SVEN OVE HANSSON

Back to Basics: Belief Revision Through Direct Selection

Abstract. Traditionally, belief change is modelled as the construction of a belief set that satisfies a success condition. The success condition is usually that a specified sentence should be believed (revision) or not believed (contraction). Furthermore, most models of belief change employ a select-and-intersect strategy. This means that a selection is made among primary objects that satisfy the success condition, and the intersection of the selected objects is taken as outcome of the operation. However, the select-and-intersect method is difficult to justify, in particular since the primary objects (usually possible worlds or remainders) are not themselves plausible outcome candidates. Some of the most controversial features of belief change theory, such as recovery and the impossibility of Ramsey test conditionals, are closely connected with the select-and-intersect method. It is proposed that a selection mechanism should instead operate directly on the potential outcomes, and select only one of them. In this way many of the problems that are associated with the select-and-intersect method can be avoided. This model is simpler than previous models in the important Ockhamist sense of doing away with intermediate, cognitively inaccessible objects. However, the role of simplicity as a choice criterion in the direct selection among potential outcomes is left as an open issue.

Keywords: Belief change, Select-and-intersect, Recovery, Expansion property, Finiteness, Ramsey test, Direct selection, Simplicity, Choice function, Selection function, Support function, General input assimilation, Descriptor revision.

1. Introduction

The logic of belief change has been the subject of extensive studies since the 1980s, and a large number of models of belief change operations have been developed [4]. Most of these models have several features in common, including the following:

- Belief states are represented by sets of belief-representing sentences. In most models these sets are closed under logical consequence.

Presented by **Jacek Malinowski**; *Received* May 16, 2018

- Changes either take the form of the removal of some specified sentence(s) from these sets or the incorporation of some new sentence(s).
- The choice which beliefs to retain and which to give up in these operations is represented by some selection function (choice function).
- The selection function typically selects a large number of options, and the actual outcome is obtained by set-theoretical intersection of these options.

The present contribution has two purposes: First, it will show that some of these features, in particular the use of set-theoretical intersection, are difficult to justify and indeed closely connected with some of the most criticized properties of the standard model. Secondly, it will show how a model can be built and justified that avoids the problematic properties.

Section 2 provides a brief introduction to the standard model (the AGM model). Sections 3, 4, 5 and 6 investigate the use of intersection to solve ties and identify some problems in the application of that method. Section 7 shows that this construction gives rise to some of the more controversial properties of the AGM model. Section 8 presents a general model for belief change in which neither belief states nor inputs are represented by sentences. In Section 9, sentential representation is reintroduced in an arguably more cautious way that avoids the problems discussed in the previous sections. In the concluding Section 10, the properties of the new construction are compared to the problems of previous constructions identified in the previous sections.

2. A Crash Course in AGM Theory

In belief change theory, a belief state is usually represented by a belief set, i.e. a logically closed set of sentences (usually denoted K). All changes result from inputs, usually inputs specifying a sentence either to be believed or disbelieved. There are three basic types of operations, corresponding to three types of such instructions: “remove this sentence”, “incorporate this sentence”, and “incorporate this sentence and retain consistency”. In the belief revision literature, these instructions are (1) supplemented with various rationality constraints, and/or (2) specified in terms of the constructions by which they should be performed.

The instruction “remove this sentence” is performed with an operation of *contraction*, usually denoted $\div$. Thus $K \div p$ is the outcome of removing p from K . Contraction is assumed to satisfy the following postulates:

$$\begin{array}{ll}
K \div p \subseteq K & \text{(inclusion)} \\
K \div p = \text{Cn}(K \div p) & \text{(closure)} \\
p \notin K \div p, \text{ unless } p \text{ is a logical truth} & \text{(success)}
\end{array}$$

The instruction “incorporate this sentence” is performed with the operation of *expansion*, denoted $+$. It is a simple set-theoretical operation, defined as follows:

$$K + p = \text{Cn}(K \cup \{p\})$$

Expansion has the virtue of simplicity, but it also has the damaging property of yielding an inconsistent outcome whenever we incorporate some information that contradicts what we believed before. (If $\neg p \in K$ then $K + p$ is inconsistent). Therefore we need the more sophisticated operation of *revision* that corresponds to the instruction “add this sentence and retain consistency”. Revision is denoted $*$ and it is assumed to have the following properties:

$$\begin{array}{ll}
K * p = \text{Cn}(K * p) & \text{(closure)} \\
p \in K * p & \text{(success)} \\
K * p \text{ is consistent if } p \text{ is consistent} & \text{(consistency)}
\end{array}$$

In 1985, Carlos Alchourrón (1931–1996), Peter Gärdenfors, and David Makinson published a paper that became the starting-point of modern research on the logic of belief change [1]. The model they proposed is usually called “AGM” after their initials. When constructing contraction, $K \div p$, they started with the observation that among the many subsets of K not implying p , some are inclusion-maximal, i.e. they are as large as they can be without implying p . These sets are called p -remainders, and the set of p -remainders of K is denoted $K \perp p$.

Intuitively, when contracting K by p we want to keep as much of K as we can while still removing p . This could lead us to take one of the elements of $K \perp p$ as the contraction outcome. However, it may be impossible to single out one of these elements as better than all the others. If several p -remainders share the top position, then our post-contraction beliefs will be those that are held in all the top-ranked p -remainders. Formally, this is achieved by introducing a selection function γ that selects from $K \perp p$ a subset $\gamma(K \perp p)$ consisting of its “best” elements. The outcome of contracting K by p is the intersection of all elements of $\gamma(K \perp p)$, i.e.

$$K \div p = \bigcap \gamma(K \perp p)$$

This construction is called *partial meet contraction*. One way to construct γ is to base it on a transitive relation covering all subsets of K that are

q -remainders for some sentence q . If γ selects the elements of $K \perp p$ that are highest ranked according to such a relation, then the resulting contraction is a *transitively relational partial meet contraction*.

The AGM paper reported axiomatic characterizations of these operations. An operation $\div$ on a belief set K is a partial meet contraction if and only if it satisfies the following six axioms:

$$\begin{array}{ll}
 K \div p = \text{Cn}(K \div p) & \text{(closure)} \\
 K \div p \subseteq K & \text{(inclusion)} \\
 \text{If } p \notin K \text{ then } K \div p = K & \text{(vacuity)} \\
 p \notin K \div p, \text{ unless } p \text{ is a logical truth} & \text{(success)} \\
 \text{If } p \leftrightarrow q \text{ is a logical truth then } K \div p = K \div q & \text{(extensionality)} \\
 K \subseteq (K \div p) + p & \text{(recovery)}
 \end{array}$$

Furthermore, such an operation is transitively relational if and only if it also satisfies the following two axioms:

$$\begin{array}{ll}
 (K \div p) \cap (K \div q) \subseteq K \div (p \& q) & \text{(conjunctive overlap)} \\
 \text{If } p \notin K \div (p \& q) \text{ then } K \div (p \& q) \subseteq K \div p & \text{(conjunctive inclusion)}
 \end{array}$$

The construction of revision in AGM is based on the simple observation that if p cannot be consistently added to K , then that is because $\neg p$ is in K . ($K + p$ is inconsistent if and only if K implies $\neg p$). Therefore, we can make p consistently addable by first removing $\neg p$. This line of reasoning can be found in early work by Isaac Levi [33]. It gives rise to the following construction of revision in terms of contraction and expansion:

$$K * p = (K \div \neg p) + p \quad \text{(the Levi identity)}$$

If revision is defined in this way, then the contraction operation on which the revision $*$ was based can be regained as follows:

$$K \div p = K \cap (K * \neg p) \quad \text{(the Harper identity)}$$

An operation is called a *partial meet revision* if and only if it is obtainable via the Levi identity from some partial meet contraction, and it is a *transitively relational partial meet revision* if and only if it is obtainable in that way from some transitively relational partial meet contraction. The AGM trio showed that partial meet revision is exactly characterized by the following six axioms:

$$\begin{array}{ll}
 K * p = \text{Cn}(K * p) & \text{(closure)} \\
 K * p \subseteq K + p & \text{(inclusion)} \\
 \text{If } \neg p \notin K \text{ then } K + p \subseteq K * p & \text{(vacuity)} \\
 p \in K * p & \text{(success)}
 \end{array}$$

