Sven Ove Hansson

A Monoselective Presentation of AGM Revision

Abstract. A new equivalent presentation of AGM revision is introduced, in which a preference-based choice function directly selects one among the potential outcomes of the operation. This model differs from the usual presentations of AGM revision in which the choice function instead delivers a collection of sets whose intersection is the outcome. The new presentation confirms the versatility of AGM revision, but it also lends credibility to the more general model of direct choice among outcomes (descriptor revision) of which AGM revision is shown here to be a special case.

Keywords: Choice function, Selection function, AGM, Belief revision, Descriptor revision, Outcome set, Monoselective choice, Sphere mode, Grove model.

1. Introduction

Choice has a central role in the theory of belief change. In the standard formal framework an epistemic agent's belief state is represented by a logically closed set, called the belief set, that corresponds (simplistically) to the sentences she believes in or (in a more sophisticated interpretation) to the sentences she is committed to believe in. Operations of change take the form of replacing one such belief set by another that satisfies a given success condition. There are two major types of operations. In belief contraction, the success condition is that the new belief set does not contain a specified sentence, the sentence we contract by. In belief revision, the success condition is instead that some specified sentence has to be included in the new belief set. Clearly, there are many belief sets satisfying this condition, and we have to settle on one of them.

For the formal representation of choice, belief change theory uses the notion of a choice function that has been taken over from social choice theory [15]. Choice functions are also used in several other branches of logic [9].

Presented by **Heinrich Wansing**; *Received* July 20, 2014

Studia Logica (2015) 103: 1019–1033
DOI: 10.1007/s11225-015-9604-5

A choice function is defined over a set $\mathcal{A}$ of alternatives. It can be used to make a selection among any subset of $\mathcal{A}$. The formal definition is as follows:

Definition 1. C is a *choice function* for a set $\mathcal{A}$ if and only if for each subset $\mathcal{B}$ of $\mathcal{A}$:

(1) $C(\mathcal{B}) \subseteq \mathcal{B}$, and

(2) $C(\mathcal{B}) \neq \varnothing$ if $\mathcal{B} \neq \varnothing$.

A choice function C is *based on* a relation $\multimap$ if and only if for all $\mathcal{B}$ and all $X \in \mathcal{B}$:

$X \in C(\mathcal{B})$ if and only if $X \multimap Y$ for all $Y \in \mathcal{B}$.

It is important to note that a choice function can have multiple outcomes, i.e. $C(\mathcal{B})$ may have more than one element. When this happens in social choice theory, all elements of $C(\mathcal{B})$ are (considered to be) equally choiceworthy. It is then left to the decision-maker to further narrow down the choice to one single object. Which element of $C(\mathcal{B})$ she ends up with is presumed to be arbitrary from the viewpoint of rationality. Therefore, strictly speaking, the choice function only represents the first of two stages in the choice process. The second stage that slims down the outcome to one single element is often described as a matter of picking rather than choosing [16]. We can call this the *select-and-pick* method.

Choice functions with multiple outputs are also used in belief change. However, it is in general assumed that an adequate model of belief change must result in a single, determinate outcome.[1] Therefore, just as in social choice theory, a second stage has to be added in order to obtain a determinate outcome. In belief change, the second stage consists in forming the intersection of the sets chosen in the first stage, and this intersection is taken to be the outcome. This has been called the *select-and-intersect method* [8]. It is a conservative way to deal with a tie between several maximally choiceworthy alternatives: The outcome contains (only) those beliefs that are held in all those maximally choiceworthy alternatives.

There are two major versions of the select-and-intersect method. One is partial meet contraction that was introduced in the original AGM article [1]. This is a method to remove a sentence p from a belief set K. A choice function (usually denoted γ) is applied to the remainder set $K \perp p$ that

[1] A few studies have been devoted to indeterministic belief change operations. These are operations that deliver, for each input, a set that may contain more than one possible outcome [3,12].

is defined as the set of inclusion-maximal p-excluding subsets of K.[2] The outcome of contracting K by p is the intersection of all elements of $\gamma(K \perp p)$, i.e.

Partial meet contraction:
$K \div p = \bigcap \gamma(K \perp p)$

Revision, i.e. consistency-preserving addition of a sentence, is defined in terms of contraction via the so-called Levi identity:

Partial meet revision:
$K * p = \mathrm{Cn}((K \div \neg p) \cup \{p\})$ (where $\div$ is partial meet contraction)

These operations are said to be *transitively relational* if the selection function is based on a transitive relation in the manner shown in Definition 1. Transitively relational partial meet contraction and revision have been axiomatically characterized (the "AGM postulates") and are often seen as the gold standard of belief change theory.

The other version of the select-and-intersect method is most lucidly presentable for belief revision. When revising a belief set K by a sentence p we are looking for a new belief set that contains p. The largest (consistent) such belief sets are the possible worlds including p. By a possible world, in the logical sense, is meant a subset of the language that is so large that any addition of a new sentence to it will result in inconsistency (i.e. X is consistent but if $q \notin X$ then $X \cup \{q\}$ is inconsistent). In the first stage we apply a choice function to select the most choiceworthy of the possible worlds that contain p. In the second stage we form their intersection and take it to be the outcome $K * p$. This approach was introduced by Grove [4], who based it on Lewis's [11] account of counterfactuals. It is usually presented as a simple geometrical model, as illustrated in Fig. 1. Each point in the square represents a possible world. The circle in the middle contains exactly those possible worlds that are compatible with the current belief state K, and their intersection is equal to K. The whole system of circles ("spheres") represents an ordering of the possible worlds in terms of their choiceworthiness. The central circle that corresponds to K consists of the most choiceworthy worlds, the ring immediately surrounding it contains the worlds that come second in terms of choiceworthiness, etc. The outcome of revision by p is equal to the intersection of the set of p-worlds in the innermost circle that

[2] $X \in K \perp p$ if and only if X is a subset of K that does not imply p, and every set X' such that $X \subset X' \subseteq K$ implies p.

