Chapter 10
Alternative Operations of Change

In the original AGM model there are three major types of operations: contraction, revision, and expansion. Subsequently a large number of additional types of operations have been proposed. In this section we summarize some of them.

10.1 Update

In 1992, Katsuno and Mendelzon presented a type of operation of change that they called update [206]. Whereas revision operations are suited to capture changes that reflect evolving knowledge about a static situation, update operations are intended to represent changes in beliefs that result from changes in the objects of belief. The difference was pointed out for the first time by Keller and Winslett [207] (in the context of relational databases) and is captured in the following example [338]:

Example 10.1. Initially the agent knows that there is either a book on the table (p) or a magazine on the table (q), but not both.

Case 1: The agent is told that there is a book on the table. She concludes that there is no magazine on the table. This is revision.

Case 2: The agent is told that subsequently a book has been put on the table. In this case she should not conclude that there is no magazine on the table. This is update.

This difference is evident in the possible worlds approach. Katsuno and Mendelzon proposed that when the world changes, the agent changes each of the worlds that (s)he considers to be possible in order to accommodate the input while changing as little else as possible. They constructed update as follows (see Figure 10.1 for a graphical representation):

Definition 10.1. A *local faithful assignment* is a function mapping each possible world ω to a total preorder $\leq_\omega$ such that if $\omega \neq \omega'$, then $\omega <_\omega \omega'$.

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_10

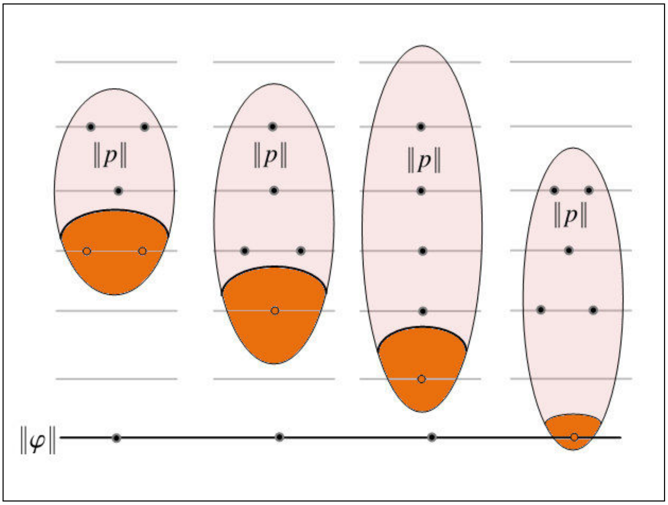

Fig. 10.1: Dots represent possible worlds and the line at the bottom represents $\|\varphi\|$. Each column represents the distribution and ordering of the possible worlds regarding the *local faithful assignment* of one world ω_i in $\|\varphi\|$. $\|\varphi \diamond p\|$ is the union of all the dark orange regions.

Definition 10.2. Let K be a finite-based belief set. Let $\varphi \in \mathcal{L}$ be such that $Cn(\varphi) = K$. An operation $\diamond$ on φ is an *update* if and only if there is a local faithful assigment such that:

$$\|\varphi \diamond p\| = \bigcup\{\min(\|p\|, \leq_\omega) : \varphi \in \omega\}$$

Update on finite-based belief sets has been axiomatically characterized as follows:

Theorem 10.2. [206] An operation $\diamond$ is an update operation if and only if it satisfies:

(U1) $\varphi \diamond p \vdash p$
(U2) If $\varphi \vdash p$, then $\vdash \varphi \diamond p \leftrightarrow \varphi$.
(U3) If $\varphi \nvdash \bot$ and $p \nvdash \bot$, then $\varphi \diamond p \nvdash \bot$.
(U4) If $\vdash \varphi_1 \leftrightarrow \varphi_2$ and $\vdash p_1 \leftrightarrow p_2$, then $\vdash \varphi_1 \diamond p_1 \leftrightarrow \varphi_2 \diamond p_2$.
(U5) $(\varphi \diamond p_1) \wedge p_2$ implies $\varphi \diamond (p_1 \wedge p_2)$.
(U6) If $\varphi_1 \diamond p_1 \vdash p_2$ and $\varphi_2 \diamond p_2 \vdash p_1$, then $\vdash \varphi_1 \diamond p_1 \leftrightarrow \varphi_2 \diamond p_2$.
(U7) If φ is complete,[1] then $(\varphi \diamond p) \wedge (\varphi \diamond q)$ implies $\varphi \diamond (p \vee q)$.
(U8) $\vdash (\varphi_1 \vee \varphi_2 \diamond p) \leftrightarrow (\varphi_1 \diamond p) \vee (\varphi_2 \diamond p)$.

