

A Note on Prototype Revision

Renata Wassermann*

ILLC

University of Amsterdam

The Netherlands

renata@wins.uva.nl

Eduardo Fermé

Departamento de Computación

Universidad de Buenos Aires

Argentina

ferme@dc.uba.ar

1 Introduction

Belief revision (for an overview, see [Gär88, GR95]) deals with the problem of how to accommodate new assertions into an existent body of knowledge. Traditionally, the body of knowledge is represented by a belief set, a set of formulas closed under logical implication. A belief set can also be represented by the set of possible worlds where all its formulas hold.

The idea of this paper is to apply methods developed in the context of belief revision in order to accommodate new information about concepts. If we think of the representation of a belief set in terms of possible worlds, it contains the set of worlds that the agent holds for possible, given his (partial) knowledge. We will be talking about partial descriptions of an object, instead of descriptions of the world, and of sets of possible objects instead of possible worlds. The intuition behind it is that, given a description, there is a set of objects that the agent holds for possible, that is, a set of objects that could be the one being described.

For the sake of simplicity, we will concentrate on the semantical aspects of the definitions. Concepts are defined by the properties an object must have in order to be classified as an instance of the concept.

A belief set can be represented by the set of worlds consistent with the beliefs in it, that is, by the set of worlds that the agent holds for possible given his beliefs. Grove [Gro88] has given a model for belief revision based on a family of concentric spheres around a given belief set K . Each sphere represents a possible weakening of K , that is, the set of possible worlds obtained by giving up some information in K . Each sphere is also called a fallback.

*Supported by a grant of the Brazilian funding agency CAPES.

In this paper we are concerned with revising not assertions about the world, but a description of a particular object. Such a description may be given by some formula of a description language. We will not go into details about the language here, but just assume that it contains the usual truth-functional connectives.

Since we are interested in revision of concepts, we would like to change perspective. We are dealing with (partial) representations of objects. Instead of associating to a concept description a set of worlds where the concept satisfies that description (like in [Leh98]), we find it more intuitive to associate with the description a set of **objects** satisfying it. These are all the objects the agent takes for possible given the description.

Following Gärdenfors [Gär99], we will consider multidimensional spaces for representing concepts. Each dimension is a quality, roughly equivalent to an attribute name in a frame structure. An object is a point in a multidimensional space, representing the values it has for each attribute. A concept is a convex region of a multidimensional space.

Each concept comes equipped with a family of fallbacks, that represent possible “weakenings” of the concept, that is, sets of objects that satisfy only part of the description. As in the case of Grove’s models for belief revision, there is an assumption here that the fallbacks can be completely ordered. This is a simplification. The general case, where the fallbacks are not necessarily nested, will be addressed in future work.

One of the main characteristics of human reasoning is the ability of jumping to conclusions even on the absence of complete information. People can communicate without filling in all the details about the object they are talking about. When I hear that someone has a bird, I imagine a more or less typical bird, imagine it flies, etc. But of course I may be wrong, the person may be talking about a bird that does not fly. This tendency to pick up most typical objects has been studied in prototype theory ([Ros73, Lak87]). Gärdenfors has also addressed prototype effects on conceptual spaces [Gär99].

The use of multidimensional spaces instead of frame-like structures makes some aspects of concept descriptions clearer. One can talk about concepts being close to each other, and most important, about objects being more or less central in a concept. Typical objects are more central than others. Desk-chairs are more typical than rocking chairs and should be represented by points near the center of the region representing the concept chair.

Gärdenfors defines the prototypical region of a concept as a sphere around the center of the concept. The radius of the sphere has to be determined empirically.

Throughout the paper, we will represent a concept complex by an ordered pair formed by the description of the concept and the description of the prototypical region.

Definition 1.1 *Let K and P be sets of formulas. $\langle K, P \rangle$ is a concept complex*

iff:

1. K and P are belief sets, and
2. $K \subseteq P$

The set K describes objects of a certain concept. We will also use K as the name of the concept. The set P describes the prototypical K -objects.

The set of possible objects that satisfy description K is denoted by $[K]$. Similarly, the set of objects that satisfy P is denoted by $[P]$.

Suppose we have the concept tiger defined by means of properties like being a vertebrate, having a certain DNA structure, etc. Prototypical tigers are big, yellow with black stripes, carnivore, have four legs, etc. We want to revise this concept due to new information acquired. Five situations are possible, as can be seen in figure 1 (α is the new piece of information being added).




Figure 1: Revising concept complexes

In situation (a), the new information is completely incompatible with even the essential properties of tigers. Suppose the new input is that the object is invertebrate. There is no such thing as an invertebrate tiger. In this case, the

new concept assumed should be the intersection of the invertebrate objects with the inclusion minimal fallback of tiger that is compatible to it. This is equivalent to a revision of the concept of tiger by the property of being invertebrate. The resulting concept is that of invertebrates that look as much as a tiger as possible. There is no information about the new prototypical region.

