Revision vs. Update: Taking a Closer Look

P. Peppas[1] and **A. Nayak**[2] and **M. Pagnucco**[2] and **N. Y. Foo**[2] and **R. Kwok**[2] and **M. Prokopenko**[3]

Abstract. Belief revision and belief update are two of the most important types of belief change. Although similar in some ways, these two processes are believed to be fundamentally different. In this paper we take a closer look at belief revision (as formalized by Alchourrón, Gärdenfors and Makinson [1]) and belief update (as modelled by Katsuno and Mendelzon [4]). The results we obtain herein show that the two processes differ, not so much in the way they change a fixed belief state K, but rather on how changes at *different* theories are related to each other.

1 Introduction

Two of the most fundamental types of belief change are *belief revision* and *belief update* (or revision and update for short). Belief revision is the process by which a rational agent changes her beliefs about a *static* world in the light of new information. Much of the research in the field is based on the work of Alchourrón, Gärdenfors, and Makinson [1] who have developed a framework, which we call the *AGM paradigm*, for studying this process. In the AGM paradigm belief revision is modelled as a function * over logical theories, called a *revision function*, which satisfies a set of rationality postulates (known as the *AGM postulates*). Belief update on the other hand is the process by which an agent keeps her beliefs up to date with an *evolving* world. The difference between revision and update, although first recognized by Keller and Winslett in [5], was explored and formally captured by Katsuno and Mendelzon in [4]. They proposed a set of postulates for belief update which, although similar in some respects, are in fact quite distinct from the AGM postulates for belief revision.

In this article we take a closer look at the difference between revision and update. More precisely, we compare the two processes in terms of their construction from preorders on possible worlds [3, 4], and we show that, essentially, the results of revision can be duplicated using the construction for update. Based on this observation we prove that for a *fixed* theory K, revising K is much the same as updating K. This was evident for the limiting case where K is a consistent *complete* theory; herein however we show that it holds for *any* consistent theory K. In interpreting this result we argue that the essential difference between revision and update is not in the way they change a fixed theory K, but rather in the way changes at different theories relate to each other. In other words, when one focuses on one-step transitions between theories, revision and update are almost indistinguishable; it is only when one considers entire *sequences* of theory transitions, that successive revisions are clearly distinguishable from successive updates. This in turn, we argue, is due to an even more fundamental difference between the two processes, namely the different interpretation that the two processes ascribe to the concept of *similarity* between possible worlds.

This article is structured as follows. In the following section we introduce some notation. In section 3 we briefly review the basic definitions and results for revision and update. In section 4 we prove a representation result which essentially shows that the constructive model for update can also be used for revision. In section 5 we interpret this result and discuss the difference between revision and update. Finally in section 7 we make some concluding remarks.

2 Preliminaries

Throughout this article we shall be working with a formal language $\mathcal{L}$, governed by a logic which we identify with its consequence relation $\vdash$. We shall leave the details of $\mathcal{L}$ and $\vdash$ open, assuming only certain structural properties. In particular, for $\mathcal{L}$, we assume that it is closed under all boolean connectives, and for $\vdash$ that it satisfies the following conditions (see [2]): (i) $\vdash \varphi$ for all truth-functional tautologies φ, (ii) If $\vdash (\varphi \rightarrow \psi)$ and $\vdash \varphi$, then $\vdash \psi$, (iii) $\vdash$ is consistent, i.e. $\not\vdash \mathcal{L}$, (iv) $\vdash$ satisfies the deduction theorem, and (v) $\vdash$ is compact.

