Chapter 20
Belief Change

Sven Ove Hansson

Abstract All formal models of belief change involve choices between different ways to accommodate new information. However, the models differ in their loci of choice, i.e. in what formal entities the choice mechanism is applied to. Four models of belief change with different loci of choice are investigated in terms of how they satisfy a set of important properties of belief contraction and revision. It is concluded that the locus of epistemic choice has a large impact on the properties of the resulting belief change operation.

20.1 Requirements on Belief Change

The rationality of beliefs can be discussed either in a static or a dynamic perspective. In a static perspective, we discuss what rationality requires of a person's state of beliefs, taken as a snapshot at a particular point in time. In a dynamic perspective, the topic is instead how one should rationally change one's state of belief in response to new information. The logical investigation of belief change came to light in the 1980s and is therefore a comparatively new area of formal philosophy.

In the static approach to belief rationality we can distinguish between on the one hand substantial requirements of rationality and on the other hand formal (or structural) requirements. For instance, suppose that after his university studies in biology, Donald still believe that dinosaurs and humans have once lived side by side. Then we would consider his belief system to be irrational, for substantial rather than formal (structural) reasons. If he believes that all snakes are poisonous, and he also believes that the black-tailed python is a non-poisonous snake, then this provides us with another, structural, type of reason to consider his current belief system to be irrational. The crucial criterion why this is a structural rather than a substantial failure of rationality is that we can discern the irrationality without even knowing

S. O. Hansson (✉)
Division of Philosophy, Royal Institute of Technology (KTH), Stockholm, Sweden
e-mail: soh@kth.se

S. O. Hansson, V. F. Hendricks (eds.), *Introduction to Formal Philosophy*, Springer Undergraduate Texts in Philosophy, https://doi.org/10.1007/978-3-319-77434-3_20

the meaning of the three terms “snake”, “black-tailed python” and “poisonous”. Structural rationality falls within the purview of logic and will therefore be at our focus here.

Can structural requirements be made also in the dynamic approach to belief rationality? More precisely, are there sensible rationality requirements on the process of belief change? The following examples are intended to show that there may indeed be such requirements, but also that these requirements may not be uncontroversial:

Example 1 For many years, Derek was confident that his wife was faithful to him. But one day a neighbour told him stories that convinced him that she was cheating on him. When he confronted her, she could explain everything, and he regained his previous strong belief in her faithfulness. But something strange happened. He never regained his belief that she loved him. He could not explain why. All misunderstandings had been cleared, and everything else was as before, but still he was unable to believe in her love any more.

Derek’s pattern of belief change contradicts the following, seemingly quite compelling requirement on rational belief change:

The recovery principle: If the agent first gives up and then fully regains a belief, then she will also regain all the beliefs that she had before this sequence of loss and regain took place.

But obvious as it may seem, the recovery principle is far from uncontroversial. The following example has been put forward to show that it does not hold in general:

Example 2 [10] First I believed that Cleopatra had a son, and therefore of course also that she had a child. But then a person whom I rely on told me that the book in which I learned this was a historical novel. Accepting this, I gave up my belief that Cleopatra had a child. But soon afterwards I heard a highly respected scholar mention in passing that Cleopatra was a mother. Then I again believed that Cleopatra had a child. However, I did not regain my previous belief that she had a son.

The pattern exhibited in this example contradicts the principle of recovery. After first losing and then taking back my belief that Cleopatra had a child, I still lack one of the beliefs that I had originally, namely that she had a son. There has been a considerable debate on whether examples like this disprove the recovery principle or they are so untypical that the principle is still a useful idealization [8, 15].

Our next example concerns the acquisition of new beliefs:

Example 3
LOGICIAN: Yesterday you told me that you had no idea whatsoever whether Mohammed has any children. Now you profess to be firmly convinced that he is a father. What has happened?
JESSICA: Yesterday evening I saw him in the supermarket with baby food in his shopping cart.