In situation (b), the new information is compatible with tiger but not with prototypical tiger. Suppose it is that the object has three legs. The new concept is that of three-legged tiger, that is a subset of tiger. The new prototypical region contains those three-legged tigers that are as close to prototypical tigers as possible (in that they are big, yellow with black stripes, etc.).

In situation (c), the new information is compatible with tiger as well as with prototypical tiger. It could be that the object is female. The concept of female tiger is the intersection of tiger and female and the prototypical female tiger is the intersection of prototypical tiger with female, that is, those female tigers that are big, yellow with black stripes, etc.

In situation (d), the new information is valid for all prototypical tigers, but not for all tigers. It could be that the object has four legs. In this case, the new concept is that of a four-legged tiger, that is the intersection of tiger and four-legged. The prototypical region stays the same, prototypical four-legged tigers are prototypical tigers.

Finally, in situation (e), the new information is already valid for all tigers, for example, that they are vertebrates. Nothing changes, vertebrate tigers are tigers.

We will examine two approaches for the revision of prototypes. First we will consider that, besides the spheres around K , the description of the concept, we have one internal sphere which represents the prototypical region, described by P . We give AGM-style characterizations for the revision of concepts and prototypes in terms of spheres, rationality postulates and partial meet revision. Next, we examine what happens when we allow for more spheres between P and K .

2 First Approach

In this section we will introduce rationality postulates for the revision of concept complexes and will present two constructive methods for obtaining the revision.

2.1 Postulates

Definition 2.1 *Let $\langle K, P \rangle$ be a concept complex and $\langle K, P \rangle @ \alpha = \langle K_\alpha, P_\alpha \rangle$. An operation $@$ of complex revision satisfies the following postulates:*

- (KP@1) $\langle K_\alpha, P_\alpha \rangle$ is a complex.
- (KP@2) $\alpha \in K_\alpha$.
- (KP@4.1) If $\neg \alpha \notin K$, then $K_\alpha = K + \alpha$.

(**KP@4.2**) If $\neg\alpha \notin P$, then $P_\alpha = P + \alpha$.

(**KP@5.1**) If $\vdash \neg\alpha$, then $K_\alpha = L$.

(**KP@5.2**) If $P_\alpha = L$ then $\vdash \neg\alpha$.

(**KP@6**) If $\vdash \alpha \leftrightarrow \beta$, then $\langle K_\alpha, P_\alpha \rangle = \langle K_\beta, P_\beta \rangle$.

(**KP@7**) If $\neg\alpha \in P$, then $P_\alpha \subseteq K_\alpha$.¹

$$(\mathbf{KP@DF}) P_{\alpha \vee \beta} = \begin{cases} P_\alpha, \text{ or} \\ P_\beta, \text{ or} \\ P_\alpha \cap P_\beta \end{cases}$$

Proposition 2.2 Let $\langle K, P \rangle$ be a complex and $@$ a operation on $\langle K, P \rangle$ that satisfies (**KP@1**) - (**KP@6**) and (**KP@DF**). Then $@$ satisfies:

(**KP@5.1'**) $K_\alpha = L$ if and only if $\vdash \neg\alpha$.

(**KP@5.2'**) $P_\alpha = L$ if and only if $\vdash \neg\alpha$.

$$(\mathbf{KP@DF.1}) K_{\alpha \vee \beta} = \begin{cases} K_\alpha, \text{ or} \\ K_\beta, \text{ or} \\ K_\alpha \cap K_\beta \end{cases}$$

2.2 Sphere-System

The model we suggest here is based on Grove's modelling for belief revision. We use a family of spheres centered in P , where $\llbracket P \rrbracket$ and $\llbracket K \rrbracket$ are spheres of the family and the other spheres contain $\llbracket K \rrbracket$ (see Figure 2).

Definition 2.3 (adapted from [Gro88]) Let M be the set of all possible objects. A family of spheres centered on $\langle K, P \rangle$ is a collection S of subsets of M such that:

(S1) If $U, V \in S$, then $U \subseteq V$ or $V \subseteq U$.

(S2) $\llbracket P \rrbracket$ is the \subseteq -minimum of S .

(S3) M is the \subseteq -maximum of S .

(S4) If α is a sentence in \mathcal{L} and there is any sphere in S intersecting $\llbracket \alpha \rrbracket$, then there is an inclusion minimal sphere intersecting $\llbracket \alpha \rrbracket$.

(S5) $\llbracket K \rrbracket \in S$

(S6) For every sphere $U \in S$: if $\llbracket P \rrbracket \subseteq U \subset \llbracket K \rrbracket$, then $U = \llbracket P \rrbracket$.²

Definition 2.4 Let $\langle K, P \rangle$ be a concept complex. Let S be a family of spheres around $\langle K, P \rangle$ satisfying (S1)-(S6). Let $f_S(\alpha) = \llbracket \alpha \rrbracket \cap c_S(\alpha)$, where $c_S(\alpha)$ is the minimal sphere in S which intersects $\llbracket \alpha \rrbracket$. The revision $@$ of $\langle K, P \rangle$ by a sentence α is given by:

$$\langle K, P \rangle @ \alpha = \langle Th(f_S(\alpha) \cup (\llbracket K \rrbracket \cap \llbracket \alpha \rrbracket)), Th(f_S(\alpha)) \rangle$$

¹Note that the converse, i.e. $K_\alpha \subseteq P_\alpha$ follows from (**KP@1**).