For a set of sentences Γ of $\mathcal{L}$, we denote by $Cn(\Gamma)$ the set of all logical consequences of Γ, i.e. $Cn(\Gamma) = \{\varphi \in \mathcal{L}\colon \Gamma \vdash \varphi\}$. A *theory* K of $\mathcal{L}$ is any set of sentences of $\mathcal{L}$ closed under $\vdash$, i.e. $K = Cn(K)$. We shall denote the set of all theories of $\mathcal{L}$ by $\mathcal{T}_{\mathcal{L}}$. A theory K of $\mathcal{L}$ is *complete* iff for all sentences $\varphi \in \mathcal{L}$, $\varphi \in K$ or $\neg\varphi \in K$. We shall denote the set of all consistent complete theories of $\mathcal{L}$ by $\mathcal{M}_{\mathcal{L}}$. In the context of systems of spheres, consistent complete theories essentially play the role of possible worlds. Following this convention, from now on we shall use the term "possible world" (or simply "world") to refer to an element of $\mathcal{M}_{\mathcal{L}}$. For a set of sentences Γ of $\mathcal{L}$, $[\Gamma]$ denotes the set of all consistent complete theories of $\mathcal{L}$ that contain Γ. Often we shall use the notation $[\varphi]$ for a sentence $\varphi \in \mathcal{L}$, as an abbreviation of $[\{\varphi\}]$. For a theory K and a set of sentences Γ of $\mathcal{L}$ we shall denote by $K + \Gamma$ the closure under $\vdash$ of $K \cup \Gamma$, i.e. $K + \Gamma = Cn(K \cup \Gamma)$. For a sentence $\varphi \in \mathcal{L}$ we shall often write $K + \varphi$ as an abbreviation for $K + \{\varphi\}$.

3 Belief Revision and Update

Within the AGM paradigm belief states are represented as theories of $\mathcal{L}$ and changes of beliefs are modelled as functions over theories. A *revision function* * is a function from $\mathcal{T}_{\mathcal{L}} \times \mathcal{L}$ to $\mathcal{T}_{\mathcal{L}}$, mapping

[1] Knowledge Systems Group, Department of Computing, School of Mathematics, Physics, Computing and Electronics, Macquarie University, NSW 2109, Australia.
Email: pavlos@mpce.mq.edu.au.
[2] Knowledge Systems Group, Basser Department of Computer Science, Madsen Building, F09, Sydney University, NSW 2006, Australia.
Email: {abhaya, morri, norman, rkwok}@karl.cs.su.oz.au.
[3] CSIRO Division of Information Technology, Locked Bag 17, North Ryde, NSW 2113, Australia.
Email: mikhail@syd.dit.csiro.au.

ECAI 96. *12th European Conference on Artificial Intelligence*
Edited by W. Wahlster
Published in 1996 by John Wiley & Sons, Ltd.

$\langle K, \varphi\rangle$ to $K * \varphi$ that satisfies the AGM postulates listed below [2].

(K*1)	$K * \varphi$ is a theory of $\mathcal{L}$.
(K*2)	$\varphi \in K * \varphi$.
(K*3)	$K * \varphi \subseteq K + \varphi$.
(K*4)	If $\neg\varphi \notin K$ then $K + \varphi \subseteq K * \varphi$.
(K*5)	$K * \varphi = \mathcal{L}$ iff $\vdash \neg\varphi$.
(K*6)	If $\vdash \varphi \leftrightarrow \psi$ then $K * \varphi = K * \psi$.
(K*7)	$K * (\varphi \wedge \psi) \subseteq (K * \varphi) + \psi$.
(K*8)	If $\neg\psi \notin K * \varphi$ then $(K * \varphi) + \psi \subseteq K * (\varphi \wedge \psi)$.

Apart from this axiomatic approach to belief revision, a constructive model has also been proposed by Grove [3], based on *systems of spheres*. For a subset Υ of $\mathcal{M}_{\mathcal{L}}$, a system of spheres S centered on Υ is a class of subsets of $\mathcal{M}_{\mathcal{L}}$ satisfying the following conditions:

(S1)	S is totally ordered with respect to set inclusion.
(S2)	Υ is smallest sphere in S.
(S3)	$\mathcal{M}_{\mathcal{L}} \in S$.
(S4)	For every consistent $\varphi \in \mathcal{L}$, there is a smallest sphere in S intersecting $[\varphi]$.