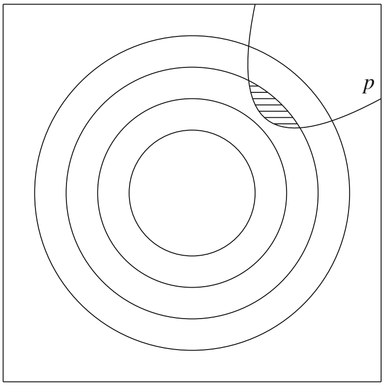

Figure 1. Revision by a sentence p that is incompatible with the present belief set K. The area covered by the central circle represents those possible worlds that contain K. The worlds in the central circle all have the same degree of choiceworthiness. Similarly, the worlds in each of the surrounding rings have the same degree of choiceworthiness. The degree of choiceworthiness decreases with the distance of the ring from the central circle. The area covered by the parabola represents those possible worlds that contain p. The shaded area represents the selected worlds, namely all the p-worlds in the sphere that contains the most choiceworthy p-worlds

contains some p-worlds. If p is logically compatible with K (i.e. $K \cup \{p\}$ is consistent), then the selected worlds will all be elements of the central circle. If the central circle does not contain any p-world, then we will instead select the p-worlds in the circle that contains the most choiceworthy p-worlds.

The two ways to construct revision, the possible worlds construction and the partial meet construction, yield exactly the same result. In other words, an operation on a belief set K is a transitively relational partial meet revision if and only if it can be constructed in the way indicated in Fig. 1. This equivalence is based on a one-to-one correspondence called "Grove's bijection" between remainders and possible worlds. [4,5, pp. 53–55]

However, both stages in these select-and-intersect procedures are problematic from an intuitive point of view. In the first stage, a selection is made among either remainders or possible worlds, but neither of these is itself a plausible belief set. If $K \div p$ is a remainder then it has the implausible property that for all q: either $p \vee q \in K \div p$ or $p \vee \neg q \in K \div p$ ([2, p. 29], [5, p. 124]). If $K * p$ is a possible world then it has the even more implausible property that for all sentences q either $q \in K * p$ or $\neg q \in K * p$. Therefore, in both cases the choice is indirect in the sense of being a choice, not among potential outcome candidates, but instead among other objects from which an outcome is constructed in a subsequent stage.

The second stage, intersection, is also debatable. It would certainly not be considered in social choice:

> GAME SHOW HOST: Congratulations! You have won the first prize. This means that you now have a choice between two options. One is a Porsche 991 and 50 litres of petrol. The other is a Lamborghini Huracán and 50 litres of petrol. Which of them do you choose?
>
> CONTESTANT: I am unable to choose between them. The two alternatives are exactly equally good.
>
> GAME SHOW HOST: Thanks for telling us. We will now follow our standard procedure for such cases of indecision, and give you the intersection between the two sets you could not choose between. One of the sets contained a Porsche 991 and 50 litres of petrol, and the other a Lamborghini Huracán and 50 litres of petrol. Let me congratulate you once more. You are now the happy owner of the intersection of those two sets, namely 50 litres of petrol, of the highest quality.

Perhaps more to the point, Sandqvist [14] has argued convincingly that similar situations may arise in belief change, i.e. the intersection of two or more optimal sets of beliefs need not itself be optimal. Based in part on these considerations, another approach to belief change has recently been proposed, namely descriptor revision [7]. It is a single-stage procedure that directly selects one of the potential outcomes of the operation. This requires that the set of potential outcomes (the outcome set) is available and that the choice function is monoselective in the following sense:

DEFINITION 2. [6] A choice function is *monoselective* if and only if it satisfies:

If $X \in C(\mathcal{B})$ and $Y \in C(\mathcal{B})$, then $X = Y$.

In this approach, a plausible version of sentential revision (revision by a sentence) can be introduced as follows :

DEFINITION 3. [7] An operation $*$ on a belief set K is a *linear sentential descriptor revision* if and only if there is a set $\mathbb{X}$ of belief sets with $K \in \mathbb{X}$ (its outcome set), and a total ordering[3] $\leqq$ on $\mathbb{X}$, such that (i) $K \leqq X$ for all

[3] A total ordering is a binary relation that is transitive, complete ($X \leqq Y$ or $Y \leqq X$, synonyms: connected, total) and anti-symmetric (If $X \leqq Y$ and $Y \leqq X$, then $X = Y$).

$X \in \mathbb{X}$, and (ii) for all sentences p: $K * p$ is the unique $\leqq$-minimal element of $\mathbb{X}$ that contains $*$, unless $p \notin \bigcup \mathbb{X}$, in which case $K * p = K$.