Figure 1. Revision by a sentence p that is compatible with the present belief set

If $p \leftrightarrow q$ is a logical truth then $K * p = K * q$ (extensionality)

If p is consistent then so is $K * p$ (consistency)

In order to characterize transitively relational partial meet revision, the following two axioms have to be added:

$K * (p \& q) \subseteq (K * p) + q$ (superexpansion)

If $\neg q \notin K * p$ then $(K * p) + q \subseteq K * (p \& q)$ (subexpansion)

Alternatively, the AGM model can be expressed in terms of possible worlds, i.e. maximally consistent subsets of the language. A belief set K is compatible with a possible world W if and only if nothing in the belief set contradicts it or, equivalently, if and if $K \subseteq W$. Furthermore, every belief set is equal to the intersection of all possible worlds that include it. (For a proof, see [19, p. 52]). We can therefore replace belief sets by sets of possible worlds in our deliberations. The agent's belief state is then represented by a set of possible worlds (whose intersection is equal to the belief set).

A simple geometrical representation can be used to aid our intuitions [14]. In Figure 1, think of each point in the square as a possible world (and note that a smaller area represents a larger belief set). The circle in the middle contains exactly those possible worlds that are compatible with the current belief state. The area covered by the parabola represents those possible worlds in which p holds. Now consider the revision $K * p$. Its outcome should be represented by a set of possible worlds in which p is true. Since we want to change as little as possible, the obvious solution is to let it consist

Figure 2. Revision by a sentence p that is incompatible with the present belief set

of the worlds included in the intersection of the circle and the parabola, i.e. those of the currently unrejected worlds in which p is true.

In this case the new information was compatible with what was already believed. In cases when this is not so, there are no p -worlds among the K -worlds. We can then resort to the p -worlds that are as close, or similar, to K -worlds as possible. For that purpose we can think of K as surrounded by a system of spheres, with the worlds most similar to it in the sphere closest to K itself, those second-most similar in the next sphere, etc., as in Figure 2. The outcome of revision by p is then equal to the set of p -worlds in the innermost sphere that contains some p -worlds. (Such a system of spheres corresponds, of course, to an ordering of the possible worlds).

Contraction is somewhat less intuitive than revision in possible world models. To contract by p means to allow for the possibility that $\neg p$, i.e. to allow for some possible worlds in which $\neg p$ holds. In a spheres model, these should be the $\neg p$ -worlds that are closest to the belief set, i.e. situated in the innermost sphere that contains some $\neg p$ -worlds. The contraction outcome will then be the union of these $\neg p$ -worlds and the original belief set, as shown in Figure 3.

The possible worlds construction and the partial meet construction yield exactly the same contractions and revisions. In other words, an operation on a belief set K is a transitively relational partial meet contraction if and only if it can be constructed in the way indicated in Figure 3, and it is a transitively relational partial meet revision if and only if it can be constructed

Figure 3. Contraction by p

as shown in Figures 1 and 2. This surprising result is based on a one-to-one correspondence called “Grove’s bijection” between the remainder set $K \perp p$ and the set of possible worlds not containing p [14], [19, pp. 53–55]. (Two crucial facts show why this can be so. First, if K is a belief set, $p \in K$, and $X \in K \perp p$, then $\text{Cn}(X \cup \{\neg p\})$ is a maximal consistent subset of the language.¹ Secondly, if K is a belief set, $p \in K$, and Y is a maximal consistent subset of the language not containing p , then $Y \cap K \in K \perp p$.²)

Hence, we obtain exactly the same results from applying a selection function to remainders of a belief set as from applying such a function to possible worlds. It is a strength of the AGM model that it can be characterized in multiple ways, of which these two are arguably the most important ones. But obviously, this does not make the model immune against criticism.

¹Proof: Suppose not. Then there is some z such that $z \notin \text{Cn}(X \cup \{\neg p\})$ and $\neg z \notin \text{Cn}(X \cup \{\neg p\})$. It follows from $z \notin \text{Cn}(X \cup \{\neg p\})$, i.e. $X \cup \{\neg p\} \not\models z$, that $X \not\models \neg z \rightarrow p$. Since $\neg z \rightarrow p \in K$ it then follows from $X \in K \perp p$ that $X \cup \{\neg z \rightarrow p\} \models p$, thus $X \models (\neg z \rightarrow p) \rightarrow p$, thus $X \models \neg p \rightarrow \neg z$, contrary to $\neg z \notin \text{Cn}(X \cup \{\neg p\})$.

²Proof: Suppose to the contrary that there is some $z \in K$ such that $z \notin Y \cap K$ and $(Y \cap K) \cup \{z\} \not\models p$. Then $z \notin Y$, and since Y is a maximal consistent subset of the language we have $\neg z \in Y$, thus $\neg z \vee p \in Y$. Since $p \in K$ we also have $\neg z \vee p \in K$, thus $\neg z \vee p \in Y \cap K$, thus $(Y \cap K) \cup \{z\} \models p$, contrary to the assumption.

3. The Select-and-Intersect Method

As should be clear from the previous section, belief change theory relies heavily on set-theoretical intersection. Both partial meet contraction and the spheres model employ what can be described as a two-step *select-and-intersect* method. In the first step of these operations several belief sets (logically closed sets) are selected, all of which satisfy the success criterion of the operation. In contraction by a non-tautologous sentence p , the success condition is not to imply p , and each element of $K \perp p$ satisfies that condition. In revision by a sentence p , the success condition is to imply p , and each possible world selected as in Figures 1 or 2 satisfies that condition. In the second step the intersection of those sets is formed, and it is taken to be the outcome of the operation.

At first glance, the select-and-intersect method may seem to be an almost impeccable way to deal with ties. When we hesitate between two or more potential outcomes, then it would seem natural to use their intersection, i.e. the part that they have in common, as the output. But closer inspection will reveal that the select-and-intersect method can be questioned on at least three accounts:

- (i) The *preservation of optimality under intersection*: In the first step of the select-and-intersect process, options are chosen that are in some sense optimal. Is that optimality retained after intersection?
- (ii) The *preservation of success under intersection*: The options chosen in the first step all satisfy the success condition of the operation. Is the success condition always satisfied by their intersection?
- (iii) The *adequacy of the options selected for intersection*: Intersection can be justified as a way to adjudicate between equally plausible outcomes. But do the options selected in the first stage of the process at all belong to the plausible outcomes? If not, can the select-and-intersect method yet be justified?

The following three sections are devoted to these three questions.³

³The first two of these issues are also relevant for models employing belief bases, i.e. sets of sentences not closed under logical consequence, as representations of the belief state. The third issue is less relevant for belief bases since maxichoice contraction (that yields a remainder as outcome) is much more plausible for belief bases than for belief sets [19, p. 77, 36].

4. Is Optimality Preserved Under Intersection?

The very idea of choosing or selecting among potential outcomes of belief change is epistemologically somewhat problematic. Many, arguably most, belief changes seem to be uncontrollable effects of external influences rather than the outcomes of voluntary choices made by the subject [24, pp. 143–145]. Therefore, the use of selection mechanisms in belief change cannot credibly be justified by claims that our changes in beliefs are actually the outcomes of choices among alternative ways to accommodate a new belief pattern. However, their use can be justified by a move that has been made with some success in decision theory: It may be claimed that when you are induced (by some input) to perform a belief change that could potentially end in several alternative outcomes, then that operation will conclude *as if* it was performed by some process that selected the most suitable among the potential outcomes.