LOGICIAN: I don't see how that can prove him to be a father. He could for instance be shopping for a friend or a relative. Is this all the new information that made you change your mind?
JESSICA: Yes it is.
LOGICIAN: Is there any way in which you can logically derive that he is a father from the fact that he bought baby food, perhaps in combination with something else you knew before?
JESSICA: No, I don't think so.
LOGICIAN: You disappoint me. You received new information that does not contradict what you believed before. Then you can conclude whatever follows logically from the new information in combination with your previous beliefs. But that is all. If you go beyond that, how can you be trusted as a rational thinker?
JESSICA: I'm sorry if I disappoint you, but given what I saw, it seems so plausible that he is a father that I can't help believing it.

Our Logician advances the following guideline for rational belief change:

> *The principle of deductivism:* If the agent adopts a new belief that does not contradict her previous beliefs, then she comes to believe in everything that follows logically from the combination of her old beliefs and the new belief, but nothing beyond that.

As should be clear from the example, this is a contestable principle since it requires, essentially, that we refrain from making any non-deductive inferences.

The theory of belief change is concerned with how human beings change their beliefs. Since our brains (and minds) are finite, we should expect them only to have room for a finite number of beliefs. However, since our language is unlimited, we have to be careful about how we express the requirement of finitude. For instance, I believe in each of the sentences on the following infinite list:

Beethoven completed less than 10 symphonies.
Beethoven completed less than 11 symphonies.
Beethoven completed less than 12 symphonies.
...
Beethoven completed less than 1.000 symphonies.
...

Each of these sentences differs in meaning from all the others, so it follows from my assent to all of them that my set of beliefs contains infinitely many sentences that are unique in terms of meaning. But obviously, this does not make my set of beliefs infinite in any interesting way. All of these sentences follow from the first. Instead of requiring that the set of beliefs be finite we should require it to be *finite-based*, i.e. everything that it contains should follow logically from some finite set of beliefs.

Finite-basedness is of course a static requirement. The corresponding dynamic requirement is that finite-basedness should be preserved under belief change. This should apply both when we add a new belief and when we remove an old one. Thus the following two principles should both apply:

The principle of finite-based contraction: If an agent with a finite-based set of beliefs gives up one of these beliefs, then her new set of beliefs is also finite-based.

The principle of finite-based revision: If an agent with a finite-based set of beliefs adopts a single new belief, then her new set of beliefs is also finite-based.

We now have four principles of rationality for belief change that are all expressible in natural language: Recovery, Deductivism, Finite-based contraction, and Finite-based revision. In the remainder of this chapter, we will introduce some formal approaches to belief change and use these four postulates to compare their properties.

20.2 A Basic Framework for Belief Change

To begin with, we need a general framework in which we can express different approaches to belief change. Following the well-established tradition in the field, we will assume that the agent's (static) belief state at each point in time is represented by a set of sentences, called the *belief set* and denoted K. The belief set is logically closed (closed under logical consequence), by which is meant that it contains everything that it logically implies. This can be expressed with a consequence relation Cn, such that for any set X of sentences, $\mathrm{Cn}(X)$ is the set of its logical consequences. Our requirement that K is logically closed can then be expressed with the simple formula $K = \mathrm{Cn}(K)$.

Logical closure is an idealization, and obviously not a realistic property of a person's set of beliefs. No one has sufficient logical and mathematical competence to believe in everything that follows logically from what she believes. However, as an idealization it is quite useful, since it allows us to work with much simpler formal models than what we would otherwise have needed.

In the logic of belief change, all changes result from inputs. An input usually consists in a sentence and an instruction saying what to do with that sentence. Standardly there are three such instructions, namely "remove this sentence", "add this sentence", and "add this sentence and retain consistency".

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:

$$K \div p \subseteq K \text{ (inclusion)}$$

$$K \div p = \mathrm{Cn}(K \div p) \text{ (closure)}$$

$$p \notin K \div p, \text{ unless } p \text{ is a logical truth (success)}$$

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

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

Expansion has the virtue of simplicity, but it also has the damaging property of bringing us to inconsistency whenever we assimilate 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 assumed to have the following properties:

$K * p = \text{Cn}(K * p)$ (closure)

$p \in K * p$ (success)

$K * p$ is consistent if p is consistent (consistency)

We now have sufficient notation to express the four conditions from Sect. 20.1 in formal language:

$K \subseteq K \div p + p$ (Recovery)
If $\neg p \notin K$ then $K * p = K + p$ (Deductivism)
If K is finite-based, then so is $K \div p$ (Finite-based contraction)
If K is finite-based, then so is $K * p$ (Finite-based revision)

Let us now turn to the construction of belief change operations. Expansion is set-theoretically defined, but how should contraction and revision be constructed? Beginning with contraction, there are many ways to remove a sentence p from a belief set K that contains it. Obviously, p itself has to be thrown out, but that is not enough. We also have to make sure that we do not preserve elements of K that together imply p. For instance, if $q \in K$, then we also have $q \rightarrow p \in K$ (since $p \in K$ and K is logically closed). Since $q \in K$ and $q \rightarrow p \in K$ together imply p, at least one of them has to go in the construction of $K \div p$.

Similar considerations apply to revision. If $\neg p$ is in K, then it has to be removed in the construction of $K * p$ in order to achieve consistency. For every sentence q in K we also have $q \rightarrow \neg p$ in K, and either q or $q \rightarrow \neg p$ has to be removed in the construction of $K * p$.

Thus, both contraction and revision involve choices. We can frame these choices in various ways: as a choice which sentences to retain, a choice which sentences to remove, a choice among possible outcomes, etc. In the next four sections we will investigate four alternative frames for the choices involved in belief change. These frames turn out to induce different properties in the operations, as we will see by testing them against our four postulates from Sect. 20.1.

20.3 AGM: The Standard Approach

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. 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. A set X is an inclusion-maximal p-excluding subset of K (in short: a p-remainder of K) if and only if it is a subset of K that does not imply p, but if we extend it with any additional element from K, then it will imply p. The set of p-remainders of K is denoted $K \perp p$.

Intuitively, we want to retain as much of K as we can without obtaining p. That 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 should be those that are held in all these top remainders. Formally, this is achieved by introducing a selection function γ that selects a subset $\gamma(K \perp p)$ of $K \perp p$ that consists, intuitively speaking, of the "best" elements of $K \perp p$. 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 x-remainders for some sentence x. 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. A sentential operation on a belief set K is a partial meet contraction if and only if it satisfies the following six axioms:

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

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

$(K \div p) \cap (K \div q) \subseteq K \div (p \& q)$ (conjunctive overlap)
If $p \notin K \div (p \& q)$ then $K \div (p \& q) \subseteq K \div p$ (conjunctive inclusion)

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, all we have to do to make p consistently addable is to first remove $\neg p$. This line of reasoning (which can also be found in earlier work by Isaac Levi) gives rise to the following construction of revision in terms of contraction and expansion:

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

It turns out that if revision is defined in this way, then the contraction operation on which the revision operation $*$ was based can be regained as follows:

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

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

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

Let us now return to the four principles that we introduced in Sect. 20.1 and reformulated as formal postulates in Sect. 20.2. As we saw above, Recovery is among the postulates for contraction, so it is satisfied. Deductivism is also satisfied; it follows directly from two of the revision postulates (inclusion and vacuity). However, both the remaining two postulates, Finite-based contraction and Finite-based revision, fail for the AGM operations.[1]

[1]See [12] for a proof that Finite-based contraction does not hold. The proof that Finite-based revision does not hold has not been published, and is therefore given here: Let S be an infinite set of logical atoms in the language, let p be another such atom, and let $K = \text{Cn}(\{\neg p\})$. Then $\{\neg p \vee s \mid s \in S\}$ is a subset of K that does not imply $\neg p$. It follows from compactness and the axiom of choice that there is some X such that $\{\neg p \vee s \mid s \in S\} \subseteq X \in K \perp \neg p$ [1]. Let γ be a selection function such that $\gamma(K \perp \neg p) = \{X\}$ and let $*$ be the revision based on γ. Then $\{\neg p \vee s \mid s \in S\} \cup \{p\} \subseteq K * p$, and since $K * p$ is logically closed we have $S \subseteq K * p$ and consequently $K * p$ is not finite-based.

20.4 Choosing Among Possible Worlds

A set of sentences is maximally consistent if and only if it is consistent but there is no sentence in the language that can be added to it without making it inconsistent. Maximally consistent sets are often called possible worlds since their structure is considered suitable for a total description of a state of the world.