This is a special case of a much more general framework that easily accommodates a wide variety of other success conditions than that of sentential revision [7]. However, the focus here will be on comparisons with AGM revision, and therefore we will only have use for the case presented in Definition 3. We are going to show that the standard AGM operator of revision (transitively relational partial meet revision) is in fact a special case of linear sentential descriptor revision. After some formal preliminaries have been presented in Sect. 2, the new formal result will be presented in Sect. 3 and further discussed in Sect. 4. The formal proof of the main result is deferred to an Appendix.

2. Formal Preliminaries

The belief-representing sentences form a language $\mathcal{L}$. Sentences, i.e. elements of this language, are represented by lowercase letters $(a, b, \ldots)$ and sets of sentences by uppercase letters $(A, B, \ldots)$. The language contains the usual truth-functional connectives: negation $(\neg)$, conjunction $(\&)$, disjunction $(\vee)$, implication $(\rightarrow)$, and equivalence $(\leftrightarrow)$.

A Tarskian consequence operator Cn expresses the logic. It is a function from and to sets of sentences. Intuitively speaking, for any set A of sentences, $\mathrm{Cn}(A)$ is the set of logical consequences of A. Cn satisfies the standard conditions: inclusion $(A \subseteq \mathrm{Cn}(A))$, monotony (If $A \subseteq B$, then $\mathrm{Cn}(A) \subseteq \mathrm{Cn}(B)$), and iteration $(\mathrm{Cn}(A) = \mathrm{Cn}(\mathrm{Cn}(A)))$. Furthermore, Cn is supraclassical (if p follows from A by classical truth-functional logic, then $p \in \mathrm{Cn}(A)$) and compact (if $p \in \mathrm{Cn}(A)$ then there is a finite subset A' of A such that $p \in \mathrm{Cn}(A')$), and it satisfies the deduction property ($q \in \mathrm{Cn}(A \cup \{p\})$ if and only if $(p \rightarrow q) \in \mathrm{Cn}(A)$).

$\mathrm{Cn}(\varnothing)$ is the set of tautologies. $X \vdash p$ is an alternative notation for $p \in \mathrm{Cn}(X)$ and $\vdash p$ for $p \in \mathrm{Cn}(\varnothing)$.

A set A of sentences is a (consistent) *belief set* if and only if it is consistent and logically closed, i.e. $A = \mathrm{Cn}(A)$. K denotes a belief set.

For any set A of sentences and sentence p, the *remainder set* $A \perp p$ consists of those sets X for which $X \subseteq A$, $X \nvdash p$, and there is no X' such that $X \subset X' \subseteq A$ and $X' \nvdash p$.

3. AGM Operations in Monoselective Guise

The following theorem provides an equivalent presentation of transitively relational partial meet revision. In doing so it shows that transitively relational partial meet revision is a special case of linear sentential descriptor revision that was introduced in Definition 3.

THEOREM 1. *Let $*$ be a sentential operation on a belief set K, with the outcome set $\mathbb{X}$. Then the following two conditions are equivalent:*

(I) $*$ *is a transitively relational partial meet revision.*

(II) $\mathbb{X}$ *satisfies:*

($\mathbb{X}$1) *If $X \in \mathbb{X}$ then $X = \mathrm{Cn}(X)$.*
($\mathbb{X}$2) *For all $X, Y \in \mathbb{X}$: if $X \cup Y \nvdash \bot$ then $X \cap Y \in \mathbb{X}$.*

and there is a relation $\leqq$ on $\mathbb{X}$ such that:

($\leqq$1) *$\leqq$ is a total ordering (transitive, complete, and antisymmetric).*
($\leqq$2) *$K \leqq X$ for all $X \in \mathbb{X}$.*
($\leqq$3) *If $X \subseteq Y$ then $X \leqq Y$.*
($\leqq$4) *If $X \cap Z \in \mathbb{X}$ and $X \leqq Y \leqq Z$ then $X \cap Y \in \mathbb{X}$ and $Y \cap Z \in \mathbb{X}$.*
($\leqq$5) *$K * p$ is the unique $\leqq$-minimal p-containing element of $\mathbb{X}$.*

PROOF. See the Appendix. ■

Properties ($\mathbb{X}$1), ($\leqq$1), ($\leqq$2), and ($\leqq$5) follow from Definition 3. They hold for all linear sentential descriptor revisions. It is the remaining three properties, ($\mathbb{X}$2), ($\leqq$3), and ($\leqq$4), that characterize transitively relational AGM revision in contradistinction to other linear sentential descriptor revisions. These properties may be somewhat opaque but they can be understood in relation to the AGM postulates for revision. The following are five postulates that hold for transitively relational AGM revision:

$p \in K * p$ (success)

$K*(p \vee q) = K*p$ or $K*(p \vee q) = K*q$ or $K*(p \vee q) = (K*p) \cap (K*q)$ (disjunctive factoring)

If $\neg p \notin K * (p \vee q)$ then $K * (p \vee q) \subseteq K * p$ (disjunctive inclusion)

$(K * p) \cap (K * q) \subseteq K * (p \vee q)$ (disjunctive overlap)

$K * p = K * q$ if and only if $q \in K * p$ and $p \in K * q$ (reciprocity)

(For more information on these postulates, see [5, pp. 270–274] and [13, pp. 107–111].)