This modelling is a slight modification of Lewis' [6] modelling for counterfactual conditionals. The main difference between the two proposals is that Grove's systems of spheres are centered on a set of possible worlds (i.e., Υ which we shall identify with $[K]$) whereas Lewis' spheres are centered on a single possible world.[4]

For a system of spheres S and a consistent sentence $\varphi \in \mathcal{L}$, we shall denote by $C_S(\varphi)$ the smallest sphere in S intersecting $[\varphi]$. When φ is inconsistent we take $C_S(\varphi)$ to be $\mathcal{M}_{\mathcal{L}}$. We associate with any system of spheres S a function $f_S : \mathcal{L} \mapsto 2^{\mathcal{M}_{\mathcal{L}}}$ defined as follows: for every $\varphi \in \mathcal{L}$, $f_S(\varphi) = [\varphi] \cap C_S(\varphi)$.

Given a theory K of $\mathcal{L}$, and a system of spheres S centred on $[K]$, Grove defines the revision of K with respect to a sentence $\varphi \in \mathcal{L}$ as follows:

$$\text{(S*)} \qquad K * \varphi = \begin{cases} \bigcap f_S(\varphi) & \text{if } [\varphi] \neq \emptyset \\ \mathcal{L} & \text{otherwise} \end{cases}$$

Intuitively a system of spheres centered on $[K]$ is an ordering of relative *similarity* over possible worlds (i.e. consistent complete theories). The closer a world r is to the center of the system, the more similar it is to the worlds in $[K]$. Based on this reading, condition (S*) defines the revision of K by a sentence φ, as the theory determined by the φ-worlds that are most similar to the worlds in $[K]$ (i.e. the members of $f_S(\varphi)$). Grove's representation result [3] shows that the functions constructed from systems of spheres by means of (S*), are precisely those satisfying the AGM postulates.

Let us now turn to belief update as modelled by Katsuno and Mendelzon in [4]. In their formalism, a belief state is represented as a sentence of some propositional language and belief update is modelled as a function over sentences, called an *update operator*, satisfying a set of postulates numbered (U1) - (U8). Katsuno and Menedelzon also propose a constructive model for belief update based on *partial* preorders on possible worlds (*interpretations*, to be more precise) and they prove that the functions induced by these preorders are precisely those satisfying the axioms (U1) - (U8). An additional postulate (U9) is also considered in [4], which essentially imposes connectivity at the associated preorders on possible worlds; i.e. the update operators satisfying the axioms (U1) - (U9) are precisely those induced by *total* preorders on interpretations. For the sake of comparison with belief revision, we shall reproduce Katsuno and Mendelzon's account for belief update, using theories to represent belief states, and moreover we shall confine ourselves to update operators that are induced by total preorders on possible worlds. An update operator $\diamond$ is therefore defined as a function from $\mathcal{T}_{\mathcal{L}} \times \mathcal{L}$ to $\mathcal{T}_{\mathcal{L}}$, mapping $\langle K, \varphi\rangle$ to $K \diamond \varphi$, that satisfies the axioms (K$\diamond$1) - (K$\diamond$8) listed below:

(K$\diamond$1)	$K \diamond \varphi$ is a theory.
(K$\diamond$2)	$\varphi \in K \diamond \varphi$.
(K$\diamond$3)	If $\varphi \in K$ then $K \diamond \varphi = K$.
(K$\diamond$4)	$K \diamond \varphi = \mathcal{L}$ iff K or φ is inconsistent.
(K$\diamond$5)	If $\vdash \varphi \leftrightarrow \psi$ then $K \diamond \varphi = K \diamond \psi$.
(K$\diamond$6)	$K \diamond (\varphi \wedge \psi) \subseteq (K \diamond \varphi) + \psi$.
(K$\diamond$7)	If K is complete and $\neg\psi \notin K \diamond \varphi$ then $(K \diamond \varphi) + \psi \subseteq K \diamond (\varphi \wedge \psi)$.
(K$\diamond$8)	If $[K] \neq \emptyset$ then $K \diamond \varphi = \bigcap_{w \in [K]} w \diamond \varphi$.