It follows from (U2) that if $\varphi \vdash \bot$, then $\varphi \diamond p \vdash \bot$ for all p. In other words, if a belief set is inconsistent, then consistency cannot be regained with an update.[2]

[1] φ is complete if and only if for all $p \in \mathcal{L}$, $\varphi \vdash p$ or $\varphi \vdash \neg p$.

[2] With respect to this property, Katsuno and Mendelzon said [206, p. 190]: "We can never repair an inconsistent theory using update, because update specifies a change in the world. If there is no

If we compare update with AGM revision (see postulates (R1)-(R6) in Section 3.6), we can note some interesting formal differences. In particular, postulate R2 (vacuity) does not hold for update. Update and its relation with revision have been further studied by Becher [28] and others.

10.2 Changes in the Strengths of Beliefs

Sometimes when a statement is presented to us, this makes us consider it to be more credible than before, but we still do not believe it. Such a change may not affect the belief set, but it will affect how the belief state responds to new inputs. This kind of belief change was studied by Cantwell. He introduced the operations of *raising* and *lowering*, whereby the degree of plausibility required for a sentence to be included into the belief set is changed in either direction [57].

One important class of such operations is the improvement operations investigated by Konieczny and Pérez [213]. These operations do not (necessarily) satisfy the success postulate, although they increase ("improve") the agent's estimate of the plausibility of the new information [211, 213]. For instance, if you see what looks like wolf tracks in your garden, this makes it more plausible than before that a wolf has visited your garden, but you will presumably still not believe it.

In the construction of improvement operations it has often been assumed that if the agent receives the same new information sufficiently many times, then (s)he will finally believe it. For an epistemic state Ψ, an operation $\circ$ and a natural number n, $\circ^n$ is defined by recursion in the following way:

$$\begin{aligned} \Psi \circ^0 p &= \Psi \\ \Psi \circ^{n+1} p &= (\Psi \circ^n p) \circ p \end{aligned}$$

and the operation $\star$ is defined as $\Psi \star p = \Psi \circ^n p$, where n is the first integer such that $B(\Psi \circ^n p) \vdash p$. B is a function that takes us from a belief state to its associated belief set. The key property of this operation is:

There exists an integer n such that $B(\Psi \circ^n p) \vdash p$. (iterative success)

Another interesting property is

If $B(\Psi) \nvdash p$, then there is some q such that $\nvdash B((\Psi \circ p) \star q) \leftrightarrow B(\Psi \star q)$. (non-triviality)

This property says that any revision by a formula p that is not a consequence of the epistemic state modifies the epistemic state of the agent. In a possible worlds model, improvement by p means that some or all of the p-worlds are moved to a

set of worlds that fits our current description, we have no way of recording the change in the real world."

higher position in the preorder, but this does not necessarily lead to a change in the belief set.

Improvement operations have been combined with credibility-limited revision in the following way: A new piece of information is accepted if it is judged credible by the agent. However, if it is not considered credible, then its epistemic status is nevertheless raised, in the manner that this would be done by an improvement operation [46].

In quantitative theories of belief change, such as probabilistic and ranking theories (Sections 11.1 and 11.2), the degree of acceptance of each sentence is represented by a numerical value. Changes in the strength of beliefs can then be represented as changes in these values. However, the meaning of these numbers is not entirely clear (especially not for non-probabilistic functions), and real agents are notoriously bad at reasoning with them [324]. These difficulties are largely avoided by using a preorder instead of a numerical representation. Comparative degrees of belief can then be specified by taking certain beliefs as points of reference. The operation of change will adjust the position of an input sentence in an ordering to be the same as that of a reference sentence. Such an operation requires two sentences as inputs: the sentence to be adjusted and the reference sentence to which it will be adjusted. Since two sentences are involved, Rott called such operations two-dimensional [299].