Arguably, the as-if account provides us with a reasonable justification for treating belief changes as selection processes (At any rate, no better justification seems to be available). However, it has the disadvantage of leading to a lack of intuition-guiding examples from the intended area of application. Since we do not normally make voluntary choices among alternative belief options (and in particular not among remainders or possible worlds), we cannot use practical examples of such choices to determine the appropriate properties of the selection functions of belief change models. Instead, discussions on suitable formal properties for selection functions in belief change have drawn heavily on intuitions about choices among other types of objects such as physical objects or social states of affairs. Rott [39] has shown in considerable detail how various properties of social choice functions, when applied to the selection function γ for remainders, correspond to interesting properties of the partial meet contraction based on that selection function. However, in spite of these elegant formal results it is not self-evident that choices among epistemic options should follow the same patterns as choices among social states. A potential reason why their properties might have to be different is that the selected set is used quite differently in the two areas. In social choice theory, if a social choice function leaves as outcome a set with more than one element, then the choice is considered to be indeterminate; the agent may choose any of the selected options. As we have already seen,

in belief change such an outcome is instead made deterministic by taking the intersection of the selected options to be the real outcome.⁴

It is easy to show that the select-and-intersect method would lead to absurd results if applied to social choice. For a simple example, consider the following offers, each of which contains two items, a vacation trip and a sum of money:

$A = \{\text{A week for two in Florence, } \text{€ } 100\}$

$B = \{\text{A week for two in Venice, } \text{€ } 100\}$

$C = \{\text{A week for two in Barcelona, } \text{€ } 200\}$

$D = \{\text{A week for two in Madrid, } \text{€ } 200\}$

Suppose that your choice set is $\{A, B\}$, i.e. you prefer each of these two alternatives to the others but you do not prefer A to B or the other way around. According to the select-and-intersect method, which we now (ungraciously) transfer from belief choice to social choice, the tie is solved by allotting to you $A \cap B$, i.e. € 100 and no travel. This is absurd, since you would probably prefer each of C , D and (most certainly) $C \cap D$ to $A \cap B$.

This would be no problem for belief change theory if it could be shown that contrary to other collections of objects, belief sets do not lose in choice-worthiness by being intersected with other equally choiceworthy objects. However, no such argument seems to be available. To the contrary, Tor Sandqvist has shown with an ingenious example that if each element of the collection $\mathcal{A}$ of belief sets is preferable to each element of the collection $\mathcal{B}$ of belief sets, it does not follow that the belief set $\cap \mathcal{A}$ is preferable to the belief set $\cap \mathcal{B}$. He concluded that “there may very well be two or more sets of beliefs that are each very valuable but such that their intersection is practically worthless—namely, if whatever makes each of them so valuable fails to be that which they all have in common” [43, p. 292].

5. Are all Success Conditions Preserved Under Intersection?

A success condition is a property of the outcome of a belief change, and furthermore it is the property that the operation aims at achieving. The success condition of contraction by p is that the outcome $K \div p$ should not

⁴A major reason for the difference is that in social choice the selection function operates on a set of potential outcomes whereas in belief change it usually operates on a set whose elements are not themselves plausible outcomes. See Section 6.

contain p , i.e. $p \notin K \div p$. The success condition of revision by p is that the outcome $K * p$ should contain p , i.e. $p \in K * p$. Both these properties are preserved under intersection, in the following sense:

DEFINITION 1. [28] A property of belief sets is *preserved under intersection* if and only if: If all elements of a set $\mathfrak{X}$ of belief sets satisfy the property, then so does $\bigcap \mathfrak{X}$.

Operations of belief change may have other success conditions than those of AGM contraction or revision, for instance:

Replacement has the success condition that a sentence p is an element of the outcome and another sentence q is absent from it [23].

Package contraction has the success condition that all elements of a set A of sentences are absent from the outcome [7].

Choice revision has the success condition that at least one element of a set A of sentences is an element of the outcome [5, pp. 160–161], [7].

It is easy to show that the success conditions of replacement and package contraction are preserved under intersection whereas that of choice revision is not.⁵ The select-and-intersect method has the limitation of only being workable for operations whose success conditions are preserved under intersection (See the “Appendix” for a formal result covering those success conditions.)

6. How Plausible are the Objects of Selection?

As mentioned in Section 3, if the first step of the select-and-intersect method has resulted in a set consisting of the most plausible outcome candidates, then the intersection performed in the second step yields a belief set consisting of exactly that which all the most plausible options have in common. This can be a reasonable way to resolve the deadlock, not least since this method is symmetrical in the sense of treating the options we cannot choose between equally. However, this justification of the select-and-intersect method does not work if the objects selected in the first step are not themselves plausible outcomes. In this section we are going to consider the plausibility issue for the objects used in this way in the operations discussed in Section 2, namely

⁵For the latter, let p and q be logically independent sentences, and let $A = \{p, q\}$. Then each element of the set $\{\text{Cn}(\{p\}), \text{Cn}(\{q\})\}$ satisfies the condition of containing at least one element of A , whereas their intersection $\text{Cn}(\{p\}) \cap \text{Cn}(\{q\}) = \text{Cn}(\{p \vee q\})$ does not.

possible worlds and remainders of a belief set. We will consider them from two points of view: groundedness and finitude.

6.1. Groundedness

The belief set resulting from a belief change should be based on the information contained in the previous belief set and the input. Therefore a plausible outcome candidate should not contain information with no ground in either of these. We can call this (informally defined) property *groundedness*.

From this point of view, possible worlds are not plausible outcome candidates. If W is a possible world, then it holds for each sentence q in the language that either $q \in W$ or $\neg q \in W$. The implausibility of outcomes with this property was noted by two of the AGM authors already in 1982 [2, pp. 19–21]. An example can serve to illustrate the point: On one occasion I had a belief set K containing the sentence “There is milk in my fridge” (p). When opening my fridge I found this to be wrong and revised my belief set by $\neg p$. The sphere model (as in Figure 2) depicts this change as a process in which I first selected the most plausible possible worlds in which $\neg p$ is true, and then adopted the intersection of all those worlds as my new belief set. In none of the options selected in the first stage would I suspend my belief on any topic—I would be a *Besserwisser* with a confident answer to every possible question.

Remainders of belief sets satisfy a related, weaker property:

OBSERVATION 1. [2, p. 20] *Let $\div$ be a maxichoice contraction on the belief set K . Let $p \in K$ and let q be any proposition. Then either $p \vee q \in K \div p$ or $p \vee \neg q \in K \div p$.*

PROOF. See Alchourrón and Makinson [2] or Hansson [19, p. 124]. ■

As Alchourrón and Makinson noted, this is a “rather counterintuitive” property, in particular when q is (intuitively speaking) content-wise unconnected with both p and the rest of K . (To see that, again let p denote that there is milk in my fridge and let q be a statement about something that I know nothing about).

In conclusion, both possible worlds and remainders of the belief set are inadequate outcome candidates since they add information that is not justified by the original belief set or by the input that triggers the operation.

6.2. Finitude

Since our brains (and minds) are finite, we should expect them only to have room for a finite number of beliefs. Therefore our formal representations of

belief states should satisfy some form of finiteness criterion. The specification of that requirement calls for some sophistication. As a first attempt, one might require the object language of belief revision to be logically finite, i.e. have a finite number of logically non-equivalent elements.⁶ However, such a language is bound to have gratuitous limits on its expressive power [21, 22]. Consider the following list of sentences:

Less than 50 paintings by Johannes Vermeer are extant.
 Less than 51 paintings by Johannes Vermeer are extant.
 Less than 52 paintings by Johannes Vermeer are extant.
 ...
 Less than 1.000.000 paintings by Johannes Vermeer are extant.
 ...

I believe in each of the sentences on that list, and therefore my set of beliefs contains infinitely many logically non-equivalent sentences. But alas, this does not make my set of beliefs infinite in any interesting way. All of these sentences follow from the first, and the others are logically superfluous.

Guided by examples like this we should reformulate our requirement. A reasonably realistic model of belief change should have a logically infinite language but logically finite belief sets. The following terminology will be useful:

DEFINITION 2. (1) A belief set X is *finite-based* if and only if there is some set X' such that $X = \text{Cn}(X')$.

(2) [17, p. 604] An operation on belief sets satisfies *finite-based outcome* if and only if it yields a finite-based belief set as outcome whenever the original belief set is finite-based.