It is not hard to verify that for any update operator $\diamond$, there exists a function $\oplus : \mathcal{L} \times \mathcal{L} \mapsto \mathcal{L}$ such that $Cn(\varphi \oplus \psi) = Cn(\varphi) \diamond \psi$, satisfying postulates (U1) - (U9) proposed by Katsuno and Mendelzon. Conversely, for any function $\oplus$ satisfying the postulates (U1) - (U9), there is an update operator $\diamond$ such that $Cn(\varphi \oplus \psi) = Cn(\varphi) \diamond \psi$. Therefore the axioms (K$\diamond$1) - (K$\diamond$8) are indeed the counterparts of (U1) - (U9) when sentences are replaced by theories in the representation of belief states (see Peppas [7] for details). Reproducing then Katsuno and Mendelzon's representation result in the present context, we have that the functions $\diamond$ satisfying (K$\diamond$1) - (K$\diamond$8) are precisely those constructed in terms of *total* preorders on possible worlds. Again for the sake of comparison with revision functions, we shall present this result using systems of spheres instead of total preorders (the two structures are essentially equivalent).

We define a *similarity structure* to be a function $\mathbf{S}$ that assigns to every consistent complete theory $w \in \mathcal{M}_{\mathcal{L}}$ a system of spheres $\mathbf{S}_w$ centered on $\{w\}$. A similarity structure is simply the reproduction in terms of systems of spheres, of Katsuno and Mendelzon's *faithful assignment* [4]. Given a similarity structure $\mathbf{S}$, the update of a theory K by a sentence φ can be defined constructively as follows:

$$\text{(S}\diamond\text{)} \qquad K \diamond \varphi = \begin{cases} \bigcap \left(\bigcup_{w \in [K]} f_{\mathbf{S}_w}(\varphi) \right), & \text{if } [\varphi], [K] \neq \emptyset \\ \mathcal{L} & \text{otherwise} \end{cases}$$

By rephrasing Katsuno and Mendelzon's representation result in terms of the above notation we derive that the functions $\diamond$ constructed from similarity structures by means of (S$\diamond$) are precisely those satisfying the postulates (K$\diamond$1) - (K$\diamond$8).

4 Reconstructing Revision Functions from Similarity Structures

In comparing a revision function * with update operators $\diamond$, the first apparent difference is in the way they change the inconsistent theory Θ. The former produces a consistent theory $\Theta * \varphi$ whenever φ is consistent, while for the later $\Theta \diamond \varphi = \Theta$, regardless of the sentence φ. However the inconsistent theory is a rather limiting case. We

[4] As we shall see, the sphere modelling for belief update is closer in spirit to Lewis' original proposal.

shall therefore focus on the difference between revision functions and update operators when applied to consistent theories.

Let K be a consistent theory of $\mathcal{L}$, * a revision function and $\diamond$ an update operator. The (technical) difference between $\diamond$ and * is perhaps best understood in terms of their constructive models. In particular, the behaviour of * at K is determined by a *single* system of spheres S centred on $[K]$ (which in principle has more than one element), while that of $\diamond$ is (essentially) determined by a *family* $\{\mathbf{S}_w\}_{w\in[K]}$ of systems of spheres, one for each world w in $[K]$, such that the system of spheres $\mathbf{S}_w$ associated to w is centered on $\{w\}$ (see figure 1).

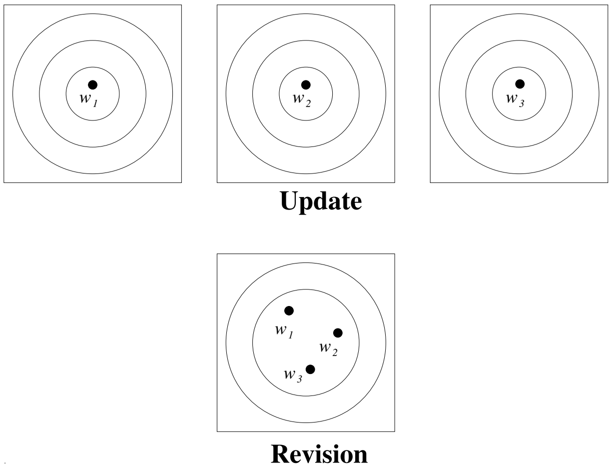

Figure 1. Constructive Models for Revision and Update with $[K] = \{w_1, w_2, w_3\}$

