Chapter 5
Criticism of the Model

Although the AGM model is commonly considered to be the standard model of belief change, it has been subject to extensive criticism. In this chapter we make a résumé of this criticism. Much of the critical discussion has referred either to the postulates for partial meet contraction and revision, or to various aspects of the use of belief sets to represent belief states.

For additional elaborations on philosophical issues relating to the AGM model, see also [165, 190, 294].

5.1 The Recovery Postulate and Minimal Change

One of the basic principles of the AGM theory [4] is that belief changes should take place with minimal loss of previous beliefs. In the opinion of the AGM trio, the postulate of *recovery* guarantees minimal losses of contents in the contraction process [133, p. 65] [238, p. 352] [241, p. 478]. *Recovery* is based on the intuition that "it is reasonable to require that we get all of the beliefs ... back again after first contracting and then expanding with respect to the same belief" [130]. This can be exemplified as follows:

Example 5.1. [163] I believed that I had my latchkey on me (p). Then I felt in my left pocket, where I usually keep it, and did not find it. I lost my belief in p (but without starting to believe in $\neg p$ instead). Half a second later, I found the key, and regained my belief in p.

However, counter-examples have been constructed in which *recovery* seems to give rise to implausible results:

Example 5.2. [148] 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 that Cleopatra had a daughter). Then I received information that made me give up my belief in c, and I contracted my belief set accordingly, forming $K - c$. Soon afterwards I learned from

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_5

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.

This pattern is incompatible with recovery, and therefore also with partial meet contraction. The problem (if we see it as such) is that every element of $K \perp (s \vee d)$ has $s \vee d \rightarrow s$ as one of its elements, and therefore $s \vee d \rightarrow s$ is also an element of $K - (s \vee d)$. Consequently, $(K - (s \vee d)) + (s \vee d)$ implies s.

Example 5.3. [157][1] I previously entertained the two beliefs, "*x is divisible by 2*" (p) and "*x is divisible by 6*" (q). When I received new information that induced me to give up the first of these beliefs (p), the second (q) had to go as well (since p would otherwise follow from q).
I then received new information that made me accept the belief "*x is divisible by 8.*" (r). Since p follows from r, $(K - p) + p$ is a subset of $(K - p) + r$; thus by *recovery* I obtain that "*x is divisible by 24*" (s), contrary to intuition.

This example shows that retaining the sentence $p \rightarrow q$ after contraction of K by p gives rise to unintuitive results. There seem to be cases when this sentence has to be removed. Due to recovery, AGM contraction cannot eliminate it.

Makinson [241, p. 478] noted that "as soon as contraction makes use of the notion 'y is believed only because of x', we run into counterexamples [to *recovery*]". He argued that this is provoked by the use of a justificatory structure that is not represented in the belief set and that, without this structure, *recovery* can be accepted; or, in Makinson's words, it can be accepted in a "naked" theory. In [161], Hansson replied: "Actual human beliefs always have such a justificatory structure (...). It is difficult if not impossible to find examples about which we can have intuitions, and in which the belief set is not associated with a justificatory structure that guides our intuitions. Against this background, it is not surprising that, as Makinson says, *recovery* 'appears to be free of intuitive counterexamples' in which the belief set is not associated with a justificatory structure. It also seems to be free of confirming examples of this kind". Glaister argued that the problem exhibited in the counterexamples dissolves if we pay sufficient attention to exactly what is to be contracted. In the Cleopatra case, he claims, the contraction is more accurately represented by a multiple contraction by the set $\{s \vee d, s \vee \neg d, \neg s \vee d\}$ than by $s \vee d$ [141].

Niederée [259] pointed out several implausible formal properties that follow from recovery:

Observation 5.1. [259] Let K be a belief set and $p \in K$. Then, regardless of whether or not q is in K, *recovery* together with *closure* implies that:

1. $q \rightarrow p \in K - (p \vee q)$,
2. $p \in (K - (p \vee q)) + q$,
3. $\neg q \in (K - (p \vee q)) + \neg p$.

[1] We use here the modified version introduced in [105].

As we mentioned in Section 3.5, operations that satisfy the other five basic contraction postulates but not *recovery* are called withdrawals [239]. In the literature several withdrawals have been proposed [58, 95, 105, 223, 224, 250, 302].

An obvious reaction to these difficulties would be to replace *recovery* by some other postulate that puts a limit to how much can be removed in a contraction. However, no plausible such alternative postulate has been presented. In particular, although the following postulates are more intuitively appealing, they are equivalent to *recovery* in the presence of the other basic AGM postulates:

If $q \in K$ and $q \notin K - p$ then there is some set K' such that $K - p \subseteq K' \subseteq K$ and $p \notin Cn(K')$ but $p \in Cn(K' \cup \{q\})$. (relevance) [146]

If $q \in K$ and $q \notin K - p$ then there is some set K' such that $K' \subseteq K$ and $p \notin Cn(K')$ but $p \in Cn(K' \cup \{q\})$. (core-retainment) [148]

If $q \in K$ and $q \notin K - p$ then $K - p \nvdash p \vee q$. (disjunctive elimination) [91]

