that satisfies ($^\diamond$1)–($^\diamond$8), when applied to a theory T can be associated with a family of such orderings, one for every consistent complete theory containing T. In contrast, a KM update is associated with a family of certain *partial* preorders on models, in particular they satisfy (M2) – (M4). Consequently, when confined to finitary propositional languages, the updates of Section 2 are a subfamily of the KM updates. This loss of generality is compensated by the connection we were able to establish between the updates of Section 2 and revision functions. Indeed, it is precisely the property of the updates of Section 2 to induce *total* preorders on models that makes the connection with revision functions possible. This connection is an important one, not only because it relates two of the most fundamental theory change operators, but it allows one to develop a constructive model for the updates of Section 2 based on *epistemic entrenchments*, which in principle have better computational properties than preorders on models, and consequently make updates more amenable to implementation. The idea of using epistemic entrenchment to implement updates in the context of Reasoning about Action is discussed in Peppas [9]. We note that for finitary propositional languages, the subfamily of KM updates that induce total preorders on models was also identified by Katsuno and Mendelson in [7] by means of a postulate named (U9) to be added to the original postulates for KM updates numbered (U1) – (U8).

We now turn to conditions under which an epistemic entrenchment, a system of spheres and a nice preorder represent the *same* revision function.

6 Explicit translations Theorem 6.1 below provides the relationship between an epistemic entrenchment ordering and a system of spheres. Rott [10] provides a related condition for contraction which concerns the relationship between transitively relational selection functions (cf. [1], [2]) and epistemic entrenchment orderings.

Theorem 6.1 *Let T be a theory of $\mathcal{L}$. If $\leq$ is an epistemic entrenchment related to T, and $\mathbf{S}$ is a system of spheres centered on $[T]$, then $\leq$ and $\mathbf{S}$ represent the same revision function by means of conditions* (E*) *and* ($\mathbf{S}^*$) *respectively, if and only if the following condition is satisfied.*

(E$\mathbf{S}$) *For every consistent $\varphi, \psi \in \mathcal{L}$ such that $\nvdash \varphi$ and $\nvdash \psi$, $\varphi \leq \psi$ if and only if $C_{\mathbf{S}}(\neg\varphi) \subseteq C_{\mathbf{S}}(\neg\psi)$.*

Proof: Let $\leq$ be an epistemic entrenchment related to T, and $\mathbf{S}$ a system of spheres centered on $[T]$. We first show that (E$\mathbf{S}$) is a necessary condition in order for $\leq$ and $\mathbf{S}$ to represent the same revision function.

Assume that $\leq$ and $\mathbf{S}$ represent the same revision function * (by means of (E*) and ($\mathbf{S}^*$) respectively), and let φ, ψ be two consistent sentences in $\mathcal{L}$ such $\nvdash \varphi$, $\nvdash \psi$. Assume that $\varphi \leq \psi$. Then by (EE1) – (EE3) it follows that $\varphi \leq \varphi \wedge \psi$, which again implies that $\varphi \vee (\varphi \wedge \psi) \leq \varphi \wedge \psi$. Then by (E*) we derive that $\varphi \notin T^*_{\neg\varphi\vee\neg\psi}$, and therefore by ($\mathbf{S}^*$) we have that $f_{\mathbf{S}}(\neg\varphi \vee \neg\psi) \cap [\neg\varphi] \neq \varnothing$. This again entails that $C_{\mathbf{S}}(\neg\varphi) \subseteq C_{\mathbf{S}}(\neg\psi)$. Conversely, assume that $C_{\mathbf{S}}(\neg\varphi) \subseteq C_{\mathbf{S}}(\neg\psi)$. Then by ($\mathbf{S}^*$) it is not hard to see that $\varphi \notin T^*_{\neg\varphi\vee\neg\psi}$, and therefore by (E*), $\varphi \vee (\varphi \wedge \psi) \leq \varphi \wedge \psi$ or equivalently, $\varphi \leq \varphi \wedge \psi$. Then since $\varphi \wedge \psi \vdash \psi$, by (EE1) – (EE2) we derive that $\varphi \leq \psi$.