Fermé and Rott proposed the operation of *revision by comparison*. In the intended case, the input sentence p is accepted to the same degree as a previously believed sentence q. However, if the negation of the input sentence p is more plausible than the reference sentence q, then q will be removed from the outcome [107]. Therefore, revision by comparison violates the DP postulates for iterated change that we mentioned in Section 7.1 (in particular DP2 since it collapses distinctions between the positions in the ordering of some $\neg p$-worlds). Rott has proposed a variant, *bounded revision*, that captures the spirit of revision by comparison and also satisfies the DP postulates [300]. As Rott pointed out, revision by comparison reduces the number of equivalence classes in the preorder, whereas bounded revision increases it.

10.3 Resource-Bounded and Local Change

AGM is a theory of changes of beliefs undertaken by highly idealized reasoners with unlimited cognitive capacities. In contrast, real reasoners such as humans, computers, and robots have limited resources. As was noted by Wassermann, it is important to distinguish between a limited implementation of a theory for ideal reasoning, and a theory for reasoners with limited resources [327].

Harman has put forward a highly useful list of principles that should be valid for any resource-bounded agent [196]:

Clutter avoidance: One should not clutter one's mind with trivialities.

Recognized implication: One has a reason to believe p if one recognizes that p is implied by one's views.

Recognized inconsistency: One has a reason to avoid believing things that one recognizes to be inconsistent.
Positive undermining: One should stop believing p whenever one positively believes that one's reasons for believing p are no good.
Conservatism: One is justified in continuing to fully accept something in the absence of a special reason not to.
Interest condition: One should add a new proposition p to one's beliefs only if one is interested in whether p is true (and it is otherwise reasonable for one to believe p).
Get back principle: One should not give up a belief one can easily (and rationally) get right back.

Doyle investigated characteristics of real agents such as mental inertia and constitutional elasticity [80]. He proposed a formal structure, a *reason maintenance system* (RMS), to capture these characteristics. Alechina, Jago, and Logan used RMS to construct a resource-bounded operation of contraction [9].

The two features of resource-boundedness that have attracted most attention among researchers are finitude and inconsistency tolerance. Both belief bases (Chapter 6) and specified meet contraction (Section 4.3) have been constructed largely in order to avoid the infinite structures of the standard AGM model.

Gabbay and Hunter maintain that there is a fundamental difference between how inconsistencies are handled by real agents and how they have usually been treated in formal logical systems. For a real agent it need not be necessary to restore consistency; it may be sufficient to have rules that specify how to act when an inconsistency arises [122]. What makes inconsistencies devastating in the AGM model is that there is only one inconsistent belief set, namely the whole language. This is an unsatisfactory feature of belief set representation, since two agents can both have inconsistent beliefs without having the same beliefs. As we saw in Section 6.1, belief bases fare much better in this respect. There are many different inconsistent belief bases, and they can reasonably be taken to represent different inconsistent belief states [147]. This feature of belief bases was employed in Hansson and Wassermann's model of *local change* [195]. Given a belief base B and a sentence r, the r-compartment of B is the subset of B that is relevant for r. In local change, revision of B by r involves changes only of the r-compartment; hence a part of the belief base can be made consistent while the belief base as a whole remains inconsistent. Wassermann has shown how these principles can be used to provide a model of change that satisfies Harman's principles. This can be accomplished with a construction involving a short-term memory in which recently computed results are temporarily stored [326]. She also showed how local change can be used for diagnosis [328], i.e., the process of finding a compartment that may have caused an abnormal behaviour of the system [279].

In a similar vein, Parikh [270] proposed a principle for relevance-sensitive change. Its basic principle is that if a belief set can be split into two independent parts (expressed in different sublanguages), then a revision of one part does not affect the other. This principle is not satisfied by AGM revision. Peppas provided a semantics for Parikh's relevance-sensitive condition in terms of systems of spheres

[271], and Kourousias and Makinson have investigated the conditions under which Parikh's relevance-sensitive condition is satisfied [216, 243].