There are therefore two major differences in the constructive models for * and $\diamond$ (at K): firstly, there is only *one* system of spheres associated with *, while there is a whole *family* of systems of spheres associated with $\diamond$ (one for each world in $[K]$), and secondly, the centers of all systems of spheres associated with $\diamond$ are *singletons*, while the center of the system of spheres associated with * contains in general more than one world. These two differences – *one vs. many* systems of spheres, and *crowded vs. singleton* centers – have widely been considered what fundamentally distinguishes (from a technical point of view) revision functions from update operators.[5] Notice however that the way in which * and $\diamond$ use systems of spheres is quite similar. More precisely, to determine the result of revising K by a consistent sentence φ, * collects the closest φ-worlds to the center of S (i.e. the members of $f_S(\varphi)$) and defines $K * \varphi$ as the theory corresponding to these worlds (i.e. $\cap f_S(\varphi)$). The update operator $\diamond$ does much the same in determining $K\diamond\varphi$, only that it collects the closest φ-worlds to the center of not just one system of spheres, but of all systems of spheres in $\{\mathbf{S}_w\}_{w\in[K]}$ (i.e. the worlds in $\cup_{w\in[K]} f_{\mathbf{S}_w}(\varphi)$). From this it should be clear why revision and update collapse to the same process when the theory K is complete (and therefore $[K]$ has only one member). If however $[K]$ contains more than one member, the above two differences in the constructive models for * and $\diamond$ appear to suggest that revision and update are fundamentally different (even from a technical point of view). We show that (at least when K is held fixed) this is in fact a fallacy; the revision function * can be constructed from a *family* of systems of spheres centered on *singletons*, in exactly the same way that $\diamond$ is constructed.

[5] Leaving aside of course the fact that in general update operators, as originally introduced by Katsuno and Mendelzon, induce *partial* preorders on possible worlds, which is also a major divergence from revision functions. As already discussed however, in this paper we focus only on the update operators that induce total preorders.

Let $\mathbf{S}$ be a similarity structure. We shall say that $\mathbf{S}$ is *loyal* to a theory K of $\mathcal{L}$ iff the following two conditions are satisfied:

(SL1) For all $w \in [K]$, $[K] \in \mathbf{S}_w$.

(SL2) For any two $w, w' \in [K]$, $\mathbf{S}_w - 2^{[K]} = \mathbf{S}_{w'} - 2^{[K]}$.

Condition (SL1) says that $[K]$ is a sphere of $\mathbf{S}_w$ for all $w \in [K]$. Condition (SL2) says that for any $w, w' \in [K]$, the two systems of spheres $\mathbf{S}_w$ and $\mathbf{S}_{w'}$ have the same structure outside $[K]$, i.e. the only spheres that may be different in $\mathbf{S}_w$ and $\mathbf{S}_{w'}$ are the ones smaller than $[K]$ – any set larger than $[K]$ either appears as a sphere in both systems or in none of the two.

We will show that for any $\varphi \in \mathcal{L}$ and any consistent $K \in \mathcal{T}_{\mathcal{L}}$, $K * \varphi$ can be constructed from a similarity structure $\mathbf{S}$ that is loyal to K, in precisely the same way that the update of K by φ is constructed; that is in terms of the following condition:

$$\text{(SL*)} \qquad K*\varphi = \begin{cases} \bigcap\left(\bigcup_{w\in[K]} f_{\mathbf{S}_w}(\varphi)\right), & \text{if } [\varphi], [K] \neq \emptyset \\ \mathcal{L} & \text{otherwise} \end{cases}$$

Theorem 1 *Let K be a consistent theory of $\mathcal{L}$. For every revision function * there exists a similarity structure $\mathbf{S}$ that is loyal to K, such that (SL*) is satisfied for all $\varphi \in \mathcal{L}$. Conversely, for any similarity structure $\mathbf{S}$ that is loyal to K, there exists a revision function * that satisfies (SL*) for all $\varphi \in \mathcal{L}$.*

Proof.