Belief revision consists in the incorporation of a single sentence into the belief set. Since a single sentence only contains a finite amount of information, it does not contain enough information to take us from a finite-based belief set to one that is logically infinite. Similarly, contraction means loss of information and therefore, when we contract a finite-based belief set by some sentence, the outcome should be finite-based. In other words, if the original belief set is finite-based, then a plausible outcome candidate for revision or contraction will also have to be finite-based.

⁶A logically finite language can be syntactically infinite, for instance since it contains the infinite series $p, p \vee p, p \vee p \vee p, \dots$ of logically equivalent sentences. Contrary to logical finiteness, syntactic finiteness is a property of the language itself (rather than a property of the logic).

In a logically infinite language, possible worlds are logically infinite. As we have just seen, this makes them implausible as outcome candidates. This is problematic for the spheres model that employs possible worlds as objects of selection in the first step of the select-and-intersect process. The following observation shows that partial meet contraction does not fare better in this respect:

OBSERVATION 2. *Let the language consist of infinitely many logically independent atoms and their truth-functional combinations. Let K be a belief set that contains some non-tautology, and let $p \in K \setminus \text{Cn}(\emptyset)$. Then:*

- (1) $K \perp p$ is infinite, and
- (2) if $X \in K \perp p$, then X is not finite-based.

PROOF. See Hansson [22, p. 33 and pp. 43–44]. ■

Thus, just like possible worlds, remainders are implausible outcome candidates since they invariably take us from a finite-based belief set to one that is logically infinite. The introduction of such intermediate infinite objects into operations on finite-based sets does not seem to be compatible with Ockhamist requirements of simplicity.

7. Properties of the Operations

In Sections 4, 5 and 6 we identified several credibility problems for the select-and-intersect process employed in standard belief revision theory. Such criticism can be countered with “black box” arguments; in other words it can be argued that it does not matter much if the process is implausible, if only the resulting operations have the right properties. In this section it will be shown that some of the controversial properties of standard belief revision models depend on the properties of the underlying process that were unfolded in the previous sections.

7.1. Contraction Properties

It follows directly from Observation 2 that partial meet contraction does not in general satisfy finite-based outcome, i.e. it does not ensure that the contraction outcome is finite-based if the original belief set is finite-based. Interestingly, although finite-based outcome is not one of the AGM postulates (and does not follow from them), all three AGM authors have gone on record endorsing what that postulate requires.

“We suggest, finally, that the intuitive processes themselves, contrary to casual impressions, are never really applied to theories as a whole, but rather to more or less clearly identified bases for them. For a theory is an infinite object, having as it does an infinite number of elements, and it is only by working on some finite generator or representative of the theory that the outcome of a process such as contraction can ever in practice be determined” [2, pp. 21–22].

“In all applications, knowledge sets [belief sets] will be *finite* in the sense that the consequence relation $\vdash$ partitions the elements of K into a finite number of equivalence classes.” [12, p. 90]

Obviously, finite-based outcome can be added as a postulate, but on the construction side rather radical changes are needed to achieve it.⁷

The most discussed property of partial meet contraction is one of the basic postulates mentioned in Section 2:

$$K \subseteq \text{Cn}((K \div p) \cup \{p\}) \quad (\text{recovery})$$

that is satisfied by all AGM contractions. The following example has been offered to show why it is problematic:

I believed that *Cleopatra had a son* (s). Therefore I also believed that *Cleopatra had a child* (c or equivalently $s \vee d$ where d denotes *Cleopatra had a daughter*). Then I received information that made me give up my belief in c , and contract my belief set accordingly, forming $K \div c$. Soon afterwards I learned from a reliable source that Cleopatra had a child. It seems perfectly reasonable for me to then add c (i.e. $s \vee d$) to my set of beliefs without also reintroducing s . [15]

Recovery has been subject to extensive discussions [9, 13, 15, 20, 36, 37]. For our present purposes it is important to note that recovery is unavoidable if $K \div p$ is the intersection of some elements of $K \perp p$. (For a proof, see [1, p. 513] or [19, p. 123].) It is in other words a direct consequence of applying the select-and-intersect method to remainders.

⁷As far as I know, there is only one construction on record by which finite-based outcome can be achieved while retaining all the AGM axioms. This construction replaces the selection function γ by a sentential selector f that takes us from one sentence in the language to another, so that $K \div p$ is identified with $\bigcap (K \perp f(p))$ (where $f(p)$ is a sentence) rather than with $\bigcap (\gamma(K \perp p))$. See [22].

7.2. Revision Postulates

It follows from the construction shown in Figure 2 that sphere-based revision does not in general satisfy finite-based outcome. The same holds of course for the equivalent operation of partial meet revision.⁸ Just as for contraction, non-satisfaction of finite-based outcome seems to be a rather persistent property that can only be removed with far-reaching changes of the framework.

Furthermore, AGM revision satisfies the following property:

If $K \neq \neg p$ then $K * p = \text{Cn}(K \cup \{p\})$ (expansion property of revision)

The expansion property of revision is just as tenacious as the recovery property of contraction. In the spheres model, it follows from the assumption that the possible worlds that are compatible with the present belief set form the innermost sphere (as in Figure 1). In the equivalent formulation of partial meet revision it follows from the Levi identity in combination with the vacuity postulate.⁹ However, examples of revision are easily found in which the expansion property is implausible.

John is a neighbour about whom I initially know next to nothing.

Case 1: I am told that he goes home from work by taxi every day (t).

This makes me believe that he is a rich man (r).

Case 2: When told t , I am also told that John is a driver by profession

(d). In this case I am not made to believe that he is a rich man (r)

[16].

In case 1 we have $r \in K * t$, and due to the expansion property $K * t = \text{Cn}(K \cup \{t\})$. Since K is logically closed it follows that $t \rightarrow r \in K$. In case 2, the expansion property yields $K * (t \& d) = \text{Cn}(K \cup \{t \& d\})$. Combining this with $t \rightarrow r \in K$ we obtain $r \in K * (t \& d)$, contrary to the description of case 2.¹⁰

This example shows that the expansion property is at variance with a common (and arguably fully rational) pattern of belief change: When we

⁸For the proof it is sufficient to consider the limiting case in Figure 2 in which only one p -containing world is selected. A proof referring to the partial meet construction can be found in [32].

⁹According to the Levi identity, $K * p = \text{Cn}((K \div \neg p) \cup \{p\})$. Since $K \neq \neg p$, vacuity yields $K \div \neg p = K$.

¹⁰For a discussion of the underlying problem in this example, see [24]. See also [11, pp. 67–68] for a discussion of the limited expressivity of belief sets; in particular they cannot express that some beliefs are reasons for other beliefs.

acquire a new belief that does not contradict our previous beliefs (such as t in the example), we often complement the outcome with some additional belief (such as r in the example) that “rounds off” the belief set and makes it more coherent, but does not follow deductively.¹¹ This seems to be connected with a form of coarse-grainedness of the outcome set (the set of possible outcomes of belief change): Only some of the belief sets that can be formed by pure set-theoretical addition to the original belief set are coherent enough to be viable revision outcomes [25].

It should also be noted that the impossibility of including Ramsey test conditionals in the AGM model [10] is closely connected with the expansion property of revision [16, 38]. This can be seen from the above example. Let $\Rightarrow$ satisfy the Ramsey condition, i.e. for all belief sets K' : $p \Rightarrow q \in K'$ if and only if $q \in K' * p$. Then (one direction of) the Ramsey condition yields $t \Rightarrow r \in K$, and due to $K * (t \& d) = \text{Cn}(K \cup \{t \& d\})$ that follows from the expansion property we can draw the implausible conclusion that $t \Rightarrow r \in K * (t \& d)$.

The expansion property also goes wrong in the opposite direction: Sometimes when we add a sentence p that does not logically contradict the belief set K , it leads to the exclusion of some sentence that the new information makes implausible but for non-deductive reasons. In such cases we can have $K \nVdash \neg p$ but $\text{Cn}(K \cup \{p\}) \notin K * p$.