²This postulate means that there are no other spheres between $\llbracket P \rrbracket$ and $\llbracket K \rrbracket$




Figure 2: One internal sphere

2.3 Partial Meet

Definition 2.5 [AM82] *The remainder operation \perp is the operation such that for all subsets B and elements α of L , $X \in B \perp \alpha$ if and only if:*

1. $X \subseteq B$,
2. $\alpha \notin Cn(X)$, and
3. $\alpha \in Cn(Y)$ for all Y such that $X \subset Y \subseteq B$.

Observation 2.6 [AM82] *If K is a belief set, then so are the elements of $K \perp \alpha$ for all $\alpha \in L$.*

Definition 2.7 [AGM85] *Let $K \subseteq L$ and $\alpha \in L$. A selection function for $K \perp \alpha$ is a function g such that:*

1. $g(K \perp \alpha) \subseteq K \perp \alpha$, and
2. If $K \perp \alpha \neq \emptyset$ then $g(K \perp \alpha) \neq \emptyset$, otherwise, $g(K \perp \alpha) = K$.

Definition 2.8 *Let $\langle K, P \rangle$ be a concept complex. Then $@$ is a partial meet revision of $\langle K, P \rangle$ by α if and only if:*

$$\langle K, P \rangle @ \alpha = \langle K_\alpha, P_\alpha \rangle$$

where:

$$K_\alpha = (\bigcap g(K \perp \neg \alpha)) + \alpha$$

and

$$P_\alpha = (\bigcap h(P \perp \neg \alpha)) + \alpha$$

and $Y \in h(P \perp \neg \alpha)$ if and only if either $Y = P$ or $Y \in g(K \perp \neg \alpha)$.

Definition 2.9 A selection function g is transitively relational if and only if it is based on some transitive relation \sqsubseteq in the sense that $g(K \perp \neg \alpha) = \{H \in K \perp \neg \alpha \mid H' \sqsubseteq H \text{ for all } H' \in K \perp \neg \alpha\}$

Proposition 2.10 Let $\langle K, P \rangle$ be a concept complex. Let $*$ be an AGM partial meet revision operation for K . Then $@$ is a partial meet revision of $\langle K, P \rangle$ if and only if:

$$\langle K, P \rangle @ \alpha = \langle K * \alpha, P_\alpha \rangle$$

where:

$$P_\alpha = \begin{cases} K * \alpha & \text{if } \neg \alpha \in P \\ P + \alpha & \text{otherwise} \end{cases}$$

2.4 Representation Result

The following theorem shows that the three definitions of the revision operation $@$ coincide:

Theorem 2.11 Let $\langle K, P \rangle$ be a concept complex and $@$ an operation on $\langle K, P \rangle$. Then the following statements are equivalent:

1. $@$ satisfies the postulates in Definition 2.1.
2. $@$ is defined as in Definition 2.4.
3. $@$ is defined as in Definition 2.8, and the selection functions g and h on which $@$ is based are transitively relational.

The problem with this approach is that there is not enough information about the new prototypical region. Proposition 2.10 implies that whenever $\neg \alpha \in P$, revision of $\langle K, P \rangle$ by α leads to a new complex where the prototypical region and the complex are the same, i.e., $\langle K', K' \rangle$. In the next section we explore another approach that does not have this drawback.

3 Second Approach

In the previous section we have characterized an operation of concept complex revision by means of postulates, systems of spheres and partial meet constructions. As we have seen, after revising a complex structure $\langle K, P \rangle$ by α , if $\neg\alpha \in P$ then the new prototypical region is equal to the new concept. This is not very natural. We would like to obtain a new prototypical region which is a proper subset of the new concept. For this, more structure is needed. In this section we observe what happens when more than one sphere internal to K is allowed.

3.1 Postulates

In this new formulation, postulate **(KP@7)** is substituted by **(KP@7.1)**. This means that the revised prototypical region is equal to the revised concept only when the new information is inconsistent with the concept. When the new information is inconsistent with the prototypical region but not with the concept, the description of the revised prototypical region is a proper superset of the description of the revised concept.

Definition 3.1 *Let $\langle K, P \rangle$ be a concept complex and $@$ an operation such that $\langle K, P \rangle @ \alpha = \langle K_\alpha, P_\alpha \rangle$. The operation $@$ is an operation of complex revision if it satisfies the following postulates:*

- (KP@1)** - **(KP@6)**, **(KP@DF)**, **(KP@DF.1)**, and
- (KP@7.1)** *If $\neg\alpha \in K$, then $P_\alpha \subseteq K_\alpha$.*

3.2 Sphere-System

Now we allow for more spheres between $[K]$ and $[P]$, as can be seen in Figure 3.