A belief set K is compatible with a possible world if and only if nothing in the belief set contradicts it. Due to the special properties of possible worlds, this is equivalent with the requirement that the belief set is a subset of the possible world. It can also be shown that every belief set is equal to the intersection of all possible worlds that includes it. 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 [9]. In Fig. 20.1, think of each point in the square as a possible world. 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. This representation is quite intuitive, once you get accustomed to the fact that a smaller area represents a larger belief set, not the other way around.

Belief revision has a remarkably simple representation in this framework. In our example, consider the revision $K * p$. Its outcome should be 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 of the intersection of the circle and the parabola, i.e. of those of the currently unrejected worlds in which p is true.

But this was a simple case, in which the new information was compatible with what was already believed. What should we do if the parabola (p) and the circle (K) have an empty intersection? Well, since there are no p-worlds in K we will then have to do with 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

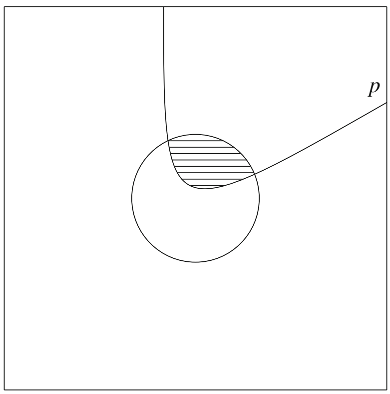

Fig. 20.1 Revision by a sentence p that is compatible with the present belief set

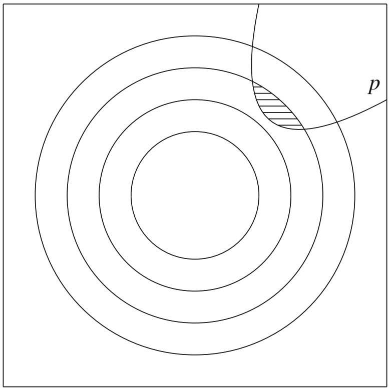

Fig. 20.2 Revision by a sentence *p* that is incompatible with the present belief set

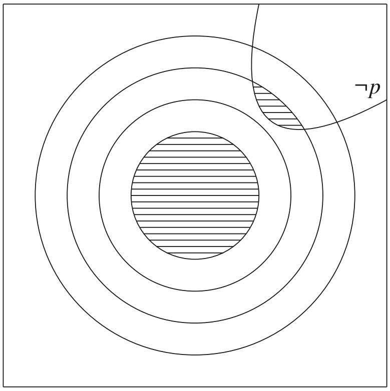

Fig. 20.3 Contraction by *p*

in the next sphere, etc., as in Fig. 20.2. 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. those that are situated in the closest 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 Fig. 20.3.

Although the possible worlds construction is quite different from the partial meet construction, the two ways to construct contraction and revision turn out to yield exactly the same result. 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 Fig. 20.3, and it is a transitively relational partial meet revision if and only if it can be constructed in the way indicated in Fig. 20.2. This surprising result is based on a one-to-one correspondence called "Grove's bijection" between remainders and possible worlds [9], [11, pp. 53–55].

Hence, it makes no difference for the belief change operations if we apply a choice mechanism to remainders or to possible worlds.[2] This might give the impression that it makes no difference what we apply the choice mechanism to. That, however, is an unjustified generalization, as we will see in the following two sections.

20.5 Belief Bases

In an infinite language, belief sets are very large entities. Since they contain all logical consequences of what the agent believes, they contain a lot of sentences that provide logical connections between epistemically unrelated beliefs. This can have rather strange consequences:

> I believe that the earth is (approximately) spherical (e). I also believe that I have my house key in my left trouser pocket (k). Consequently, I believe that the earth is spherical if and only if my house key is in my left pocket ($e \leftrightarrow k$). I put my hand in the pocket to pick up the key. It is not there! I have to give up my belief in k. I cannot then, on pain of inconsistency, retain both my belief in e and my belief in $e \leftrightarrow k$.