I believed that one of the three heirs, Amelia, Barbara, and Carol, murdered the rich eccentric ($a \vee b \vee c$). Then I received information convincing me that both Amelia and Barbara are innocent ($\neg a \& \neg b$). However, since I had no specific information binding Carol to the crime this did not make me believe that Carol was the murderer.

Valentina was uncertain whether or not her husband is faithful to her (f), but she still believed that her husband loves her (l). However, when she learnt that he is unfaithful to her, she lost her belief that he loves her.

In the first example we have $K \nVdash \neg(\neg a \& \neg b)$ but it can be seen from $c \notin K * (\neg a \& \neg b)$ that $\text{Cn}(K \cup \{\neg a \& \neg b\}) \notin K * (\neg a \& \neg b)$. In the second example we similarly have $K \nVdash f$ but $l \in \text{Cn}(K \cup \{\neg f\})$ and $l \notin K * \neg f$, thus $\text{Cn}(K \cup \{\neg f\}) \notin K * \neg f$.

In summary, the expansion property of revision is about as problematic as the recovery property of contraction. Since both properties are inseparable

¹¹Standard inductive inference is a clear example of this. The generalization from a thousand observations of black ravens (none of which contradicts previous beliefs) to the belief that all ravens are black is a case in point.

from the standard framework of belief revision, this adds to our justifications for investigating possible alternatives.

8. Back to Basics: Input Assimilation

Sections 3, 4, 5, 6 and 7 have assembled a series of problems with the select-and-intersect method in belief change. In this and the following section an alternative approach will be developed in which belief change still makes use of a selection function, but that function operates directly on the set of potential outcomes and selects exactly one of them. As a first step in developing that method, we will temporarily relinquish an idealization that has been taken for granted since it was mentioned in Section 2, namely the sentential representation of belief states and inputs.

Obviously, a person's state of belief consist of much more than sentences or that which can be expressed with sentences. You may have quite definite beliefs (or knowledge) about how Beethoven's Pastoral symphony sounds, how hydrogen sulphide smells, or what Picasso's *Guernica* looks like, without being able to convey more than rudimentary fragments of those beliefs in linguistic form. Our states of belief are largely non-sentential, and therefore a fully general model of belief change will have to operate with a set of *primitive belief states*.¹² We will assume that there is a set $\mathbb{K}$ of possible belief states and a set $\mathbb{I}$ of inputs. To express the changes brought about by the inputs, we will use a universal operation $\odot$ of *input assimilation*. For each $\mathcal{K} \in \mathbb{K}$ and $i \in \mathbb{I}$, $\mathcal{K} \odot i$ is the outcome of subjecting $\mathcal{K}$ to the input i . The outcome of this operation should be a new belief state, in other words:

$$\mathcal{K} \odot i \in \mathbb{K} \quad (\text{IA-1})$$

Just like the standard operations of belief change, input assimilation is deterministic, i.e. it always specifies exactly one outcome.¹³ The following property can be seen as an inversion of IA-1:

$$\text{For all } \mathcal{K}, \mathcal{K}' \in \mathbb{K} \text{ there is some } i \in \mathbb{I} \text{ such that } \mathcal{K}' = \mathcal{K} \odot i \quad (\text{IA-2})$$

¹²Since the belief state is a part or (rather) an aspect of the agent's overall state of mind we can justifiably go one step further and operate with states of mind rather than belief states. This generalization will be important in investigations of the relationships between the beliefs supported by a state of mind and other linguistic entities, such as value statements, that are supported by the same state of mind. However, for our present purposes, it is sufficient to consider belief states.

¹³A few studies in belief revision have been devoted to indeterministic operations that for some inputs specify a set of more than one possible outcomes, without telling us which of these will eventuate. See [8, 35].

IA-2 is much less plausible than IA-1, since there may be potential states of belief that the agent can only arrive at after receiving several inputs. For instance, if $\mathcal{K}$ is a belief state in which the agent is a devout religious believer and $\mathcal{K}'$ one in which she is a staunch atheist, then there may be no single input that would take her from $\mathcal{K}$ to $\mathcal{K}'$. It is much more plausible that a series of inputs could take her there through a mechanism whereby the earlier of these inputs facilitate her assimilation of those coming later. Thus, the following condition could be satisfied.

For all $\mathcal{K}, \mathcal{K}' \in \mathbb{K}$ there are some $i_1, \dots, i_n \in \mathbb{I}$ such that $\mathcal{K}' = \mathcal{K} \odot i_1 \odot \dots \odot i_n$
(IA-3)

IA-3 is a plausible postulate if $\mathbb{K}$ is interpreted as the set of belief states that are currently possible for the agent to arrive at in some way. However, the potential to arrive at a belief state can be lost as the agent assimilates inputs, hence $\mathcal{K}'$ can be reachable from $\mathcal{K}$ but not from $\mathcal{K} \odot i$. Therefore, the interpretation of $\mathbb{K}$ as the set of currently reachable belief states requires $\mathbb{K}$ to be revised each time the belief state is changed. From the viewpoint of formal convenience this should be avoided. Instead, $\mathbb{K}$ can be treated as an unchanging part of the framework, and interpreted as containing not only the belief states that the agent can reach from her present starting-point but also those that she could have reached at some earlier point.¹⁴

9. Reintroducing Sentences

The framework of general input assimilation introduced in the previous section has the advantage of discarding in one fell swoop all the assumptions about relations between sentence structure and operations of change that were shown in the previous sections to give rise to difficulties. However, we may have thrown out too much. Although belief states do not consist of sentences, they are closely associated with beliefs in sentences. Although many belief changes are not adequately described as the incorporation or removal of a sentence, many other belief changes are reasonably well described in that way. Furthermore, chances seem slim of building an interesting model of belief change within the impoverished framework of general input assimilation. Perhaps we can reintroduce sentences in a more

¹⁴An accessibility relation R in the style of modal logic can be formed according to the formula: $\mathcal{K}R\mathcal{K}'$ iff $(\exists i \in \mathbb{I})(\mathcal{K} \odot i = \mathcal{K}')$. Cf. [18].

cautious manner, avoiding some of the more controversial assumptions of the traditional approach.

9.1. Support Functions

Although your current belief state contains non-sentential elements and has features that you cannot fully describe with sentences, it also has implications for what sentences you believe in. We can describe this property of belief states with a *support function* $\mathfrak{s}_{\mathcal{L}}$ that takes us from elements of $\mathbb{K}$ to sets of sentences in the object language $\mathcal{L}$.¹⁵ [16, p. 525] Thus, $\mathfrak{s}_{\mathcal{L}}(\mathcal{K})$ is the set of sentences in $\mathcal{L}$ that are supported (believed by the epistemic agent) in the belief state $\mathcal{K}$. In what follows the index of $\mathfrak{s}_{\mathcal{L}}$ will be omitted.

The following property:

$$\mathfrak{s}(\mathcal{K}) = \text{Cn}(\mathfrak{s}(\mathcal{K})) \quad (\text{closure})$$

is a highly simplifying feature of the formal framework. It can be justified by interpreting $\mathfrak{s}(\mathcal{K})$ as the set of sentences that the epistemic agent is (logically) committed to believe in, rather than those that she actually believes in [33, 34]. Closure is a reasonable assumption in studies of rational belief change, and it will be made in what follows.