10.4 Paraconsistent Belief Change

Consistency preservation is a central requirement in AGM revision. The reason for this is that the underlying logic is supraclassical and therefore satisfies the *explosion principle*, namely that anything follows from a contradiction (ex contradictione quodlibet, $\{p, \neg p\} \vdash q$). Consequently there is, as we just noted, only one inconsistent belief set, namely the whole language. If we arrive at an inconsistent belief set, then we have lost all distinctions. To avoid this we have to steer clear of contradictions in all operations on belief sets in a supraclassical logic.

However, this does not seem to be how cognitive agents behave in practice. Real agents can believe in contradictory statements without believing everything and losing all distinctions. In order to model that feature of actual reasoning, we can weaken the consequence relation and make it paraconsistent (which means that the explosion principle does not hold). Relatively little work has been performed on paraconsistent belief revision, but important contributions have been made for instance by Restall and Slaney [280], Priest [273], Mares [248], Tanaka [320], and Testa, Coniglio and Ribeiro [322].

The underlying logic used by Restall and Slaney [280] avoids the explosion principle by demanding a connection between the premises and the conclusion of an inference. In a valid inference the premises have to be *relevant* to the conclusion. Mares [248] developed a model in which an agent's belief state is represented by a pair of sets. One of these is the belief set, and the other consists of the sentences that the agent rejects. A belief state is coherent if and only if the intersection of these two sets is empty, i.e., if and only if there is no statement that the agent both accepts and rejects. In this model, belief revision preserves coherence but does not necessarily preserve consistency.

Priest [273] and Tanaka [320] suggested that in a paraconsistent logic, revision can be performed by just adding sentences without removing anything. In other words, if the logic tolerates inconsistencies, then expansion can serve the function usually assigned to revision. Furthermore, Priest [273] pointed out that in a paraconsistent framework, revision on belief sets can be performed as external revision, i.e., with the reversed Levi identity. In a supraclassical framework, external revision can only be used on belief bases. (See further Section 6.3.) Testa, Coniglio and Ribeiro [322] showed that this holds for semi-revision as well. In a supraclassical system, semi-revision (defined in Section 8.2) can only be used for belief bases, but in a paraconsistent system it can also be used for belief sets. The reason for this difference is that the intermediate inconsistent belief set that arises in external revision and semi-revision extinguishes all distinctions if the underlying logic is supraclassical but not if it is paraconsistent.

10.5 Some Other Operations of Change

Indeterministic change: The AGM model and most other models of belief change are deterministic in the sense that given a belief set and an input, the resulting belief set is well-determined. There is no scope for chance in determining the outcome of the change. Although this may not be a realistic feature, it substantially simplifies the formal structure. In indeterministic belief change, an operation can have more than one admissible outcome. Indeterministic belief change has been studied by Gallier [126] and by Lindström and Rabinowicz [231]. The latter authors gave up the assumption that epistemic entrenchment satisfies connectedness. This resulted in Grove-style sphere systems with spheres ("fallbacks") that are not linearly ordered but still all include the original belief set.

Replacement is an operation that replaces one sentence by another in a belief set. An operation of replacement has two variables, such that in $K|_q^p$, p has been replaced by q. Hence, the outcome is a belief set that contains q but not p. This operation can have outcomes that are not obtainable through either partial meet contraction or partial meet revision. Replacement can also be used as a kind of Sheffer stroke for belief change, i.e., an operation in terms of which the other operations can be defined. Contraction by p can be defined as the replacement $|_\top^p$ of p by a tautology, revision by p as the replacement $|_p^\bot$ of falsum by p, and expansion by p as the replacement $|_p^\top$ of a tautology by p. (Tautologies are as usual taken to be unremovable.) Partial meet replacement has been axiomatically characterized, and a semantic account in terms of possible worlds has been provided [170].

Reconsideration, introduced by Johnson and Shapiro [204, 203], is a non-prioritized operation on belief bases. It represents changes that are performed in hindsight in order to eliminate negative effects of previously performed changes. Previously removed beliefs can be reintroduced if there are no longer any valid reasons for their removal. This operation can be seen as an optimization that eliminates the negative effects of the order in which inputs have been received. It can involve an examination of all current and previous beliefs, but the same result can also be produced by an algorithm that examines a subset of the retracted basic beliefs, using dependency relationships.