Definition 3.2 *Let $\langle K, P \rangle$ be a concept complex. Let S be a family of spheres around $\langle K, P \rangle$ satisfying (S1)-(S5)³. Let $f_S(\alpha) = [\alpha] \cap c_S(\alpha)$, where $c_S(\alpha)$ is the minimal sphere in S which intersects $[\alpha]$. The revision $@$ of $\langle K, P \rangle$ by a sentence α is given by:*

$$\langle K, P \rangle @ \alpha = \langle Th(f_S(\alpha) \cup ([K] \cap [\alpha])), Th(f_S(\alpha)) \rangle$$

3.3 Partial Meet

Definition 3.3 *Let $\langle K, P \rangle$ be a concept complex. Then $@$ is a partial meet revision of $\langle K, P \rangle$ by α if and only if:*

$$\langle K, P \rangle @ \alpha = \langle K_\alpha, P_\alpha \rangle$$

³Note that we suppress (S6) in this approach.




Figure 3: Several Internal Spheres

where:

$$K_\alpha = (\bigcap g(K \perp \neg \alpha)) + \alpha$$

and

$$P_\alpha = (\bigcap h(P \perp \neg \alpha)) + \alpha$$

and $X \in g(K \perp \neg \alpha)$ if and only if there exists $Y \in h(P \perp \neg \alpha)$ such that $(Y \cap K) = X$.

3.4 Representation Result

The following theorem shows the relation between the three last definitions:

Theorem 3.4 *Let $\langle K, P \rangle$ be a concept complex and $@$ an operation on $\langle K, P \rangle$. Then the following statements are equivalent:*

1. $@$ satisfies the postulates in Definition 3.1.
2. $@$ is defined as in Definition 3.2.
3. $@$ is defined as in Definition 3.3, and the selection functions g and h on which $@$ is based are transitively relational.

Recall the tiger example presented in the introduction. With the second approach, only in case (a) of Figure 1, i.e., when $\neg\alpha \in K$, we get $P_\alpha = K_\alpha$. This can be understood as follows: when $\neg\alpha \in K$, the change by adding α is a change of concept and there is no information available about the prototypical region of the new concept. More structure is needed.

4 Future Work

The nonmonotonic properties of prototypes were studied in [Leh98] and [Was98]. It would be interesting to combine the results obtained here with those articles, mirroring the relation between AGM belief revision and nonmonotonic inferences given in [MG91].

In this article we have only dealt with the case where the fallbacks can be nested. This is not a very realistic assumption. Recall the tiger example. If the prototypical tigers are described as those having three legs, being big and yellow, there are three ways of weakening this description, by giving up each of the given properties, and it is not clear how to order these weakenings. A solution to this problem could be to adopt the model for relational belief revision given in [LR91], which allow for non-nested fallbacks.

In [Pag96], a model for abductive expansion is given which uses a family of spheres internal to the belief set. It would be interesting to see the relation between the technical results obtained for abductive expansion and the ones given here.

A Appendix: Proofs

The next Lemmas will be used in the proofs.

Lemma A.1 [Mak85] *Let K be a belief set. If $\neg\alpha \in K$, then for every $\beta \in L$ and every $X \in K \perp \neg\alpha$, either $\alpha \rightarrow \beta \in X$ or $\alpha \rightarrow \neg\beta \in X$.*

Lemma A.2 *Let $\langle K, P \rangle$ be a complex. Let S be a family of spheres around $\langle K, P \rangle$ satisfying (S1)-(S6). Then there exists a family of spheres S_K around $\llbracket K \rrbracket$, satisfying (S1), (S3), (S4), and (S5) such that $U \in S_K$ iff $U \in S$ and $\llbracket K \rrbracket \subseteq U$.*

Proof: Let $S_K = \{U \in S \mid \llbracket K \rrbracket \subseteq U\}$. We have to show that S_K satisfies (S1), (S3), (S4), and (S5). Since $\llbracket K \rrbracket \in S$, (S5) and (S3) follow trivially. Moreover, $\llbracket K \rrbracket$ is the \subseteq -minimum of S_K .

(S1): If $U, V \in S_K$, since $U, V \in S$, either $U \subseteq V$ or $V \subseteq U$.

(S4): Let $c_S(\alpha)$ be the minimal sphere in S intersecting $[\alpha]$. We define $c_{S_K}(\alpha)$ to be $c_S(\alpha)$ if $c_S(\alpha) \in S_K$. Otherwise, $c_{S_K}(\alpha) = \llbracket K \rrbracket$. It is easy to see that $c_{S_K}(\alpha)$ is the minimal sphere of S_K intersecting $[\alpha]$. \square

Lemma A.3 [AGM85, Gro88]

Let K be a belief set and $*$ an operator for K and let $K_\alpha = K * \alpha$. Then the following conditions are equivalent:

1. $*$ satisfies closure (K_α is a belief set), (KP@2), (KP@4.1), (KP@5.1), (KP@6), and (KP@DF.1).
2. There exists a family S of spheres around K satisfying (S1)-(S4) such that $K * \alpha = Th(f_S(\alpha))$.
3. There exists a transitively relational selection function g such that $K * \alpha = (\bigcap g(K \perp \neg \alpha)) + \alpha$.