For ($\mathbb{X}$2), note first that since X and Y are assumed to be elements of $\mathbb{X}$ we can assume that there are p and q such that $X = K * p$ and $Y = K * q$. With this substitution, we are going to derive the contrapositive form of ($\mathbb{X}$2), i.e.:

$$\text{If } (K * p) \cap (K * q) \notin \mathbb{X}, \text{ then } (K * p) \cup (K * q) \vdash \bot. \tag{1}$$

Let $(K*p) \cap (K*q) \notin \mathbb{X}$. Then $K*(p \vee q) \neq (K*p) \cap (K*q)$. It follows from disjunctive factoring that either $K * (p \vee q) = K * p$ or $K * (p \vee q) = K * q$. Without loss of generality we can assume that $K * (p \vee q) = K * p$. Now suppose for contradiction that $\neg q \notin K*p$. Then equivalently $\neg q \notin K*(p \vee q)$, and it follows from disjunctive inclusion that $K*(p \vee q) \subseteq K*q$, equivalently $K*p \subseteq K*q$. Then it follows from $K*(p \vee q) = K*p$ and $K*p \subseteq K*q$ that $K * (p \vee q) = (K * p) \cap (K * q)$, contrary to our initial assumption. We can conclude that $\neg q \in K*p$. Due to success, $q \in K*q$, thus $(K*p) \cup (K*q) \vdash \bot$ as desired.

Next, let us turn to ($\leqq$3). Since X and Y are assumed to be elements of $\mathbb{X}$ we can replace them by $K * p$ and $K * q$. Since $\leqq$ is antisymmetric it would be sufficient to show that:

$$\text{If } K * p \subset K * q \text{ then } K * p < K * q. \tag{2}$$

However, (2) cannot be derived from the AGM postulates since the language of those postulates does not contain $\leqq$ or $<$. Instead we can show that $*$ has a property that is necessary for (2) to hold. If $q \in K*p$, then ($\leqq$5) makes it impossible for $K*p < K*q$ to hold. Therefore, in order for (2) (and ($\leqq$3)) to hold, the following condition must be satisfied:

$$\text{If } K * p \subset K * q \text{ then } q \notin K * p. \tag{3}$$

To show that (3) holds, let $K*p \subset K*q$. Success yields $p \in K*q$. Suppose for contradiction that $q \in K*p$. Then reciprocity yields $K*p = K*q$, contrary to our assumption that $K * p \subset K * q$. We can conclude that $q \notin K * p$, as desired.

Finally, ($\leqq$4) is somewhat more complex but it can be understood with the help of the following property [7]:

$$\text{If } K * z = (K * p) \cap (K * q) \text{ then } K * (p \vee q) = (K * p) \cap (K * q). \tag{4}$$

To show that (4) follows from the AGM postulates, let $K*z = (K*p) \cap (K*q)$. Success yields $p \vee q \in K * z$ and $z \in (K * p) \cap (K * q)$. Due to disjunctive overlap, $z \in K * (p \vee q)$. Finally, we apply reciprocity to $p \vee q \in K * z$ and $z \in K * (p \vee q)$, and obtain $K * z = K * (p \vee q)$.

The following equivalent form of (4) will be useful:

$$(K * p) \cap (K * q) \in \mathbb{X} \text{ if and only if } K * (p \vee q) = (K * p) \cap (K * q). \tag{5}$$

We can interpret $K * (p \vee q) = (K * p) \cap (K * q)$ as saying that if we accept the information that either p or q, then we enter a state of hesitation between revising by p and revising by q, presumably because these two alternatives are equally plausible. Using (5) we can therefore interpret $(K * p) \cap (K * q) \in \mathbb{X}$ as saying that as seen from the viewpoint of K, $K * p$ and $K * q$ are equally plausible. In this perspective, ($\leqq$4) can be read as saying that if the belief sets $K * p$ and $K * r$ are equally plausible, and $K * p \leqq K * q \leqq K * r$, then $K * q$ is equally plausible as $K * p$, and also equally plausible as $K * r$.

4. Conclusion

AGM operations are commonly constructed with the select-and-intersect method, which means that the choice function chooses among a set of objects (remainders or possible worlds) none of which is a plausible outcome of the intended operation. Even if the outcome turns out to be plausible, such choice mechanisms are problematic from the viewpoint of justification. It would seem more reasonable to use a choice function that selects among the potential outcomes, and delivers one of them as the outcome of the operation.

The theorem proved in this paper partly defuses the conflict between AGM revision and the idea of such a direct choice among potential outcomes. We have shown that transitively relational AGM revision can be reconstructed with a transitively relational choice function that selects directly among the potential outcomes. This theorem lends credibility to AGM revision but also to the more general model of such direct choice (descriptor revision) of which AGM revision has now been shown to be a special case.

Appendix: Proof of the Theorem

We are going to use systems of spheres as defined by Grove [4], and the proof is therefore preceded by some definitions and preparatory results for such spheres.

Definition 4. Let K be a belief set and $*$ a sentential operation on K. The outcome set of $*$ is the set $\{K * p \mid p \in \mathcal{L}\}$.

Definition 5. For all $A \subseteq \mathcal{L}$: $[A] = \{W \in \mathcal{L} \perp \bot \mid A \subseteq W\}$.

Brackets of singletons can be omitted, thus $[p] = [\{p\}]$.