Next we prove that (E$\mathbf{S}$) is also sufficient for $\leq$ and $\mathbf{S}$ to represent the same revision function. Assume that (E$\mathbf{S}$) is satisfied and let φ be a consistent sentence of $\mathcal{L}$, such that $\nvdash \varphi$. It suffices to show that for any $\psi \in \mathcal{L}$, $\neg\varphi \vee \psi \leq \neg\varphi$ iff $\psi \notin \bigcap f_{\mathbf{S}}(\varphi)$. Let ψ be an arbitrary sentence of $\mathcal{L}$. Assume that $\neg\varphi \vee \psi \leq \neg\varphi$. Clearly then from (EE1) – (EE5), it follows that $\nvdash \psi$. If $\vdash \neg\psi$ then $\psi \notin \bigcap f_{\mathbf{S}}(\varphi)$ trivially holds (notice that since φ is consistent $\bigcap f_{\mathbf{S}}(\varphi)$ is a consistent theory). Assume therefore that ψ is consistent. Then from $\neg\varphi \vee \psi \leq \neg\varphi$ and (E$\mathbf{S}$) we derive that $C_{\mathbf{S}}(\varphi \wedge \neg\psi) \subseteq C_{\mathbf{S}}(\varphi)$, which again implies that $f_{\mathbf{S}}(\varphi) \cap [\neg\psi] \neq \varnothing$ and therefore $\psi \notin \bigcap f_{\mathbf{S}}(\varphi)$. Conversely, assume that $\psi \notin \bigcap f_{\mathbf{S}}(\varphi)$. Then clearly $\not\vdash \psi$. If on the other hand, $\vdash \neg\psi$, then $\vdash (\neg\varphi \vee \psi) \equiv \neg\varphi$, and therefore $\neg\varphi \vee \psi \leq \neg\varphi$ trivially holds. Assume therefore that ψ is consistent. From $\psi \notin \bigcap f_{\mathbf{S}}(\varphi)$ it follows that $f_{\mathbf{S}}(\varphi) \cap [\neg\psi] \neq \varnothing$ and consequently $C_{\mathbf{S}}(\varphi \wedge \neg\psi) \subseteq C_{\mathbf{S}}(\varphi)$. Then (E**S**) implies that $\neg\varphi \vee \psi \leq \neg\varphi$.

The next result provides the condition relating a system of spheres and a nice preorder on models.

Theorem 6.2 *Let T be a theory of $\mathcal{L}$. If $\mathbf{S}$ is a system of spheres centered on $[T]$ and $\preceq$ is a nice preorder on $\mathcal{M}_{\mathcal{L}}$ starting from $\mathrm{Mod}(T)$, then $\mathbf{S}$ and $\preceq$ represent the same revision function by means of conditions (**S***) and (M*) respectively, if and only if the following condition is satisfied:*

(**S**M) *For every consistent $\varphi, \psi \in \mathcal{L}$ such that $\not\vdash \varphi$ and $\not\vdash \psi$, $C_{\mathbf{S}}(\varphi) \subseteq C_{\mathbf{S}}(\psi)$ if and only if for some $M \in \mathrm{Mod}(\{\varphi\})$, $M \preceq M'$ for every $M' \in \mathrm{Mod}(\{\psi\})$.*

Proof: Let $\mathbf{S}$ be a system of spheres centered on $[T]$ and $\preceq$ a nice preorder on $\mathcal{M}_{\mathcal{L}}$ starting from $\mathrm{Mod}(T)$. We first prove that (**S**M) is a necessary condition for $\mathbf{S}$ and $\preceq$ to represent the same revision function.

Assume therefore that $\mathbf{S}$ and $\preceq$ represent the same revision function *. Then clearly, for any consistent sentence $\xi \in \mathcal{L}$, $\bigcap f_{\mathbf{S}}(\xi) = \mathrm{Th}_{\mathcal{L}}(\min(\mathrm{Mod}(\{\xi\})))$. Suppose now that φ, ψ are two consistent sentences of $\mathcal{L}$ such that $\not\vdash \varphi$ and $\not\vdash \psi$, and $C_{\mathbf{S}}(\varphi) \subseteq C_{\mathbf{S}}(\psi)$. It is not hard to verify that $C_{\mathbf{S}}(\varphi) = C_{\mathbf{S}}(\varphi \vee \psi)$, which again implies that $[\varphi] \bigcap f_{\mathbf{S}}(\varphi \vee \psi) \neq \varnothing$, and therefore $\neg\varphi \notin \bigcap f_{\mathbf{S}}(\varphi \vee \psi)$. Then from $\bigcap f_{\mathbf{S}}(\varphi \vee \psi) = \mathrm{Th}_{\mathcal{L}}(\min(\mathrm{Mod}(\{\varphi \vee \psi\})))$ we have that $\neg\varphi \notin \mathrm{Th}_{\mathcal{L}}(\min(\mathrm{Mod}(\{\varphi \vee \psi\})))$, and consequently (since $\mathrm{Mod}(\{\varphi \vee \psi\}) = \mathrm{Mod}(\{\varphi\}) \cup \mathrm{Mod}(\{\psi\})$), $\mathrm{Mod}(\{\varphi\}) \cap \min(\mathrm{Mod}(\{\varphi\}) \cup \mathrm{Mod}(\{\psi\})) \neq \varnothing$. This again entails that for some $M \in \mathrm{Mod}(\{\varphi\})$, $M \preceq M'$, for all $M' \in \mathrm{Mod}(\{\psi\})$.