Lemma A.4 Let g be a selection function for $K \perp \neg \alpha$ and h a selection function for $P \perp \neg \alpha$, such that $Y \in h(P \perp \neg \alpha)$ if and only if either $Y = P$ or $Y \in g(K \perp \neg \alpha)$. If g is transitively relational, then so is h .

Proof: If $\neg \alpha \notin P$, then $P \perp \neg \alpha = \{P\}$ and it follows trivially that h is transitively relational. Suppose $\neg \alpha \in P$. Then, by the definition of h , $h(P \perp \neg \alpha) \subseteq g(K \perp \neg \alpha)$. Since g is based on a transitive relation \sqsubseteq , we can extend \sqsubseteq to $P \perp \neg \alpha$ in such a way that $h(P \perp \neg \alpha) = \{Y \in P \perp \neg \alpha \mid Y' \sqsubseteq Y \text{ for all } Y' \in P \perp \neg \alpha\}$. \square

Lemma A.5 Let h be a selection function for $P \perp \neg \alpha$ and g a selection function for $K \perp \neg \alpha$, such that $X \in g(K \perp \neg \alpha)$ if and only if there exists $Y \in h(P \perp \neg \alpha)$ such that $(Y \cap K) = X$. If h is transitively relational, then so is g .

Proof: If h is transitively relational, then there is a transitive relation \sqsubseteq such that $h(P \perp \neg \alpha) = \{Y \in P \perp \neg \alpha \mid Y' \sqsubseteq Y \text{ for all } Y' \in P \perp \neg \alpha\}$. Let \sqsubseteq_K be a relation such that for all $X, X' \in K \perp \neg \alpha$:

$X \sqsubseteq_K X'$ if and only if $Y \sqsubset Y'$, $X = Y \cap K$ and $X' = Y' \cap K$.

It is easy to see that g is based on \sqsubseteq_K . We only have to show that \sqsubseteq_K is transitive. Let $X_1 \sqsubseteq_K X_2 \sqsubseteq_K X_3$. We have to show that $X_1 \sqsubseteq_K X_3$. From the definition of \sqsubseteq_K , we know that there are Y_1, Y_2 and Y_3 in $P \perp \neg \alpha$ such that $X_1 = Y_1 \cap K$, $X_2 = Y_2 \cap K$, $X_3 = Y_3 \cap K$, and $Y_1 \sqsubset Y_2 \sqsubset Y_3$. Since \sqsubseteq is transitive, $Y_1 \sqsubseteq Y_3$, and hence, $X_1 \sqsubseteq_K X_3$. \square

Proof of Proposition 2.2: (KP@5.1') and (KP@5.2') follow trivially from (KP@5.1) and (KP@5.2). For (KP@DF.1) we will prove by cases.

(a) $\neg \alpha \notin K$ and $\neg \beta \notin K$: Then $\neg(\alpha \vee \beta) \notin K$, from which it follows that $K_{\alpha \vee \beta} = K + \alpha \vee \beta = K_\alpha \cap K_\beta$.

(b) $\neg \alpha \notin K$ and $\neg \beta \in K$: Then $\neg(\alpha \vee \beta) \notin K$, from which it follows, since $\neg \beta \in K$ and $K \neq K_\perp$, that $K_{\alpha \vee \beta} = K + \alpha \vee \beta = (K + \neg \beta) + (\alpha \vee \beta) = K + (\neg \beta \wedge \alpha) = K + \alpha$.

(c) $\neg\alpha \in K$ and $\neg\beta \notin K$: Due the symmetry of the case, is similar to case (b).

(d) $\neg\alpha \in K$ and $\neg\beta \in K$: Then $\neg(\alpha \vee \beta) \in K$. Then $\neg\alpha \in P$, $\neg\beta \in P$ and $\neg(\alpha \vee \beta) \in P$. Hence by **(KP@1)**, **(KP@7)** and **(KP@DF)**:

$$K_{\alpha \vee \beta} = P_{\alpha \vee \beta} = \begin{cases} P_\alpha = K_\alpha, \text{ or} \\ P_\beta = K_\beta, \text{ or} \\ P_\alpha \cap P_\beta = K_\alpha \cap K_\beta \end{cases}$$

□

Proof of Proposition 2.10:

By definition, $*$ is a partial meet revision function if and only if there exists a selection function g such that $K * \alpha = (\bigcap g(K \perp \neg\alpha)) + \alpha$. We must prove only:

There exists h , where $Y \in h(P \perp \neg\alpha)$ if and only if either $Y = P$ or $Y \in g(K \perp \neg\alpha)$ and such that $P_\alpha = \bigcap h(P \perp \neg\alpha) + \alpha$. We prove it by cases:

(a) $\neg\alpha \notin P$: Then $(\bigcap h(P \perp \neg\alpha)) + \alpha = (\bigcap h(\{P\})) + \alpha = P + \alpha$.