Observation 1. $\bigcap[A] = \mathrm{Cn}(A)$

Proof. See Hansson [5, p. 52]. ∎

Lemma 1. (1) $[X] = [Y]$ *if and only if* $\mathrm{Cn}(X) = \mathrm{Cn}(Y)$

(2) $[X] \subseteq [Y]$ *if and only if* $\mathrm{Cn}(Y) \subseteq \mathrm{Cn}(X)$

(3) *If* X *and* Y *are logically closed, then* $[X \cup Y] = [X] \cap [Y]$

(4) *If* X *and* Y *are logically closed, then* $[X \cap Y] = [X] \cup [Y]$

Proof. *Parts 1–3:* See Hansson [5, pp. 52–53]. ∎

Part 4: Right-to-left: It follows from Part 2 that $[X] \subseteq [X \cup Y]$ and $[Y] \subseteq [X \cup Y]$

Left-to-right: Suppose to the contrary that $[X \cap Y] \nsubseteq [X] \cup [Y]$. Then there is a possible world W such that $X \cap Y \subseteq W$, $X \nsubseteq W$, and $Y \nsubseteq W$, and there must be some $x \in X$ such that $x \notin W$ and some some $y \in Y$ such that $y \notin W$. It follows from $x \vee y \in X \cap Y$ and $X \cap Y \subseteq W$ that $x \vee y \in W$. Furthermore, it follows from $x \notin W$ and the maximality of W that $\neg x \in W$, and $\neg y \in W$ follows in the same way. Thus W is inconsistent which it cannot be since it is a possible world. We can conclude from this contradiction that $[X \cap Y] \subseteq [X] \cup [Y]$.

Systems of spheres are defined as follows:

Definition 6. Let $\mathcal{W} \subseteq \mathcal{L} \perp \bot$. A set $\mathfrak{S}$ of subsets of $\mathcal{L} \perp \bot$ is a *system of spheres centered on* $\mathcal{W}$ if and only if:

($\mathfrak{S}$1) If $\mathcal{S}_1, \mathcal{S}_2 \in \mathfrak{S}$, then either $\mathcal{S}_1 \subseteq \mathcal{S}_2$ or $\mathcal{S}_2 \subseteq \mathcal{S}_1$.

($\mathfrak{S}$2) $\mathcal{W} \in \mathfrak{S}$, and $\mathcal{W} \subseteq \mathcal{S}$ for all $\mathcal{S} \in \mathfrak{S}$.

($\mathfrak{S}$3) $\mathcal{L} \perp \bot \in \mathfrak{S}$.

($\mathfrak{S}$4) For all sentences p and spheres $\mathcal{S} \in \mathfrak{S}$, if $[p] \cap \mathcal{S} \neq \emptyset$, then there is some $\mathcal{S}' \in \mathfrak{S}$ such that $[p] \cap \mathcal{S}' \neq \emptyset$ and that $[p] \cap \mathcal{S}'' = \emptyset$ for all $\mathcal{S}'' \in \mathfrak{S}$ with $\mathcal{S}'' \subset \mathcal{S}'$.

($\mathfrak{S}$1) tells us that spheres are concentric (totally ordered by set inclusion). ($\mathfrak{S}$2) says that $\mathcal{W}$ itself is the minimal sphere. Without loss of generality, we will assume that $\mathcal{W} = [K]$ for some belief set K. ($\mathfrak{S}$3) says that the set of all possible worlds is itself the maximal sphere. ($\mathfrak{S}4$) ensures that for all sentences p there is a minimal sphere containing some p-world.

DEFINITION 7. Let $\mathfrak{S}$ be a system of spheres. A *world-ring* in $\mathfrak{S}$ is s set $\mathcal{R}$ consisting of all elements of some $\mathcal{S} \in \mathfrak{S}$ that are not elements of any $\mathcal{S}' \in \mathfrak{S}$ with $\mathcal{S}' \subset \mathcal{S}$.

Obviously, the world-rings are equivalence classes.

DEFINITION 8. Let K be a belief set. A revision operator $*$ on K is *based on* a system $\mathfrak{S}$ of spheres centered on $[K]$ if and only if for all p, $[K * p]$ is the intersection of all p-worlds in the smallest sphere in $\mathfrak{S}$ that contains some p-world; unless there are no p-worlds, in which case $[K * p] = \varnothing$.

Grove's [4] main result was that sphere-based revision coincides exactly with transitively relational partial meet revision. We will have use for a few more results on systems of spheres.

LEMMA 2. *Let $\mathbb{X}$ be the outcome set of a sphere-based revision in a sphere system $\mathfrak{S}$, and let $X \in \mathbb{X}$. Then all worlds W with $X \subseteq W$ are elements of the same ring in $\mathfrak{S}$.*

PROOF. From Observation 1 and Definition 8. ■

LEMMA 3. *Let $\mathbb{X}$ be the outcome set of a sphere-based revision, and let $X, Y \in \mathbb{X}$. Furthermore, let $W_1, W_2 \in \mathcal{L} \perp\perp$ be such that $X \subseteq W_1$ and $Y \subseteq W_2$. Then: $X \cap Y \in \mathbb{X}$ if and only if W_1 and W_2 are elements of the same world-ring.*

PROOF. Since $X, Y \in \mathbb{X}$ there are p and q such that $X = K*p$ and $Y = K*q$.