$(\Rightarrow)$

Let * be a revision function and let $S^{\#}$ be a system of spheres centered on $[K]$ that satisfies (S*) for all $\varphi \in \mathcal{L}$ (by Grove's representation result such a system indeed exists). We shall construct from $S^{\#}$ a similarity structure $\mathbf{S}$ that is loyal to K and satisfies (SL*). For all $w \in [K]$ we define $\mathbf{S}_w$ to be $\mathbf{S}_w = S^{\#} \cup \{\{w\}\}$. For worlds w not in $[K]$ we leave the definition of $\mathbf{S}_w$ open. It is not hard to see that the similarity structure $\mathbf{S}$ so defined is indeed loyal to K. Let us now consider (SL*). Let φ be any sentence of $\mathcal{L}$. If φ is inconsistent then (SL*) trivially holds. Assume therefore that $\not\vdash \neg\varphi$. If $[\varphi] \cap [K] = \emptyset$, then by the construction of $\mathbf{S}$, $f_{\mathbf{S}_w}(\varphi) = f_{S^{\#}}(\varphi)$ for all $w \in [K]$ and consequently $\bigcup_{w\in[K]} f_{\mathbf{S}_w}(\varphi) = f_{S^{\#}}(\varphi)$, from which it follows that in this case too, (SL*) is satisfied. Finally consider the case were $[\varphi] \cap [K] \neq \emptyset$. By the construction of $\mathbf{S}$ it follows that $\bigcup_{w\in[K]} f_{\mathbf{S}_w}(\varphi) = [\varphi] \cap [K]$, which again is equal to $[K * \varphi]$ as desired.

$(\Leftarrow)$

Let $\mathbf{S}$ be a similarity structure that is loyal to K. Let w be an arbitrary element of $[K]$ (since K is consistent $[K] \neq \emptyset$). Define $S^{\#}$ as follows: $S^{\#} = \{U \in \mathbf{S}_w : [K] \subseteq U\}$. Clearly $S^{\#}$ is a system of spheres centered on $[K]$ and therefore by Grove'e representation result there is a revision function * such that $K * \varphi = \cap f_{S^{\#}}(\varphi)$ for all consistent $\varphi \in \mathcal{L}$. Moreover since $\mathbf{S}$ is loyal to K (and therefore satisfies (SL1) - (SL2)) it is not hard to see that from the construction of $S^{\#}$ it follows that $f_{S^{\#}}(\varphi) = \bigcup_{w\in[K]} f_{\mathbf{S}_w}(\varphi)$, for all consistent $\varphi \in \mathcal{L}$. Therefore the revision function * corresponding to $S^{\#}$ by means of (S*), also corresponds to $\mathbf{S}$ by means of (SL*). ■

Theorem 1 essentially provides us with an alternative constructive model for a revision function, based on a *family* of systems of spheres centered on *singletons*, as opposed to a single system of spheres centered on a family of consistent complete theories. Combining Theorem 1 with the representation result of Katsuno and Mendelzon [4, Theorem 3.4] we derive Theorem 2 below:

Theorem 2 *Let K be a consistent theory of $\mathcal{L}$. For every revision function * there exists an update operator $\diamond$ such that $K * \varphi = K \diamond \varphi$, for all $\varphi \in \mathcal{L}$.*

5 Interpreting the Results

At first glance Theorem 2 seems rather surprising. It appears to suggest that revision functions are nothing but a special kind of update operator. A closer look however will convince us that this is not the case.