The following are some other properties that may be worth considering:

$$\perp \notin \mathfrak{s}(\mathcal{K}) \quad (\text{consistency})$$

$$\mathfrak{s}(\mathcal{K}) \text{ is finite-based} \quad (\text{finite representability})$$

$$\text{If } \mathfrak{s}(\mathcal{K}) \text{ is finite-based then so is } \mathfrak{s}(\mathcal{K} \odot i). \quad (\text{finite-based outcome})$$

$$\text{If } p \not\vdash \perp \text{ then there is some } i \text{ with } p \in \mathfrak{s}(\mathcal{K} \odot i). \quad (\text{believability})$$

$$\text{If } p \notin \text{Cn}(\emptyset) \text{ then there is some } i \text{ with } p \notin \mathfrak{s}(\mathcal{K} \odot i). \quad (\text{removability, [25]})$$

$$\text{There is some } i \text{ with } \mathfrak{s}(\mathcal{K} \odot i) = \text{Cn}(\emptyset) \quad (\text{depletability, [25]})$$

$$\text{If } \mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}') \text{ then } \mathcal{K} = \mathcal{K}' \quad (\text{sentential uniqueness})$$

$$\text{If } \mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}') \text{ then } \mathfrak{s}(\mathcal{K} \odot i) = \mathfrak{s}(\mathcal{K}' \odot i) \text{ for all } i \in \mathbb{I} \quad (\text{local sententiality})$$

$$\text{If } \mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}') \text{ then } \mathfrak{s}(\mathcal{K} \odot i_1 \odot \cdots \odot i_n) = \mathfrak{s}(\mathcal{K}' \odot i_1 \odot \cdots \odot i_n) \text{ for all series } i_1, \dots, i_n \text{ of elements of } \mathbb{I} \quad (\text{global sententiality})$$

The last three of these properties are particularly interesting for the reconstruction of a sentential framework. Sentential uniqueness is the strongest

¹⁵The construction of $\mathcal{L}$ will be left open. It was shown in [30] that the well-known restrictions against including conditional, modal, and autoepistemic sentences in the language of AGM theory do not apply here.

of these; it says that differences between belief states always have sentential implications. If it holds, then the difference between non-identical belief states is always manifested on the sentential level. Global sententiality is weaker but still implies that if two belief states are sententially indistinguishable, then after any series of changes the outcomes will continue to be sententially indistinguishable. Local sententiality only guarantees this for single changes; it allows divergences to arise in iterated change.

9.2. Belief Descriptors

In order to reconstruct a sentential framework we need to represent not only belief states but also inputs in sentential terms. We will do this in a more general fashion than the traditional one. As noted above, different types of belief change operations are characterized by their success conditions. Since these conditions all refer to what the agent believes, they can be expressed in a uniform way with the help of a metalinguistic belief operator $\mathfrak{B}$. This will be done as follows:

DEFINITION 3. [27] An *atomic belief descriptor* is a sentence $\mathfrak{B}p$ with $p \in \mathcal{L}$. It is *satisfied* by a belief set K if and only if $p \in K$.

A *molecular belief descriptor* (denoted by lower-case Greek letters $\alpha, \beta \dots$) is a truth-functional combination of atomic descriptors. Conditions of satisfaction are defined inductively, hence K satisfies $\neg\alpha$ if and only if it does not satisfy α , it satisfies $\alpha \vee \beta$ if and only if it satisfies either α or β , etc.

A *composite belief descriptor* (in short: descriptor; denoted by upper-case Greek letters $\Psi, \Xi \dots$) is a set of molecular descriptors. A belief set K satisfies a composite descriptor Ψ if and only if it satisfies all its elements.

With this notation we can express not only the success conditions of the standard AGM operations but also those of a much wider range of operations:

$\mathfrak{B}p$	Revision by p
$\neg\mathfrak{B}p$	Revocation (“contraction”) by p
$\mathfrak{B}p_1 \vee \dots \vee \mathfrak{B}p_n$	Choice revision by $\{p_1, \dots, p_n\}$
$\{\neg\mathfrak{B}p, \mathfrak{B}q\}$	Replacement of p by q
$\mathfrak{B}p \vee \mathfrak{B}\neg p$	Making up one’s mind about p [44]

Moreover, a unified descriptor-based operation $\circ$ of belief change can be introduced, such that for any belief state $\mathcal{K}$ and any descriptor Ψ , if revision by Ψ is successful, then $\mathcal{K} \circ \Psi$ is one of those elements of $\mathbb{K}$ that are reachable from $\mathcal{K}$ and satisfy Ψ . In formal language:

$$\mathcal{K} \circ \Psi \in \{\mathcal{K}' \in \mathbb{K} \mid (\exists i \in \mathbb{I})(\mathcal{K}' = \mathcal{K} \odot i) \text{ and } \mathfrak{s}(\mathcal{K}') \text{ satisfies } \Psi\}$$

or equivalently:

$$\mathcal{K} \circ \Psi \in \{\mathcal{K} \odot i \mid i \in \mathbb{I} \text{ and } \mathfrak{s}(\mathcal{K} \odot i) \text{ satisfies } \Psi\}$$

We can assume that there are many belief change outcomes $\mathcal{K} \odot i$ that satisfy Ψ . For instance, if Ψ represents the belief that the old vase in my family's living-room is broken, then Ψ is satisfied in a large number of potential belief change outcomes, including far-fetched ones with various additional beliefs such as that a wild bird flew in through an open window and knocked down the vase, etc. The operation of revision by Ψ should not result in one of these far-fetched outcomes but rather in a “minimally changed” belief state that is, intuitively speaking, as close or similar to my previous belief state as is compatible with the assimilation of Ψ . We can see this as an Ockhamist simplicity requirement: in revision by Ψ no unnecessary additional information should be added to that which is contained in Ψ . The further specification of this simplicity requirement will be left open. The crucial requirement is that among the various potential outcomes satisfying Ψ , there is one that is privileged in the sense of being singled out as the outcome assigned to revision *by* Ψ . In the formal language, this singling out is most conveniently represented by a selection function that is applied to the outcomes supporting Ψ , and selects exactly one of them. More precisely:

DEFINITION 4. [27] Let K be a belief set. $K \models \Psi$ means that K satisfies Ψ and $\Psi \models \Xi$ that all belief sets satisfying Ψ also satisfy Ξ . The corresponding equivalence relation is written $\models$.

DEFINITION 5. [26] A *monoselective choice function* for a set $\mathbb{X}$ is a function s on $\wp(\mathbb{X}) \setminus \{\emptyset\}$ such that if $\emptyset \neq \mathbb{Y} \subseteq \mathbb{X}$ then $s(\mathbb{Y}) \in \mathbb{Y}$.

DEFINITION 6. Let $\mathbb{K}$ be a set of belief states, $\mathbb{I}$ a set of inputs, $\odot$ an input assimilation operation on $\mathbb{K}$ and $\mathbb{I}$, $\mathfrak{s}_{\mathcal{L}}$ a support function for $\mathbb{K}$ and a language $\mathcal{L}$, and s a monoselective choice function for $\mathbb{K}$. The *descriptor revision* based on $\langle \mathbb{K}, \mathbb{I}, \odot, \mathfrak{s}_{\mathcal{L}}, s \rangle$ is the operation $\circ$ such that for all $\mathcal{K} \in \mathbb{K}$ and all descriptors Ψ for the language $\mathcal{L}$:

- (i) If $\{\mathcal{K} \odot i \mid i \in \mathbb{I} \text{ and } \mathfrak{s}(\mathcal{K} \odot i) \models \Psi\} \neq \emptyset$, then
 $\mathcal{K} \circ \Psi = s(\{\mathcal{K} \odot i \mid i \in \mathbb{I} \text{ and } \mathfrak{s}(\mathcal{K} \odot i) \models \Psi\})$, and
- (ii) otherwise $\mathcal{K} \circ \Psi = \mathcal{K}$

Descriptor revision was introduced in [27] in a framework of belief sets, thus without mentioning an underlying operation of general input assimilation.