For one direction, let W_1 and W_2 be elements of the same world-ring. It follows from Lemma 2 that the innermost world-ring containing some $K*p$-world coincides with the innermost world-ring containing some $K*q$-world. Let $\mathcal{R}$ be that world-ring. Then $[K*(p \vee q)] = \mathcal{R} \cap [p \vee q]$, $[K*p] = \mathcal{R} \cap [p]$, and $[K*q] = \mathcal{R} \cap [q]$. We can conclude that $K*(p \vee q) = (K*p) \cap (K*q)$, thus $X \cap Y \in \mathbb{X}$.

For the other direction, let $X \cap Y \in \mathbb{X}$. Due to Lemma 1, $[X \cap Y] = [X] \cup [Y]$, and it follows from Lemma 2 that the worlds containing X and the worlds containing Y all belong to the same ring. ■

LEMMA 4. *If $\mathbb{X}$ is the outcome set of a sphere-based revision, then: If $X \cap Y \in \mathbb{X}$ and $Y \cap Z \in \mathbb{X}$, then $X \cap Z \in \mathbb{X}$.*

PROOF. Directly from Lemma 3. ■

LEMMA 5. *Let $\mathbb{X}$ be the outcome set of a sphere-based revision, and let $X, Y \in \mathbb{X}$. If $X \cup Y \nvdash \bot$, then $X \cap Y \in \mathbb{X}$.*

PROOF. Since $X \cup Y \nvdash \bot$ there is some world $W \in \mathcal{L} \perp \bot$ such that $X \cup Y \subseteq W$. It follows from Lemma 3 that $X \cap Y \in \mathbb{X}$. ■

We are now ready to prove the theorem:

PROOF OF THEOREM 1. We will make use of the well-known result from Grove [4] showing that (I) is equivalent with:

(I$^+$) $*$ is a revision operator based on some sphere system $\mathfrak{S}$.

The proof will therefore proceed by showing the equivalence between (I$^+$) and (II).

FROM (I$^+$) to (II): ($\mathbb{X}$1) follows directly and ($\mathbb{X}$2) follows from Lemma 5. For the rest of the proof, we define for each world-ring $\mathcal{R}$, a *cluster* $\overline{\mathcal{R}}$ of elements of $\mathbb{X}$:

$\overline{\mathcal{R}} = \{X \in \mathbb{X} \mid (\exists W \in \mathcal{R})(X \subseteq W\}$

It follows from Lemma 2 that each belief set is an element of exactly one cluster, and from Lemma 3 that the relation $\sqsupseteq\!\!\!\sqsubseteq$ on $\mathbb{X}$ such that $X \sqsupseteq\!\!\!\sqsubseteq Y$ iff $X \cap Y \in \mathbb{X}$ is an equivalence relation with the clusters as equivalance classes.

Next we define the relation $\sqsubseteq$ on $\mathbb{X}$ such that $X \sqsubseteq Y$ if and only if the world-ring corresponding to the cluster containing X is included in every sphere that contains the world-ring corresponding to the cluster containing Y. Clearly, $\sqsupseteq\!\!\!\sqsubseteq$ is the symmetric part of $\sqsubseteq$. Its strict part is denoted $\sqsubset$.

Let $\breve{\subset}$ be a strict linear ordering on $\mathbb{X}$ satisfying the condition: If $X \subset Y$ then $X \breve{\subset} Y$. (The existence of such a relation is guaranteed by the order extension principle that follows from the axiom of choice, see Jech [10], p. 19.) We let $<$ be the relation on $\mathbb{X}$ such that for all $X, Y \in \mathbb{X}$:

$X < Y$ if and only if either (i) $X \sqsubset Y$ or (ii) both $X \sqsupseteq\!\!\!\sqsubseteq Y$ and $X \breve{\subset} Y$.

Furthermore, let $X \leqq Y$ be the total ordering such that $X \leqq Y$ iff either $X < Y$ or $X = Y$.

We need to show that the conditions ($\leqq$1), ($\leqq$2), ($\leqq$3), ($\leqq$4), and ($\leqq$5) are satisfied. ($\leqq$1) and ($\leqq$2) follow directly.

For ($\leqq$3), note that if $X \subseteq Y$ then it follows from Lemma 2 that X and Y are included only in worlds in one and the same world-ring, thus they belong to the same cluster, i.e. $X \sqsupseteq\!\!\!\sqsubseteq Y$. Since $X \subseteq Y$ we have either $X = Y$ or $X \breve{\subset} Y$, and in both cases $X \leqq Y$ follows directly.

For ($\leqq$4), we use Lemma 3 to conclude that X and Z belong to the same cluster. Due to the definition of $\leqq$, Y belongs to that same cluster. Again using Lemma 3 we find that $X \cap Y \in \mathbb{X}$ and $Y \cap Z \in \mathbb{X}$.

For ($\leqq 5$): $K * p$ is the intersection of the p-worlds in the innermost world-ring that has p-worlds. Therefore it is an element of the corresponding cluster. Let X be a p-containing element of $\mathbb{X}$ such that $X \neq K * p$. If X belongs to the same cluster as $K * p$, then $K * p \sqsubseteq X$ and $K * p \subset X$, and if it belongs to some other cluster then $K * p \sqsubset X$. Thus in both cases $K * p < X$.

FROM (II) TO (I^+): We are first going to define an equivalence relation on the set consisting of those worlds that contain at least one element of $\mathbb{X}$, namely the relation $\sim$ such that:

$W_1 \sim W_2$ iff there are $X_1, X_2 \in \mathbb{X}$ such that $X_1 \subseteq W_1$, $X_2 \subseteq W_2$, and $X_1 \cap X_2 \in \mathbb{X}$.