Conversely, assume that φ, ψ are two consistent sentences of $\mathcal{L}$ such that $\not\vdash \varphi$ and $\not\vdash \psi$, and for some $M \in \mathrm{Mod}(\{\varphi\})$, $M \preceq M'$, for all $M' \in \mathrm{Mod}(\{\psi\})$. Then it is not hard to see that $\neg\varphi \notin \mathrm{Th}_{\mathcal{L}}(\min(\mathrm{Mod}(\{\varphi \vee \psi\})))$, and given that $\bigcap f_{\mathbf{S}}(\varphi \vee \psi) = \mathrm{Th}_{\mathcal{L}}(\min(\mathrm{Mod}(\{\varphi \vee \psi\})))$, we have that $\neg\varphi \notin \bigcap f_{\mathbf{S}}(\varphi \vee \psi)$. This again entails that $C_{\mathbf{S}}(\varphi) \subseteq C_{\mathbf{S}}(\psi)$ as desired.

Finally, we show that (**S**M) is a sufficient condition for $\mathbf{S}$ and $\preceq$ to represent the same revision function. Assume therefore that $\mathbf{S}$ and $\preceq$ satisfy (**S**M). Moreover let $\mathbf{S}'$ be a system of spheres representing the same revision function as $\preceq$. Then, from the first part of the proof it follows that $\mathbf{S}'$ and $\preceq$ satisfy (**S**M). From the above assumptions we derive that for any consistent $\varphi, \psi \in \mathcal{L}$ such that $\not\vdash \varphi$ and $\not\vdash \psi$, $C_{\mathbf{S}}(\varphi) \subseteq C_{\mathbf{S}}(\psi)$ iff $C_{\mathbf{S}'}(\varphi) \subseteq C_{\mathbf{S}'}(\psi)$. This again entails that $\mathbf{S}$ and $\mathbf{S}'$ represent the same revision function, which by the definition of $\mathbf{S}'$ is the revision function represented by $\preceq$.

Finally, the following theorem, which is a consequence of Theorem 6.1 and Theorem 6.2, provides the condition that captures the connection between an epistemic entrenchment and a nice preorder on models such that they construct the same revision function.

Theorem 6.3 *Let T be a theory of $\mathcal{L}$. If $\leq$ is an epistemic entrenchment related to T and $\preceq$ is a nice preorder on models starting from $\mathrm{Mod}(T)$, then $\leq$ and $\preceq$ represent* *the same revision function by means of conditions (*E**) and (*M**) respectively if and only if the following condition is satisfied:*

(EM) *For every consistent* $\varphi, \psi \in \mathcal{L}$ *such that* $\not\vdash \varphi$ *and* $\not\vdash \psi$, $\varphi \leq \psi$ *if and only if for some* $M \in \mathrm{Mod}(\{\neg\varphi\})$, $M \preceq M'$ *for every* $M' \in \mathrm{Mod}(\{\neg\psi\})$.

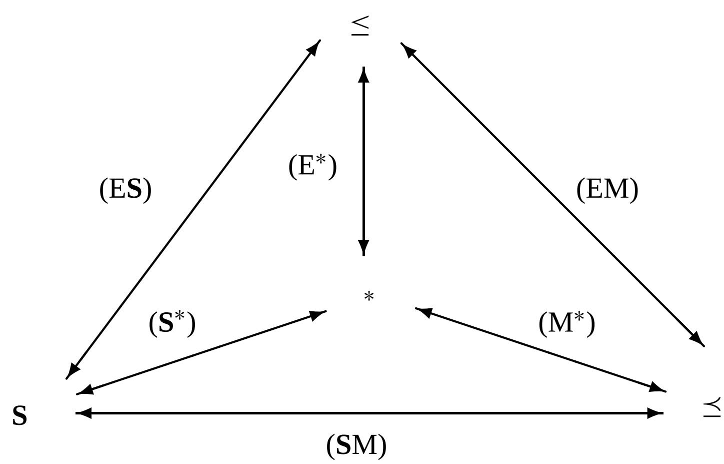

Figure 3: Interrelationships among Constructive Modelings for Revision Functions

Intuitively, this theorem says that if ψ is more entrenched than φ, then some model of $\neg\varphi$ is more plausible, than every model of $\neg\psi$. Indeed, interpreting $\preceq$ as an ordering of disbelief on models as in [10], we obtain by (M4) for instance that the models that are minimal with respect to $\preceq$ are those not disbelieved at all. Moreover, the "closer" a model is to the minimal ones, the less disbelieved this model is. In view of the (EM) condition, if ψ is more entrenched than φ, then we are more willing to give up our belief in φ in preference to ψ, because a model where φ is not true seems less disbelieved then models where ψ is not true. This interpretation engenders an interesting view of epistemic entrenchment.