Given a fixed consistent theory K and a revision function *, Theorem 2 guarantees that we can find an update operator $\diamond$ that behaved exactly like * at K, i.e. $K * \varphi = K \diamond \varphi$ for all $\varphi \in \mathcal{L}$. What Theorem 2 *does not* guarantee (for good reasons too) is that, if we now allow K to vary, the *same* update operator $\diamond$ will keep on "simulating" *. More precisely, let K' be a consistent theory which differs from K. By Theorem 2, there exists an update operator $\diamond'$ such that $K' * \varphi = K' \diamond' \varphi$ for all $\varphi \in \mathcal{L}$. However $\diamond'$ may not necessarily be the same as $\diamond$; as we change theories we may have to change update operators as well. This leads us to ask whether there is a revision function * that can be simulated by a *single* update operator $\diamond$? In other words, is there a revision function * for which there exists an update orepator $\diamond$ such that, for every *consistent* theory K and every sentence $\varphi \in \mathcal{L}$, $K * \varphi = K \diamond \varphi$? Let us call such a revision function *universal*. In view of Katsuno and Mendelzon's representation theorem and the results presented herein, it is not hard to see that the existence of a universal revision function presupposes (in fact is equivalent to) the existence of a similarity structure $\mathbf{S}$ that is loyal to *every* consistent theory K of $\mathcal{L}$. Let us call such a similarity structure *universally loyal*. Do universally loyal similarity structures exist? The answer is no, unless the object language $\mathcal{L}$ is *trivial*. We say that a language $\mathcal{L}$ (governed by a logic $\vdash$) is trivial iff $\mathcal{L}$ has at most two consistent complete theories.

Lemma 5.1 *Let $\mathbf{S}$ be a similarity structure defined over a language $\mathcal{L}$. If $\mathbf{S}$ is universally loyal then $\mathcal{L}$ is trivial.*

Proof. Assume on the contrary that $\mathbf{S}$ is universally loyal and $\mathcal{L}$ has at least three consistent complete theories w_1, w_2, and w_3 that are mutually distinct. Let K be the theory $K = w_1 \cap w_2$ and K' the theory $K' = w_2 \cap w_3$. Clearly $[K] = \{w_1, w_2\}$ and $[K'] = \{w_2, w_3\}$. Since $w_2 \in [K]$, by (SL1) we derive that $\{w_1, w_2\} \in S_{w_2}$. On the other hand, since w_2 is also a member of $[K']$, by (SL1) again we derive that $\{w_2, w_3\} \in S_{w_2}$. However $\{w_1, w_2\}$ and $\{w_2, w_3\}$ are incomparable with respect to set inclusion, which of course contradicts (S1). ■

From Lemma 5.1 we can now derive Theorem 3 below:

Theorem 3 *Every universal revision function is defined over a trivial language.*

In view of all the results presented herein, let us take a final look at the difference between revision functions and update operators. For a fixed consistent theory K we have seen that both a revision function * and an update operator $\diamond$ can be constructed from a similarity structure $\mathbf{S}$ in exactly the same way. However, this resemblance ends the moment we allow K to vary (unless the object language is trivial). This is because, contrary to $\diamond$, the revision function * requires the similarity structure associated with it, to be loyal to the current belief state, and given that a similarity structure cannot be loyal to every theory (on pain of triviality), it follows that, as the belief state changes, the similarity structure associated with * changes with it. On the other hand, the similarity structure associated with $\diamond$ is fixed for all belief sets K, unresponsive to changes of belief. This is a crucial point and we consider it to be one of the most fundamental (technical) differences between a revision function * and an update operator $\diamond$. In short, we believe that it is this concept of *fixed versus variable* similarity structures that gives a semantic basis upon which to differentiate between update operators and revision functions. This technical distinction has also an interesting intuitive interpretation. In particular, intuitively, a similarity structure $\mathbf{S}$ represents the comparative similarity of possible worlds. Since, in the context of belief update, $\mathbf{S}$ remains fixed, the notion of similarity adopted appears to be one that relates to an intrinsic, eternal property of the universe, independent of any cognitive attribute; it is what we may call *ontological similarity*. On the other hand, the fact that in the context of belief revision, $\mathbf{S}$ changes with the current belief state, appears to suggest that similarity in this context should be interpreted as the agent's *perception* of (ontological) similarity – what we may call *epistemic similarity*. This reading fits quite nicely with the original distinction made by Katsuno and Mendelson: belief update is a process triggered by changes in a dynamic world and therefore heavily dependent on the physics of the world (which explains to a certain extent the ontological interpretation of similarity), while belief revision is almost relativistic, independent of the factual connections to an external reality (which accounts for the epistemic reading of similarity).