5.2 The Success Postulates

Partial meet revision satisfies the following postulate:

Revision success: $p \in K * p$

Several authors have found this to be an implausible feature of belief revision, even if p is not a contradiction. Hence Cross and Thomason pointed out that a system obeying this postulate

> "is totally trusting at each stage about the input information; it is willing to give up whatever elements of the background theory must be abandoned to render it consistent with the new information. Once this information has been incorporated, however, it is at once as susceptible to revision as anything else in the current theory.
>
> Such a rule of revision seems to place an inordinate value on novelty, and its behaviour towards what it learns seems capricious." [65]

Similarly, one of the AGM postulates for partial meet contraction:

Contraction success: If $\nvdash p$, then $p \notin K - p$,

has been contested on the grounds that we should "allow a reasoner to refuse the withdrawal of p not only in the case where p is a logical truth. There may well be other sentences ('necessary truths') which are of topmost importance for him" [290, p. 54]. Both with respect to *revision success* and *contraction success*, a common strategy among critics has been to construct AGM-style operations that do not always give primacy to the new information. (See Chapter 8.)

5.3 Remainder Sets: Information vs. Informational Value

Some researchers have argued that remainder sets retain “too much information”. According to Levi [222, 223], some of the information in the belief set may have no value for the inquiring agent; consequently, the agent tries to retain as much of the valuable information as possible, instead of as much of the information as possible. He argued that measures of information should be replaced by measures of informational value,[2] and proposed an alternative construction, based on saturatable sets: H is a saturatable set with respect to p if and only if $H = Cn(H)$ and $Cn(H \cup \{\neg p\}) \in \mathcal{L} \perp\perp$. Furthermore, $S(K,p)$ is the set of saturatable sets with respect to p that are subsets of K. As Alchourrón and Makinson proved, $K \perp p \subseteq S(K,p)$ [5]. *Partial meet Levi contraction*, based on a selection among all the saturatable subsets of K with respect to p, is defined as $K - p = \bigcap \gamma(S(K,p))$, where γ is a selection function defined in the same way as in the AGM account. Hansson and Olsson [193] proved that an operation $-$ on K is a partial meet Levi contraction if and only if it satisfies *closure, inclusion, vacuity, success, extensionality*, and *failure*.

The main problem with this construction is that it allows for very drastic contractions. For instance, the following rather extreme operation is a partial meet Levi contraction: [193, pp. 111–112]

$$K - p = \begin{cases} K & \text{if } p \notin K \text{ or } \vdash p \\ Cn(\varnothing) & \text{otherwise} \end{cases}$$

5.4 The Expansion Property

It follows from the basic AGM postulates that the revision operation satisfies the following property:

If $K \nvdash \neg p$ then $K * p = Cn(K \cup \{p\})$. (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 4.7). In the original AGM formulation of partial meet revision it follows from the Levi identity in combination with the contraction postulate *vacuity*.[3] However, it is easy to find examples in which it does not seem plausible:

[2] “... when seeking to answer a question, not all new information is relevant to the question being asked. This is, perhaps, the chief of several reasons why measures of informational value ought to be carefully distinguished from measures of information” [222, p. 123].

[3] According to the Levi identity, $K * p = Cn((K - \neg p) \cup \{p\})$. When $K \nvdash \neg p$, *vacuity* yields $K - \neg p = K$.

Example 5.4. [151] 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).

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

This example exemplifies a quite common pattern of belief change: When we acquire a new belief that does not contradict our previous beliefs (such as t in the example), we often include in the outcome some additional belief (such as r in the example) that does not follow deductively but nevertheless serves to make the belief set more complete and/or more coherent.

The expansion property can also go wrong in the opposite direction, as illustrated in the following example:

Example 5.5. Valentina was uncertain whether or not her husband is unfaithful to her (u), 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 this case we have $l \in K$ and $K \nvdash \neg u$. The *expansion property* of revision requires that $K * u = Cn(K \cup \{u\})$; thus $l \in K * u$, contradicting the plausible pattern of belief change exhibited in the example.

The expansion property of revision has been much less discussed than the recovery property of contraction, but it is no less problematic and no less difficult to remove from the AGM framework. Both properties have provided impetus for the development of alternative frameworks.

In [268] and [269], Pagnucco and coworkers introduced a new belief change operation, *abductive expansion*. Unlike AGM expansion (consisting in $K+p = Cn(K \cup \{p\}))$, in abductive expansion the agent also incorporates a justification or explanation of the new belief. The justification is the "abduction" of a formula p and it can be defined as follows:

Definition 5.1. [269] Let K be a belief set. An *abduction function* for K is a function f such that for each sentence p:

1. If $K \cup \{p\}) \nvdash \bot$, then $K \cup \{f(p)\} \vdash p$ and $K \cup \{f(p)\} \nvdash \bot$
2. If $K \cup \{p\}) \vdash \bot$, then $f(p) = \top$
3. If $\vdash p \leftrightarrow p'$, then $\vdash f(p) \leftrightarrow f(p')$

Definition 5.2. $\oplus$ is an *abductive expansion* for K if and only if there is an abduction function f fot K such that $K \oplus p = K + \{f(p)\}$ for all p.