(b) $\neg\alpha \notin K$ and $\neg\alpha \in P$: Then $g(K \perp \neg\alpha) = \{K\}$ and (since $\not\vdash \neg\alpha$) $P \notin h(P \perp \neg\alpha)$. Hence, $h(P \perp \neg\alpha) = \{K\}$ and it follows that $P_\alpha = K + \alpha = K * \alpha$.

(c) $\neg\alpha \in K$: We will prove by double inclusion. For the first direction let $\beta \in K * \alpha$. Then it follows that $\beta \in (\bigcap g(K \perp \neg\alpha)) + \alpha$, from which it follows that $\alpha \rightarrow \beta \in \bigcap g(K \perp \neg\alpha)$. For all $Y \in h(P \perp \neg\alpha)$ either $Y = P$ or $Y \in g(K \perp \neg\alpha)$. In both cases it follows that $\alpha \rightarrow \beta \in \bigcap h(P \perp \neg\alpha)$; from which we conclude that $\beta \in (\bigcap h(P \perp \neg\alpha)) + \alpha$.

For the other direction, let $\beta \in (\bigcap h(P \perp \neg\alpha)) + \alpha$. Then $\alpha \rightarrow \beta \in \bigcap h(P \perp \neg\alpha)$. Suppose by reductio that $\alpha \rightarrow \beta \notin \bigcap g(K \perp \neg\alpha)$; then there exists $X \in g(K \perp \neg\alpha)$ such that $\alpha \rightarrow \beta \notin X$. By previous lemma A.1 $\alpha \rightarrow \neg\beta \in X$; from which it follows (since $X \in h(P \perp \neg\alpha)$) that $\alpha \rightarrow \beta \notin \bigcap h(P \perp \neg\alpha)$. Absurd, then $\alpha \rightarrow \beta \in \bigcap g(K \perp \neg\alpha)$; hence $\beta \in (\bigcap g(K \perp \neg\alpha)) + \alpha = K * \alpha$.

□

Proof of Theorem 2.11:

Part 1: Postulates to Partial Meet:

Due to lemma A.3, if K_α satisfies *closure*; (i.e., K_α is a belief set), **(KP@2)**, **(KP@4.1)**, **(KP@5.1)**, **(KP@6)**, and **(KP@DF.1)**, then there exists a transitively relational function g , such that $K_\alpha = (\bigcap g(K \perp \neg\alpha)) + \alpha$. Due to proposition 2.10 it is suffices to prove that

$$P_\alpha = \begin{cases} K_\alpha & \text{if } \neg\alpha \in P \\ P + \alpha & \text{otherwise} \end{cases}$$

since Lemma A.4 yields that h is transitively relational.

Let $\neg\alpha \in P$, then by **(KP@1)** and **(KP@7)** it follows that $P_\alpha = K_\alpha$. If $\neg\alpha \notin P$, then **(KP@4.2)** yields that $P_\alpha = P + \alpha$.

Part 2: Partial Meet to Spheres:

From Lemma A.3 it follows that:

(1) $P_\alpha = (\bigcap h(P \perp \neg\alpha)) + \alpha$ and h transitively relational imply that there is a system of spheres S satisfying **(S1)**-**(S4)** such that $P_\alpha = Th(f_S(\alpha))$.

(2) $K_\alpha = (\bigcap g(K \perp \neg\alpha)) + \alpha$ and g transitively relational imply that there is a system of spheres S_K satisfying **(S1)**, **(S3)**-**(S5)** such that $K_\alpha = Th(f_{S_K}(\alpha))$, where $c_{S_K}(\alpha)$ is the minimal sphere of S_K intersecting $[\alpha]$ and $f_{S_K}(\alpha) = c_{S_K}(\alpha) \cap [\alpha]$.

We have to prove that (a) $S_K = \{U \in S \mid [K] \subseteq U\}$, (b) S satisfies **(S6)**, and (c) $Th(f_{S_K}(\alpha)) = Th(f_S(\alpha) \cup ([K] \cap [\alpha]))$.

(a) Let $U \in S$ and let β be a sentence such that $U = c_S(\beta)$. We have to show that if $[K] \subseteq U$, then $U \in S_K$. From (1) we know that $Th(c_S(\beta)) = \bigcap h(P \perp \neg\beta)$. Since Y in $h(P \perp \neg\beta)$ if and only if either $Y = P$ or Y in $g(K \perp \neg\beta)$, we have that $\bigcap h(P \perp \neg\beta) \cap K = \bigcap g(K \perp \neg\beta)$. This means that $c_S(\beta) \cup [K] = c_{S_K}(\beta)$ and hence, if $[K] \subseteq U$, then $c_S(\beta) \cup [K] = U = c_{S_K}(\beta)$ and so, $U \in S_K$.

(b) Let $[P] \subseteq U \subseteq [K]$ and β be a sentence such that $U = c_S(\beta)$. Then $Th(c_S(\beta)) = \bigcap h(P \perp \neg\beta)$. But since $\bigcap h(P \perp \neg\beta) = P$ or $\bigcap h(P \perp \neg\beta) \subseteq K$, we have that either $U = [P]$ or $[K] \subseteq U$. Since we assumed $U \subseteq [K]$, $U = [P]$.