6 Discussion

Cognoscenti might suspect that the results of the previous section (Theorem 3 in particular) are related to Gärdenfors [2] triviality (or impossibility) result. Briefly, this result shows that, if we were to extend our language to include Ramsey test conditionals of the form $x > y$ (i.e., $x > y$ if and only if $y \in K_x^*$), then there do not exist any revision functions defined over a non-trivial language. Gärdenfors points the finger at the so-called *Preservation* property (essentially (K*4)) and the following *Monotonicity* property.

(K∗M) If $K \subseteq H$, then $K_\varphi^* = H_\varphi^*$.

He suggests that it is not possible for these two properties to coexist (in the presence of other, considerably weak, conditions).

Now, it is quite easy to show that loyalty (of a similarity structure) implies Preservation. Similarly, it is easy to show that use of a similarity structure of fixed (or universal) similarity structure implies Monotonicity. Therefore, any attempt to combine these two semantic properties, as we showed in Lemma 5.1, is doomed to failure. This suggests that a semantic reason for Gärdenfors triviality result may now be at hand. Triviality arises exactly because of the attempt to conflate loyal similarity structures with those of fixed similarity. These structural or semantic properties are not compatible however — considered individually they give rise to interesting (and important) belief change strategies but together they cancel each other out. In order to show that this structural clash is the semantic culprit of Gärdenfors triviality result we would need to demonstrate that Preservation gives rise to loyal similarity structures and Monotonicity implies structures of fixed similarity (in the presence of some basic postulates on belief change of course). We are currently pursuing this direction of research by considering a basic set of rationality postulates common to both revision and update operators and examining what occurs (semantically) when additional properties, like Preservation and Monotonicity for instance, are taken into consideration. Interestingly enough, such a result would have repercussions for reasoning with conditionals $x > y$ which was Lewis' original motivation for the systems of spheres modelling.

7 Conclusion

In this article we have examined the differences between two fundamental types of belief change operator — revision functions and update functions — from a mainly technical point of view. We proved that for a fixed consistent theory K, revising K is much the same as updating K using an argument based on the systems of spheres constructive modelling. This is true regardless of whether K is complete or not. From this we concluded that the difference between revision and update is not so much on how to change a fixed theory K, but rather in the way that changes at different theories relate to one another. This in turn, we argued, is due to an even more fundamental difference between the two types of theory change, namely in the ontological versus the epistemic reading of similarity that update operators and revision function adopt respectively.

Acknowledgements

The authors would like to thank David Makinson, Hans Rott and Mary-Anne Williams for interesting discussions concerning the contents of this paper.

REFERENCES

[1] Carlos E. Alchourron, Peter Gardenfors, and David Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *The Journal of Symbolic Logic*, **50**(2), 510–530, (1985).

[2] Peter Gardenfors, *Knowledge in Flux – Modeling the Dynamics of Epistemic States*, MIT Press, Cambridge, England, 1988.

[3] Adam Grove, 'Two modelings for theory change', *Journal of Philosophical Logic*, **17**, 157–170, (1988).

[4] Hirofumi Katsuno and Alberto O. Mendelzon, 'On the difference between updating a knowledge base and revising it', in *Belief Revision*, ed., Peter Gardenfors, 183–203, Cambridge University Press, (1992).

[5] Arthur Keller and Marianne Winslett, 'On the use of an extended relational model to handle changing incomplete information', *IEEE Transactions on Software Engineering*, **SE-11**(7), 620–633, (1985).

[6] David Lewis, *Counterfactuals*, Harvard University Press, Cambridge, Massachusetts, 1973.

[7] Pavlos Peppas, *Belief Change and Reasoning about Action — An Axiomatic Approach to Modelling Dynamic Worlds and the Connection to the Logic of Theory Change*, PhD Thesis, University of Sydney, 1993.