This relation is obviously reflexive and symmetric. To prove that it is an equivalence relation it remains to show that it is transitive. Let $W_1 \sim W_2$ and $W_2 \sim W_3$. Then there are $X_1, X_2, Y_2, Y_3 \in \mathbb{X}$ such that $X_1 \subseteq W_1$, $X_2, Y_2 \subseteq W_2$, $Y_3 \subseteq W_3$, $X_1 \cap X_2 \in \mathbb{X}$ and $Y_2 \cap Y_3 \in \mathbb{X}$.

It follows from ($\mathbb{X}$2) that $X_2 \cap Y_2 \in \mathbb{X}$. Since $X_1 \cap X_2 \in \mathbb{X}$ another use of ($\mathbb{X}$2) yields $X_1 \cap X_2 \cap Y_2 \in \mathbb{X}$. We also have $Y_2 \cap Y_3 \in \mathbb{X}$, and a third use of ($\mathbb{X}$2) yields $X_1 \cap X_2 \cap Y_2 \cap Y_3 \in \mathbb{X}$. Combining this with $X_1 \cap X_2 \subseteq W_1$ and $Y_2 \cap Y_3 \subseteq W_3$, we obtain $W_1 \sim W_3$.

We will use the equivalence classes over worlds based on $\sim$ as world-rings, and define the following relations over these world-rings:

$\mathcal{R} \lhd \mathcal{R}'$ if and only if it holds for all $X, X' \in \mathbb{X}$ and all $W, W' \in \mathcal{L} \perp \perp$ that if $X \subseteq W \in \mathcal{R}$ and $X' \subseteq W' \in \mathcal{R}'$ then $X < X'$.

$\mathcal{R} \unlhd \mathcal{R}'$ if and only if either $\mathcal{R} \lhd \mathcal{R}'$ or $\mathcal{R} = \mathcal{R}'$.

($<$ is the strict part of $\leqq$.) If there are any worlds not containing any element of $\mathbb{X}$ then they are added as the lowest-ranked world-ring.

We need to show that $\unlhd$ is a total ordering of the world-rings, i.e. that it is transitive, complete and antisymmetric.

To show the transitivity of $\unlhd$, let $\mathcal{R}_1 \unlhd \mathcal{R}_2$ and $\mathcal{R}_2 \unlhd \mathcal{R}_3$. Excluding trivial cases we can assume that $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{R}_3$ are pairwise non-identical. Thus $\mathcal{R}_1 \lhd \mathcal{R}_2$ and $\mathcal{R}_2 \lhd \mathcal{R}_3$, and we suppose for contradiction that $\mathcal{R}_1 \not\lhd \mathcal{R}_3$. Then there are $X_1, X_3 \in \mathbb{X}$ and $W_1, W_3 \in \mathcal{L} \perp \perp$ such that $X_1 \subseteq W_1 \in \mathcal{R}_1$ and $X_3 \subseteq W_3 \in \mathcal{R}_3$ and $X_1 \not\lhd X_3$. ($\leqq$1) yields $X_3 \leqq X_1$. Since $\mathcal{R}_1$ and $\mathcal{R}_2$ are disjoint, $X_3 < X_1$. Let $X_2 \subseteq W_2 \in \mathcal{R}_2$. Due to $\mathcal{R}_1 \lhd \mathcal{R}_2$ and $\mathcal{R}_2 \lhd \mathcal{R}_3$ we have $X_1 < X_2$ and $X_2 < X_3$. This makes $<$ cyclic, contrary to ($\leqq$1).

Antisymmetry of $\trianglelefteq$: Suppose to the contrary that $\mathcal{R} \trianglelefteq \mathcal{R}'$, $\mathcal{R}' \trianglelefteq \mathcal{R}$ and $\mathcal{R} \neq \mathcal{R}'$. Then $\mathcal{R} \lhd \mathcal{R}' \lhd \mathcal{R}$. Let $X \subseteq W \in \mathcal{R}$ and $X' \subseteq W' \in \mathcal{R}$. Then $X < X' < X$, contrary to ($\leqq$1).

As a preparation for the proving the completeness of $\trianglelefteq$ we prove the following:

(X) If $\mathcal{R}_1 \neq \mathcal{R}_2$, $X_1 \subseteq W_1 \in \mathcal{R}_1$, $X_1' \subseteq W_1' \in \mathcal{R}_1$, $X_2 \subseteq W_2 \in \mathcal{R}_2$, $X_2' \subseteq W_2' \in \mathcal{R}_2$ and $X_1 < X_2$, then $X_1' < X_2'$.

Proof of (X): First suppose that $X_2' < X_1$. We then have $X_2' < X_1 < X_2$. Due to our definition of world-rings as $\sim$-equivalence classes, this contradicts ($\leqq$4). Since $\leqq$ is complete due to ($\leqq$1) we can conclude that $X_1 \leqq X_2'$. Next suppose that $X_2' < X_1'$. Then due to the result we just obtained we have $X_1 \leqq X_2' < X_1'$, again contrary to ($\leqq$4). Thus not $X_2' < X_1'$, and since $\mathcal{R}_1$ and $\mathcal{R}_2$ are disjoint, it follows from the completeness and antisymmetry of $\leqq$ that $X_1' < X_2'$ as desired.