(c) There are two cases:

Case 1: If $c_S(\alpha) \subseteq K$, then $c_{S_K} = [K]$ and $f_{S_K} = [K] \cap [\alpha]$. Since $f_S(\alpha) = c_S(\alpha) \cap [\alpha] \subseteq [K] \cap [\alpha]$, we have $f_S(\alpha) \cup ([K] \cap [\alpha]) = [K] \cap [\alpha] = f_{S_K}(\alpha)$.

Case 2: If $K \subseteq c_S(\alpha)$, then $c_{S_K}(\alpha) = c_S(\alpha)$ and $[K] \cap [\alpha] = \emptyset$. Hence, $f_{S_K}(\alpha) = f_S(\alpha) \cup ([K] \cap [\alpha])$.

Part 3: Spheres to Postulates:

Due to lemma A.2 and A.3, \textcircled{a} satisfies **(KP@2)**, **(KP@4.1)**, **(KP@5.1)**, K_α is a belief set, and $K_\alpha = K_\beta$ if $\vdash \alpha \leftrightarrow \beta$. By A.2, since S is a sphere system centered on $[P]$, we obtain **(KP@4.2)**, **(KP@5.2)**, that P_α is a belief set, that $P_\alpha = P_\beta$ if $\vdash \alpha \leftrightarrow \beta$, and **(KP@DF)**. **(KP@6)** follows immediately. We must prove only **(KP@7)** and that $K_\alpha \subseteq P_\alpha$. That $K_\alpha \subseteq P_\alpha$ follows directly from the definition, since $Th(f_S(\alpha) \cup ([K] \cap [\alpha])) \subseteq Th(f_S(\alpha))$. For **(KP@7)**, let $\neg\alpha \in P$. We have two cases: If $\neg\alpha \in K$, then $([K] \cap [\alpha]) = \emptyset$, from which it follows that $K_\alpha = P_\alpha$. If $\neg\alpha \notin K$, and due to $\neg\alpha \in P$, **(S6)** yields that $[K]$ is the minimal α -sphere, from which it follows that $Th(f_S(\alpha)) = ([K] \cap [\alpha])$, hence $K_\alpha = P_\alpha$. \square

Proof of Theorem 3.4:

Part 1: Postulates to Partial Meet:

From lemma A.3 it follows that if @ satisfies **(KP@1)**, **(KP@2)**, **(KP@4.1)**, **(KP@5.1')**, **(KP@6)**, and **(KP@DF.1)**, then there exists a transitively relational function g , such that $K_\alpha = (\bigcap g(K \perp \neg \alpha)) + \alpha$. By the same lemma, if @ satisfies **(KP@1)**, **(KP@2)**, **(KP@4.2)**, **(KP@5.2')**, **(KP@6)**, and **(KP@DF)**, then there is a transitively relational function h , such that $P_\alpha = (\bigcap h(P \perp \neg \alpha)) + \alpha$.

It suffices then to prove that $X \in g(K \perp \neg \alpha)$ if and only if there exists $Y \in h(P \perp \neg \alpha)$ such that $(Y \cap K) = X$.

There are three cases to be considered:

Case 1: If $\neg \alpha \notin P$, then $g(K \perp \neg \alpha) = \{K\}$ and $h(P \perp \neg \alpha) = \{P\}$, and the condition follows trivially.

Case 2: If $\neg \alpha \in P$ and $\neg \alpha \notin K$, then since $K \subseteq P$, for every $Y \in h(P \perp \neg \alpha)$ it holds that $K \subseteq Y$. Since $g(K \perp \neg \alpha) = \{K\}$, the condition follows immediately.

Case 3: If $\neg \alpha \in K$, then by postulate **(KP@7.1)** $\bigcap g(K \perp \neg \alpha) + \alpha = \bigcap h(P \perp \neg \alpha) + \alpha$. Since $K \subseteq P$, it is easy to see that $X \in K \perp \neg \alpha$ if and only if there is $Y \in P \perp \neg \alpha$ such that $Y \cap K = X$.

If $X \in g(K \perp \neg \alpha)$, then $\bigcap g(K \perp \neg \alpha) + \alpha \subseteq X + \alpha$. From this it follows that $\bigcap h(P \perp \neg \alpha) + \alpha \subseteq X + \alpha$. Let $Y \in P \perp \neg \alpha$ such that $Y \cap K = X$. Then $\bigcap h(P \perp \neg \alpha) + \alpha \subseteq Y + \alpha$ and hence, $Y \in h(P \perp \neg \alpha)$.

For the other inclusion, let $Y \in h(P \perp \neg \alpha)$ and let $X = Y \cap K$. We have $X \in K \perp \neg \alpha$. Then $\bigcap h(P \perp \neg \alpha) + \alpha \subseteq Y + \alpha$ and hence, $\bigcap g(K \perp \neg \alpha) + \alpha \subseteq Y + \alpha$. Since $\bigcap g(K \perp \neg \alpha) \subseteq K$, we have $\bigcap g(K \perp \neg \alpha) + \alpha \subseteq X + \alpha$ and hence, $X \in g(K \perp \neg \alpha)$.