Results discussed herein are summarized in Figure 3, which illustrates the three constructions associated with a revision and the interrelationships among them.

7 Discussion We have formalized the relationship between revision functions and update operators within the AGM paradigm, and we have described a new construction for a revision function, namely a nice preorder on models.

We have developed a unified view of a nice preorder on models, an epistemic entrenchment ordering, and a system of spheres, by providing explicit and perspicuous conditions under which all three constructions yield the same revision function.

In view of the relationship between revision and contraction, via the Levi Identity, and the relationship between revision and update, via the Winslett Identity, the nice preorder on models provides a construction for both contraction and update functions.

In future work we will explore analogous constructions of revision functions based on partial preorders on senteces/models, and the consequences of placing restrictions, such as well-orderedness and finiteness, on the various underlying preference relations.

Acknowledgments The authors gratefully acknowledge discussions with members of the Knowledge System Group, and wish to thank Norman Foo, David Makinson, and Hans Rott whose astute observations and comments on an earlier version of this paper led to numerous improvements.

REFERENCES

[1] Alchourron, C., P. Gärdenfors and D. Makinson, "On the logic of theory change: Partial meet contraction and revision functions," *The Journal of Symbolic Logic*, vol. 50 (1985), pp. 510–530. Zbl 0578.03011 MR 87c:03020 1, 2, 6

[2] Gärdenfors, P., *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, Cambridge, 1988. MR 89k:00015 1, 1, 2, 2, 2, 3, 3, 4, 6

[3] Gärdenfors, P. and D. Makinson, "Revisions of knowledge systems using epistemic entrenchment," pp. 83–95 in *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, edited by M. Y. Vardi, Morgan Kaufmann, Pacific Grove, 1988. Zbl 0711.03009 MR 1011079 1, 1, 2, 2, 3, 3, 3.1

[4] Gärdenfors, P. and D. Makinson, "Nonmonotonic inference based on expectations," *Artificial Intelligence*, vol. 65 (1994), pp. 197–245. Zbl 0803.68125 MR 95d:03045 1, 5

[5] Grove, A., "Two modelings for theory change," *Journal of Philosophical Logic*, vol. 17 (1988), pp. 57–170. Zbl 0639.03025 MR 89e:03027 1, 1, 1, 2, 4, 4, 4.1, 5

[6] Katsuno, H. and A. Mendelzon, "Propositional knowledge base revision and minimal change," *Artificial Intelligence*, vol. 52 (1991), pp. 263–294. Zbl 0792.68182 MR 93h:68133 1, 1, 1, 5

[7] Katsuno, H. and A. Mendelzon, "On the difference between updating a knowledge base and revising it," pp. 183–203 in *Belief Revision*, edited by P. Gärdenfors, Cambridge University Press, Cambridge, 1992. Zbl 0765.68197 MR 1142168 1, 2, 5, 5

[8] Makinson, D., "Five faces of minimality," *Studia Logica*, vol. 52 (1993), pp. 339 – 379. Zbl 0790.03026 MR 95b:03029 5, 5

[9] Peppas, P., *Belief Change and Reasoning about Action: An Axiomatic Approach to Modelling Inert Dynamic Worlds and the Connection to the Logic of Theory Change*, PhD thesis, Department of Computer Science, University of Sydney, 1993. 5

[10] Rott, H., "Two methods of constructing contractions and revisions of knowledge systems," *Journal of Philosophical Logic*, vol 20 (1991), pp. 149–173. Zbl 0723.03009 MR 92g:68136 6, 6

[11] Williams, M., *Transmutations of Knowledge Systems*, PhD thesis, Department of Computer Science, University of Sydney, 1993. 2

[12] Winslett, M., "Reasoning about action using a possible models approach," pages 89–93 in *Proceedings of the Seventh National (U.S.) Conference on Artificial Intelligence*, Morgan Kaufmann, St. Paul, 1988. 2

Department of Computing
School of Mathematics, Physics, Computing and Electronics
Macquarie University, NSW 2109

Australia
email: *pavlos@macadam.mpce.mq.edu.au*

Information Systems Group
Department of Management
Faculty of Information Science
University of Newcastle, NSW 2308
Australia
email: *maryanne@frey.newcastle.edu.au*