9.3. Combining Support Functions and Descriptors

If sentential uniqueness holds, then we can construct descriptor revision as based on the set $\{s(\mathcal{K}) \mid \mathcal{K} \in \mathbb{K}\}$ of belief sets instead of the set $\mathbb{K}$ of primitive belief states. The following observation shows that descriptor revision can then completely replace general input assimilation:

OBSERVATION 3. *Let $\circ$ be the descriptor revision based on $\langle \mathbb{K}, \mathbb{I}, \odot, \mathfrak{s}_{\mathcal{L}}, s \rangle$. If sentential uniqueness holds, then for each $\mathcal{K} \in \mathbb{K}$ and $i \in \mathbb{I}$ there is a descriptor Ψ with $\mathcal{K} \odot i = \mathcal{K} \circ \Psi$.*

PROOF. For each $i \in \mathbb{I}$, use the descriptor $\{\mathfrak{B}p \mid p \in \mathfrak{s}(\mathcal{K} \odot i)\} \cup \{\neg \mathfrak{B}p \mid p \notin \mathfrak{s}(\mathcal{K} \odot i)\}$. ■

The selection function of Definition 6 can be based on an ordering or a distance measure. We can think of the potential outcomes as dispersed in some kind of metric space. $K \circ \Psi$ can then be identified as the potential outcome satisfying Ψ that is closest to K . It was shown in [27] that such *relational descriptor revision* is axiomatically characterized by the following properties:

- If $\Psi \dashv\vdash \Psi'$ then $K \circ \Psi = K \circ \Psi'$ (extensionality)
- $K \circ \Psi = \text{Cn}(K \circ \Psi)$ (closure)
- If $K \Vdash \Psi$ then $K \circ \Psi = K$ (confirmation)
- $K \circ \Psi \Vdash \Psi$ or $K \circ \Psi = K$ (relative success)
- If $K \circ \Xi \Vdash \Psi$ then $K \circ \Psi \Vdash \Psi$ (regularity)
- If $K \circ \Psi \Vdash \Xi$ then $K \circ \Psi = K \circ (\Psi \cup \Xi)$ (cumulativity)

In the AGM framework, an important alternative formalization of transitively relational partial meet contraction is based on a binary relation on sentences, called epistemic entrenchment [11, 12, 39, 40].¹⁶ A sentence p is said to be less entrenched than q if it is more easily given up. With the descriptor terminology we can also express this by saying that the agent is more inclined to adopt a belief state satisfying $\{\neg \mathfrak{B}p\}$ than one that satisfies $\{\neg \mathfrak{B}q\}$. This can be generalized to a relation $\leq$ on descriptors in general, with $\Psi \leq \Xi$ interpreted as saying that the agent is at least as inclined to have a belief state satisfying Ψ as one satisfying Ξ . In [28] this relation was shown to be interchangeable with the distance-based construction just referred to.

¹⁶Epistemic entrenchment is closely related to the spheres model [19, pp. 227–228 and 300–304], and also to another construction based on a binary relation on sentences, safe contraction [3, 42].

In [30] iterated descriptor revision was introduced. Its introduction requires that distances are appropriately defined for all pairs of potential outcomes, not only for such pairs in which the original belief set is one of the elements. Hence, we obtain $K \circ (\mathfrak{B}_p \vee \mathfrak{B}_{\neg p}) \circ (\mathfrak{B}_q \vee \mathfrak{B}_{\neg q})$ by going first from K to the closest belief set containing either p or $\neg p$, and then from there to the closest belief set containing either q or $\neg q$. (This corresponds to making up one's mind first about p and then about q .) It was also shown in [30] that Ramsey test conditionals can be introduced in descriptor revision without giving rise to the paradoxical results that they generate in other systems. In addition to standard (sentential) Ramsey test conditionals, a more general variant was defined, representing statements of the form “if the belief state is changed to satisfy Ψ then it will satisfy Ξ ”.

Another important development of descriptor revision is its restriction to various more specified types of descriptors [27]. We can define a sentential revision operation $*$ such that $K * p = K \circ \mathfrak{B}_p$ for all p . If $*$ is distance-based in the way described above, then $*$ will be weaker than AGM revision; in particular it will not satisfy the supplementary AGM postulates (see above, Section 2) but instead the following much weaker postulate:

$$\text{If } q \in K * p \text{ then } K * p = K * (p \& q) \quad (\text{cumulativity})^{17}$$

Analogously, a sentential operation $\div$ can be based on the success condition of contraction, thus $K \div p = K \circ \neg \mathfrak{B}_p$ for all p . The resulting operation is called revocation. It is not a contraction in the usual sense since it does not satisfy the inclusion postulate ($K \div p \subseteq K$ for all p). This may not be a disadvantage since the inclusion property of contraction is far from unproblematic. In real life, belief changes that lead to the removal of a certain belief are normally prompted by the acquisition of some new information that is added to the belief set. The only credible examples of pure contraction (i.e. removal of a sentence without acquisition of some other belief that pushes it out) that have been presented in the literature are hypothetical contractions such as contractions for the sake of argument [6, 34]. In this respect the operation of revocation that can be derived from descriptor revision is arguably more realistic than the standard operations of contraction in the belief change literature [27].

Finally it should be mentioned that transitively relational AGM revision is reconstructible as descriptor revision, i.e. for all such AGM revisions $*$ there is some descriptor revision $\circ$ such that $K * p = K \circ \mathfrak{B}_p$ for all p [29].

¹⁷On this postulate, see [39, 41].

However, there is no way to reconstruct AGM contraction as descriptor revision [31].

10. Conclusion

In Sections 3, 4, 5, 6 and 7 we identified a number of problems in traditional models of belief revision. In Sections 8 and 9 an alternative approach, descriptor revision, was stepwise developed from basic principles. In conclusion, let us briefly review how it deals with the major problems of the traditional framework that justified its introduction.

The *select-and-intersect method* (Sections 3 and 4) has been dispensed with in descriptor revision, and is replaced by direct selection among the potential outcomes of belief change. The main precondition for doing this was the introduction of an *outcome set* consisting of all the outcomes that can be reached by an operation of change. It can be interpreted as the collection of all those belief sets within reach that are coherent, stable, and/or plausible enough to be suitable as outcomes of an operation of belief change.

Success conditions that are not preserved under intersection (Section 5) cannot be dealt with in models employing the select-and-intersect method, but in descriptor revision they pose no special problem. Two examples are the success conditions of choice revision ($\mathfrak{B}p_1 \vee \dots \vee \mathfrak{B}p_n$) and making up one's mind ($\mathfrak{B}p \vee \mathfrak{B}\neg p$).

There are at least two reasons why possible worlds and remainder sets are *cognitively implausible* as outcome candidates (Section 6). First, they contain beliefs that have no ground in either the original belief set or the input. Secondly, they are logically infinite whenever the language is so, even if the original belief set was logically finite. In descriptor revision, both these problems are easily avoided by constructing the outcome set as a collection of cognitively realistic belief sets. In realistic applications, all its elements should be finite-based. The fact that descriptor revision does not require the introduction of intermediate, cognitively inaccessible objects into the belief change process is clearly an advantage from the viewpoint of Ockhamist requirements on simplicity.

Finally, neither the *recovery postulate* for contraction (Section 7.1) nor the *expansion property* of revision (Section 7.2) hold in descriptor revision. As a consequence of the latter, the *impossibility of Ramsey test conditionals* in the AGM framework does not hold for descriptor revision.

Note added in proof: Some of the ideas presented here have been further developed in Sven Ove Hansson, *Descriptor Revision. Belief Change Through Direct Choice* (Cham: Springer 2017), in particular Chapter 4.

Open Access. This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix: Preservation of Success Conditions Under Intersection

The following result holds for all success conditions that are expressible as the requirement that the belief set satisfies some descriptor Ψ .

DEFINITION 7. A descriptor Ψ is *preserved under intersection* if and only if it holds for all sets $\mathbb{Y}$ of belief sets that if $K \Vdash \Psi$ for all $K \in \mathbb{Y}$, then $\bigcap \mathbb{Y} \Vdash \Psi$.

OBSERVATION 4. A descriptor is preserved under intersection if each of its elements has one of the three forms

- (i) $\mathfrak{B}p$,
- (ii) $\neg \mathfrak{B}p$, or
- (iii) $\mathfrak{B}p_1 \vee \dots \vee \mathfrak{B}p_n \vee \neg \mathfrak{B}q$, with $q \vdash p_1 \vee \dots \vee p_n \rightarrow p_k$ for some p_k .

PROOF. Let Ψ consist of elements of the forms (i), (ii), and (iii). Let $X \Vdash \Psi$ for all $X \in \mathbb{Y}$. Then it holds for each $\alpha \in \Psi$ that $X \Vdash \alpha$ for all $X \in \mathbb{Y}$. We are going to show that $\bigcap \mathbb{Y} \Vdash \alpha$. There are three cases:

Case (i): α has the form $\mathfrak{B}p$: Then $p \in X$ for all $X \in \mathbb{Y}$, thus $p \in \bigcap \mathbb{Y}$, thus $\bigcap \mathbb{Y} \Vdash \alpha$.