Part 2: Partial Meet to Spheres:

From Lemma A.3 it follows that:

(1) $P_\alpha = (\bigcap h(P \perp \neg \alpha)) + \alpha$ and h transitively relational imply that there is a system of spheres S satisfying *(S1)*-*(S4)* such that $P_\alpha = Th(f_S(\alpha))$.

(2) $K_\alpha = (\bigcap g(K \perp \neg \alpha)) + \alpha$ and g transitively relational imply that there is a system of spheres S_K satisfying *(S1)*-*(S3)* and *(S5)* and having $[K]$ as minimum such that $K_\alpha = Th(f_{S_K}(\alpha))$, where $c_{S_K}(\alpha)$ is the minimal sphere of S_K intersecting $[\alpha]$ and $f_{S_K}(\alpha) = c_{S_K}(\alpha) \cap [\alpha]$.

We have to prove that (a) $S_K = \{U \in S \mid [K] \subseteq U\}$, and (b) $Th(f_{S_K}(\alpha)) = Th(f_S(\alpha) \cup ([K] \cap [\alpha]))$.

(a) Let $U \in S$ and let β be a sentence such that $U = c_S(\beta)$. We have to show that if $[K] \subseteq U$, then $U \in S_K$. From (1) we know that $Th(c_S(\beta)) = \bigcap h(P \perp \neg \beta)$. Since $X \in g(K \perp \neg \beta)$ if and only if there is $Y \in h(P \perp \neg \beta)$ such that $Y \cap K = X$, we have that $\bigcap h(P \perp \neg \beta) \cap K = \bigcap g(K \perp \neg \beta)$. This means that $c_S(\beta) \cup [K] = c_{S_K}(\beta)$ and hence, if $[K] \subseteq U$, then $c_S(\beta) \cup [K] = U = c_{S_K}(\beta)$ and so, $U \in S_K$.

(b) There are two cases:

Case 1: If $c_S(\alpha) \subset K$, then $c_{S_K} = [K]$ and $f_{S_K} = [K] \cap [\alpha]$. Since $f_S(\alpha) = c_S(\alpha) \cap [\alpha] \subseteq [K] \cap [\alpha]$, we have $f_S(\alpha) \cup ([K] \cap [\alpha]) = [K] \cap [\alpha]$

$[\alpha] = f_{S_K}(\alpha)$.

Case 2: If $K \subseteq c_S(\alpha)$, then $c_{S_K}(\alpha) = c_S(\alpha)$ and $[K] \cap [\alpha] = \emptyset$. Hence, $f_{S_K}(\alpha) = f_S(\alpha) \cup ([K] \cap [\alpha])$.

Part 3: Spheres to Postulates:

Since in theorem 2.11(S6) was used only to prove **(KP@7)**, we must only show that @ satisfies **(KP@7.1)**. If $\neg\alpha \in K$, then $[K] \cap [\alpha] = \emptyset$ and thus, $P_\alpha = K_\alpha$.

□

References

- [AGM85] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [AM82] Carlos Alchourrón and David Makinson. On the logic of theory change: contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.
- [Gär88] Peter Gärdenfors. *Knowledge in Flux - Modeling the Dynamics of Epistemic States*. MIT Press, 1988.
- [Gär99] Peter Gärdenfors. *Conceptual Spaces*. MIT Press - to appear, 1999.
- [GR95] Peter Gärdenfors and Hans Rott. Belief revision. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume IV, chapter 4.2. 1995.
- [Gro88] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [Lak87] George Lakoff. *Women, Fire and Dangerous Things*. The University of Chicago Press, 1987.
- [Leh98] Daniel Lehmann. Stereotypical reasoning: Logical properties. *Journal of the IGPL*, 6(1):49–58, 1998.
- [LR91] Sten Lindström and Włodzimierz Rabinowicz. Epistemic entrenchment with incomparabilities and relational belief revision. In André Fuhrmann and Michael Morreau, editors, *The Logic of Theory Change*, volume 465 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1991.
- [Mak85] David Makinson. How to give it up: a survey of some recent work on formal aspects of the logic of theory change. *Synthese*, 62, 1985.

- [MG91] David Makinson and Peter Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In André Fuhrmann and Michael Morreau, editors, *The Logic of Theory Change*, volume 465 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1991.
- [Pag96] Maurice Pagnucco. *The Role of Abductive Reasoning Within the Process of Belief Revision*. PhD thesis, Department of Computer Science, University of Sydney, 1996.
- [Ros73] Eleanor H. Rosch. Natural categories. *Cognitive Psychology*, 4:328–350, 1973.
- [Was98] Renata Wassermann. Revising concepts. In *Proceedings of the Fifth Workshop on Logic, Language, Information and Communication (WoLLIC)*, São Paulo, 1998.