The postulates that characterize abductive expansion are:

$K \oplus p$ is a belief set. (closure)
If $\neg p \notin K$, then $p \in K \oplus p$. (limited success)
$K \subseteq K \oplus p$. (inclusion)
If $\neg p \in K$, then $K \oplus p = K$. (vacuity)
If $\neg p \notin K$ then $\neg p \notin K \oplus p$. (consistency)
If $\vdash p \leftrightarrow q$, then $K \oplus p = K \oplus q$. (extensionality)

For the supplementary level, disjunctive factoring is added:

Either $K \oplus (p \vee q) = K \oplus p$, or $K \oplus (p \vee q) = K \oplus q$, or $K \oplus (p \vee q) = (K \oplus p) \cap (K \oplus q)$. (disjunctive factoring)

A semantic account of abduction can be based on Grove's systems of spheres. The construction is very similar to the one used in AGM belief revision. However, whereas $\|K\|$ is the innermost sphere in AGM revision, in abductive expansion there may be spheres that are proper subsets of $\|K\|$.

5.5 Are Belief Sets Too Large?

Belief sets have been criticized for being too extensive in two important respects that are both problematic from the viewpoint of cognitive realism: their logical closure and their infinite structure.

The use of a logically closed belief set to represent the belief state has important implications. In particular it means that all beliefs are treated as if they have independent status. Suppose you believe that you have your keys in your pocket (p). It follows that you also believe that either you have your keys in your pocket or the Archbishop of York is a Quranist Muslim ($p \vee q$). However, $p \vee q$ has no independent standing; it is in the belief set only because p is there. Therefore, if you give up your belief in p we should expect $p \vee q$ to be lost directly, without the need for any mechanism to select it for removal. In the AGM framework, however, "merely derived" beliefs such as $p \vee q$ have the same status as independently justified beliefs such as p. Belief base models (to be discussed in Chapter 6) have largely been constructed in order to distinguish between these two types of beliefs.

The logical closure of belief sets is also problematic from another point of view. In a study of the philosophical foundations of AGM, Hans Rott pointed out that the theory is unrealistic in its assumption that epistemic agents are "ideally competent regarding matters of logic. They should accept all the consequences of the beliefs they hold (that is, their set of beliefs should be logically closed), and they should rigorously see to it that their beliefs are consistent" [294].[4] In the same article he argued that the AGM model is not based on a principle of minimal change, something that has often been taken for granted.

[4] See Section 10.3 for resource-bounded agents.

However, as we noted in Section 3.2, logical closure only requires that the agent be able to draw the inferences that have been incorporated into the consequence operation *Cn*, and in the minimal case this does not go beyond classical sentential truth-functional logic. Furthermore, Isaac Levi has proposed that the belief set K should be interpreted as containing the statements that the agent is *committed to* believing, rather than those that she actually believes in [220, 222]. Such an interpretation may have other problems, but it defuses problems created by the high demands on inferential competence that seem to follow from logical closure.

Since actual human agents have finite minds, a good case can be made that a cognitively realistic model of belief change should be finitistic, and this in two senses. First, both the original belief set and the belief sets that result from a contraction should be finite-based, i.e., obtainable as the logical closure of some finite set. Secondly, the outcome set, i.e., the class of belief sets obtainable by contraction from the original belief set ($\{X \mid (\exists p)(X = K - p)\}$), should be finite [153, 179]. Partial meet contraction does not in general satisfy either of these two finitistic criteria. To the contrary, even if the original belief set K is finite-based, the standard AGM axioms do not ensure that the outcome $K - p$ of contracting it by a single sentence p is also finite-based [169].[5] And even if both K and $K - p$ are finite-based, the procedure that takes us from K to $K - p$ involves a choice among infinitely many sets (the elements of $K \perp p$), none of which is finite-based. Such a procedure does not seem to satisfy reasonable criteria of cognitive realism. This has led to the development of finitistic models such as belief base models (Chapter 6) and specified meet contraction (Section 4.3).

5.6 Lack of Information in the Belief Set

Belief sets have been criticized not only for being too large but also for lacking important information.

Most importantly, AGM contraction or revision in its original form is a "one shot" operation. After partial meet contraction of K by p we obtain a new belief set $K - p$ but we do not obtain a new selection function to be used in further operations on this new belief set. In other words, the original AGM framework does not satisfy the principle of *categorial matching*, according to which the representation of a belief state after a change should have the same format (and contain the same types of information) as the representation of the belief state before the change[6] [139]. In studies of iterated revision, various ways to extend the belief state representation to solve this problem have been investigated. (See Chapter 7.)

---

[5] Unless, of course, the language $\mathcal{L}$ is logically finite, by which is meant that it does not contain an infinite set of logically non-equivalent sentences. Arguably, this is a strong and implausible condition, since it excludes the possibility of expressing natural numbers of unlimited size.

[6] However, a fairly small rearrangement of the AGM definition is sufficient to make iterated partial meet contraction possible. We will return to that in Section 7.3.