Completeness of $\trianglelefteq$: For the two distinct world-rings $\mathcal{R}$ and $\mathcal{R}'$, let $X \subseteq W \in \mathcal{R}$ and $X' \subseteq W' \in \mathcal{R}'$, Since $\mathcal{R}$ and $\mathcal{R}'$ are disjoint, it follows from the completeness and antisymmetry of $\leqq$ that either $X < X'$ or $X' < X$. In the former case, (X) yields $\mathcal{R} \lhd \mathcal{R}'$ and in the latter case $\mathcal{R}' \lhd \mathcal{R}$.

Next, let $\mathfrak{S}$ be the set consisting of the sets $\mathcal{S}$ such that $\mathcal{S} = \bigcup\{\mathcal{R}' \mid \mathcal{R}' \trianglelefteq \mathcal{R}\}$ for some world-ring $\mathcal{R}$. We need to show that $\mathfrak{S}$ satisfies the four conditions given in Definition 6 for being a system of spheres centered on $[K]$, and that $*$ is based on $\mathfrak{S}$ in the manner described in Definition 8.

($\mathfrak{S}$1) follows directly from the construction of $\mathfrak{S}$. For ($\mathfrak{S}$2), note that it follows from ($\leqq$1) and ($\leqq$2) that $K < X$ for all $X \in \mathbb{X} \setminus \{K\}$. It follows from ($\leqq$3) that $\mathbb{X}$ contains no proper subset of K, from the construction of $\sim$ that the K-containing worlds form a world-ring of their own, and from our definition of $\mathfrak{S}$ that this world-ring is also the innermost sphere.

($\mathfrak{S}$3) follows from the construction of $\mathfrak{S}$, since all worlds are included in one of the world-rings.

For ($\mathfrak{S}$4), let $[p] \cap \mathcal{S} \neq \varnothing$. There is then some p-containing world, and according to ($\leqq$1) and ($\leqq$5) there is some $X_p \in \mathbb{X}$ such that $X_p < Y$ for all $Y \in \mathbb{X}$ with $p \in Y$. Let $\mathcal{R}_p$ be the world-ring containing the worlds including X_p. Then it holds for all world-rings $\mathcal{R}'$ that if $\mathcal{R}' \lhd \mathcal{R}_p$ then $\mathcal{R}'$ has no element containing p. It follows that the sphere $\mathcal{S}_p = \bigcup\{\mathcal{R}' \mid \mathcal{R}' \trianglelefteq \mathcal{R}_p\}$ has the desired property, i.e. it has a non-empty intersection with $[p]$ but no sphere that is its proper subset has a non-empty intersection with $[p]$.

Finally, it follows from the construction of the world rings that for all worlds W, if $X_p \subseteq W$ then $W \in \mathcal{R}_p$. Due to ($\mathbb{X}$1) and Observation 1,

$X_p = \bigcap\{W \in \mathcal{R}_p \mid p \in W\} = \bigcap\{W \in \mathcal{S}_p \mid p \in W\}$, i.e. X_p is the outcome of $\mathfrak{S}$-based revision of K by p according to Definition 8. Due to ($\leqq$ 5), $K * p = X_p$. This concludes the proof. ■

References

[1] Alchourrón, C., P. Gärdenfors, and D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, *Journal of Symbolic Logic* 50:510–530, 1985.

[2] Alchourrón, C., and D. Makinson, On the logic of theory change: Contraction functions and their associated revision functions, *Theoria* 48:14–37, 1982.

[3] Galliers, J. R., Autonomous belief revision and communication, in P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, Cambridge, 1992, pp. 220–246.

[4] Grove, A., Two modellings for theory change, *Journal of Philosophical Logic* 17:157–170, 1988.

[5] Hansson, S. O., *A Textbook of Belief Dynamics. Theory Change and Database Updating*, Kluwer, Dordrecht, 1999.

[6] Hansson, S. O., Maximal and perimaximal contraction, *Synthese* 190:3325–3348, 2013.

[7] Hansson, S. O., Descriptor revision, *Studia Logica* 102:955–980, 2014.

[8] Hansson, S. O., Back to basics: Belief revision through direct selection, submitted manuscript, 2014.

[9] Hansson, S. O., and P. Gärdenfors, David Makinson and the extension of classical logic, in S. O. Hansson (ed.), *David Makinson on Classical Methods for Non-Classical Problems*, Springer, Dordrecht, 2014, pp. 11–18.

[10] Jech, T. J., *The axiom of choice*, Dover Publications, New York, 2008.

[11] Lewis, D., *Counterfactuals*, Blackwell, Oxford, 1973.

[12] Lindström, S., and W. Rabinowicz, Epistemic entrenchment with incomparabilities and relational belief revision, in A. Fuhrmann and M. Morreau (eds.), *The Logic of Theory Change*, Springer, New York, 1991, pp. 93–126.

[13] Rott, H., *Change, Choice and Inference*, Oxford University Press, Oxford, 2001.

[14] Sandqvist, T., On why the best should always meet, *Economics and Philosophy* 16:287–313, 2000.

[15] Sen, A., *Collective Choice and Social Welfare*. Holden-Day, San Francisco, 1970.

[16] Ullmann-Margalit, E., and S. Morgenbesser, Picking and choosing, *Social Research* 44:757–785, 1977.

S. O. Hansson
Royal Institute of Technology (KTH)
Stockholm
Sweden
soh@kth.se