Case (ii): α has the form $\neg \mathfrak{B}p$: Then $p \notin X$ for all $X \in \mathbb{Y}$, thus $p \notin \bigcap \mathbb{Y}$, thus $\bigcap \mathbb{Y} \Vdash \alpha$.

Case (iii): α has the form shown in (iii) in the observation: If $q \notin X$ for some $X \in \mathbb{Y}$ then $q \notin \bigcap \mathbb{Y}$ and we are done. If $q \in X$ for all $X \in \mathbb{Y}$ then for each $X \in \mathbb{Y}$ then there is some $p_m \in \{p_1, \dots, p_n\}$ such that $p_m \in X$ and consequently $q \& (p_1 \vee \dots \vee p_n) \in X$. Thus $q \& (p_1 \vee \dots \vee p_n) \in \bigcap \mathbb{Y}$. We have $q \vdash p_1 \vee \dots \vee p_n \rightarrow p_k$ for some p_k , thus $p_k \in \bigcap \mathbb{Y}$ and $\bigcap \mathbb{Y} \Vdash \alpha$.

We can conclude that $\bigcap \mathbb{Y} \Vdash \alpha$ for each $\alpha \in \Psi$, thus $\bigcap \mathbb{Y} \Vdash \Psi$. ■

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] ALCHOURRÓN, C., and D. MAKINSON, On the logic of theory change: safe contraction, *Studia Logica* 44:405–422, 1985.
- [4] FERMÉ, E., and S.O. HANSSON, AGM 25 years. Twenty-five years of research in belief change, *Journal of Philosophical Logic* 40:295–331, 2011.
- [5] FUHRMANN, A., *Relevant Logics, Modal Logics, and Theory Change*. Doctoral thesis, Australian National University, Canberra, 1988.
- [6] FUHRMANN, A., Theory contraction through base contraction, *Journal of Philosophical Logic* 20:175–203, 1991.
- [7] FUHRMANN, A., and S.O. HANSSON, A survey of multiple contraction, *Journal of Logic, Language, and Information* 3:39–76, 1994.
- [8] GALLIER, J.R., Autonomous belief revision and communication, in P. Gärdenfors, (ed.), *Belief Revision*. Cambridge University Press, Cambridge, 1992, pp. 220–246.
- [9] GÄRDENFORS, P., Rules for rational changes of belief, in T. Pauli, (ed.), *Philosophical Essays Dedicated to Lennart Åqvist on His Fiftieth Birthday*. Philosophical studies, vol. 34. University of Uppsala, Department of Philosophy, Uppsala, 1982, pp. 88–101.
- [10] GÄRDENFORS, P., Belief revision and the Ramsey test for conditionals, *Philosophical Review* 95:81–93, 1986.
- [11] GÄRDENFORS, P., *Knowledge in Flux. Modeling the Dynamics of Epistemic States*, The MIT Press, Cambridge, 1988.
- [12] GÄRDENFORS, P., and D. MAKINSON, Revisions of knowledge systems using epistemic entrenchment, in M.Y. Vardi, (ed.), *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, Morgan Kaufmann, Los Altos, pp. 83–95, 1988.
- [13] GLAISTER, S.M., Recovery recovered, *Journal of Philosophical Logic* 29:171–206, 2000.
- [14] GROVE, A., Two modellings for theory change, *Journal of Philosophical Logic* 17:157–170, 1988.
- [15] HANSSON, S.O., Belief contraction without recovery, *Studia Logica* 50:251–260, 1991.
- [16] HANSSON, S.O., In defense of the Ramsey test, *Journal of Philosophy* 89:522–540, 1992.
- [17] HANSSON, S.O., Theory contraction and base contraction unified, *Journal of Symbolic Logic* 58:602–625, 1993.
- [18] HANSSON, S.O., Taking belief bases seriously, in D. Prawitz and D. Westerståhl, (eds.), *Logic and Philosophy of Science in Uppsala*, Kluwer, Dordrecht, 1994, pp. 13–28.
- [19] HANSSON, S.O., *A Textbook of Belief Dynamics. Theory Change and Database Updating*, Kluwer, Dordrecht, 1999.
- [20] HANSSON, S.O., Recovery and epistemic residues, *Journal of Logic, Language and Information* 8:421–428, 1999.

- [21] HANSSON, S.O., Three approaches to finitude in belief change, in T. Rønnow-Rasmussen, B. Petersson, J. Josefsson and D. Egonsson, (eds.), *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*, Lund University, Lund, 2007.
- [22] HANSSON, S.O., Specified meet contraction, *Erkenntnis* 69:31–54, 2008.
- [23] HANSSON, S.O., Replacement: a Sheffer stroke for belief revision, *Journal of Philosophical Logic* 38:127–149, 2009.
- [24] HANSSON, S.O., Contraction, revision, expansion: representing belief change operations, in R. Trypuz, (ed.), *Krister Segerberg on Logic of Actions*, Springer, Berlin, 2013, pp. 135–152 .
- [25] HANSSON, S.O., Outcome level analysis of belief contraction, *Review of Symbolic Logic* 6:183–204, 2013.
- [26] HANSSON, S.O., Maximal and perimaximal contraction, *Synthese* 190:3325–3348, 2013.
- [27] HANSSON, S.O., Descriptor revision, *Studia Logica* 102:955–980, 2014.
- [28] HANSSON, S.O., Relations of epistemic proximity for belief change, *Artificial Intelligence* 217:76–91, 2014.
- [29] HANSSON, S.O., A monoselective presentation of AGM revision, *Studia Logica* 103:1019–1033, 2015.
- [30] HANSSON, S.O., Iterated descriptor revision and the logic of Ramsey test conditionals, *Journal of Philosophical Logic* 45:429–450, 2016.
- [31] HANSSON, S.O., AGM contraction is not reconstructible as a descriptor operation, *Journal of Logic and Computation* 27(4):1133–1141, 2017.
- [32] HANSSON, S.O., Belief change, in S.O. Hansson and V.F. Hendricks, (eds.), *Introduction to Formal Philosophy*. Springer, Berlin, 2018.
- [33] LEVI, I., Subjunctives, dispositions and chances, *Synthese* 34:423–455, 1977.
- [34] LEVI, I., *The Fixation of Belief and Its Undoing*, Cambridge University Press, Cambridge, 1991.
- [35] 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.
- [36] MAKINSON, D., On the status of the postulate of recovery in the logic of theory change, *Journal of Philosophical Logic* 16:383–394, 1987.
- [37] MAKINSON, D., On the force of some apparent counterexamples to recovery, in E. Garzón Valdéz et al., (eds.), *Normative Systems in Legal and Moral Theory: Festschrift for Carlos Alchourrón and Eugenio Bulygin*, Duncker & Humblot, Berlin, 1997, pp. 475–481.
- [38] ROTT, H., Conditionals and theory change: revisions, expansions, and additions, *Synthese* 81:91–113, 1989.
- [39] ROTT, H., *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*, Oxford University Press, Oxford, 2001.
- [40] ROTT, H., Basic entrenchment, *Studia Logica* 73:257–280, 2003.
- [41] ROTT, H., Three floors for the theory of theory change, in V. Punčochář and M. Dančák, (eds.), *Logica Yearbook 2013*, College Publications, London, 2014.

- [42] ROTT, H., and S.O. HANSSON, Safe contraction revisited, in S.O. Hansson, (ed.), *David Makinson on Classical Methods for Non-Classical Problems*, Springer, Dordrecht, 2014, pp. 35–70.
- [43] SANDQVIST, T., On why the best should always meet, *Economics and Philosophy* 16:287–313, 2000.
- [44] ZHANG L., and S.O., HANSSON, How to make up one’s mind, *Logic Journal of the IGPL* 23:705–717, 2015.

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