Both e and $e \leftrightarrow k$ are elements of the belief set. Therefore, when I find out that k is false, I have to choose between retaining e and retaining $e \leftrightarrow k$. The option of keeping $e \leftrightarrow k$ and giving up e is not excluded automatically, but has to be excluded by the selection mechanism. This appears inappropriate, since $e \leftrightarrow k$ is a merely derived belief that should arguably disappear automatically when k is given up.

Considerations like this led to the construction of belief change operations in which the actual choice takes place among "real" beliefs (like e and k), and the "merely derived" beliefs (like $e \leftrightarrow k$) have no role in the selection process. The crucial construction is a *belief base* consisting of the "real" beliefs, from which the rest of the belief set can be derived. The belief base is denoted B. It satisfies the criterion $\mathrm{Cn}(B) = K$, and contrary to K it does not have to be logically closed. For most purposes we assume that B is finite.

Partial meet contraction and revision can be performed on belief bases in the same way as for belief sets: We define $B \perp p$ as the set of inclusion-maximal subsets of B that do not imply p and γ as a selection function that selects a non-empty subset of each such remainder set. This gives rise to the partial meet contraction $B \div p = \bigcap \gamma(B \perp p)$. Partial meet revision is defined as $(B \div \neg p) \cup \{p\}$.

For any given belief set K, we obtain *base-generated* partial meet contraction and revision (denoted $\widehat{\div}$ and $\widehat{*}$) by assigning to it a belief base B and a selection function for that belief set:

---

[2] AGM is also equivalent to a construction based on selection among the sentences in K, namely epistemic entrenchment [6, 7, 16, 17]. It is also close to equivalent to another such construction, safe contraction [2, 18].

$$K \mathbin{\widehat{\div}} p = \mathrm{Cn}(B \div p)$$

$$K \mathbin{\widehat{*}} p = \mathrm{Cn}(B * p)$$

These operations have been rather carefully investigated and also axiomatically characterized. It is easy to show that base-generated partial meet contraction does not satisfy Recovery. Let p, q, and r be logically independent sentences, and let $K = \mathrm{Cn}(\{p, q, r\})$. We can assign to it a belief base $B = \{p, q, r\}$. Any selection function γ for B will yield $\gamma(B \perp (q \vee r)) = \{\{p\}\}$, thus $K \mathbin{\widehat{\div}} (q \vee r) = \mathrm{Cn}(\{p\})$. It follows that $K \mathbin{\widehat{\div}} (q \vee r) + (q \vee r) = \mathrm{Cn}(\{p, q \vee r\})$, from which we can see that Recovery does not hold.

It is equally easy to show that Deductivism is satisfied. If $\neg p \notin K$ then it holds for any belief base B for K and any partial meet contraction $\div$ on B that $B \div \neg p = B$, and we can conclude that $K \mathbin{\widehat{*}} p = \mathrm{Cn}(B \cup \{p\}) = K + p$.

When we assign a belief base to a finite-based belief set, then we can always choose a finite belief base, and it would indeed be difficult to justify doing otherwise. Provided that we use finite belief bases when that is possible, both Finite-based contraction and Finite-based revision hold for the base-generated partial meet operations. We can conclude that in terms of these four postulates it makes a big difference whether we apply selection functions to the remainders of a belief set or to the remainders of a belief base for it.

20.6 Descriptor Revision

Since we now know that it is important what we apply the choice mechanism to, we have reason to ponder what are the appropriate objects of choice. From the viewpoint of cognitive realism, possible worlds seem unsuitable since they are large structures beyond our grasp. Remainders of belief sets are problematic for much the same reason. Even if the original belief set K is finite-based, if p is a non-tautologous element of K, then the remainder set $K \perp p$ has an infinite number of elements, none of which is finite-based.[3] Remainders of a finite belief base are somewhat more plausible objects of choice, but there is an additional problem with the partial meet construction that applies to belief sets and belief bases alike: It is difficult to justify that we intersect the remainders chosen by the selection function. If the remainders are the top candidates for being the contraction outcome, then their intersection is not one of the top candidates.[4] Then why should it be chosen? [19].

[3]Provided that the language has an infinite number of non-equivalent sentences. For a proof, see [12].

[4]More precisely: if there is more than one top candidate and the top candidates are all p-remainders for some sentence p, then their intersection is not itself a top candidate.

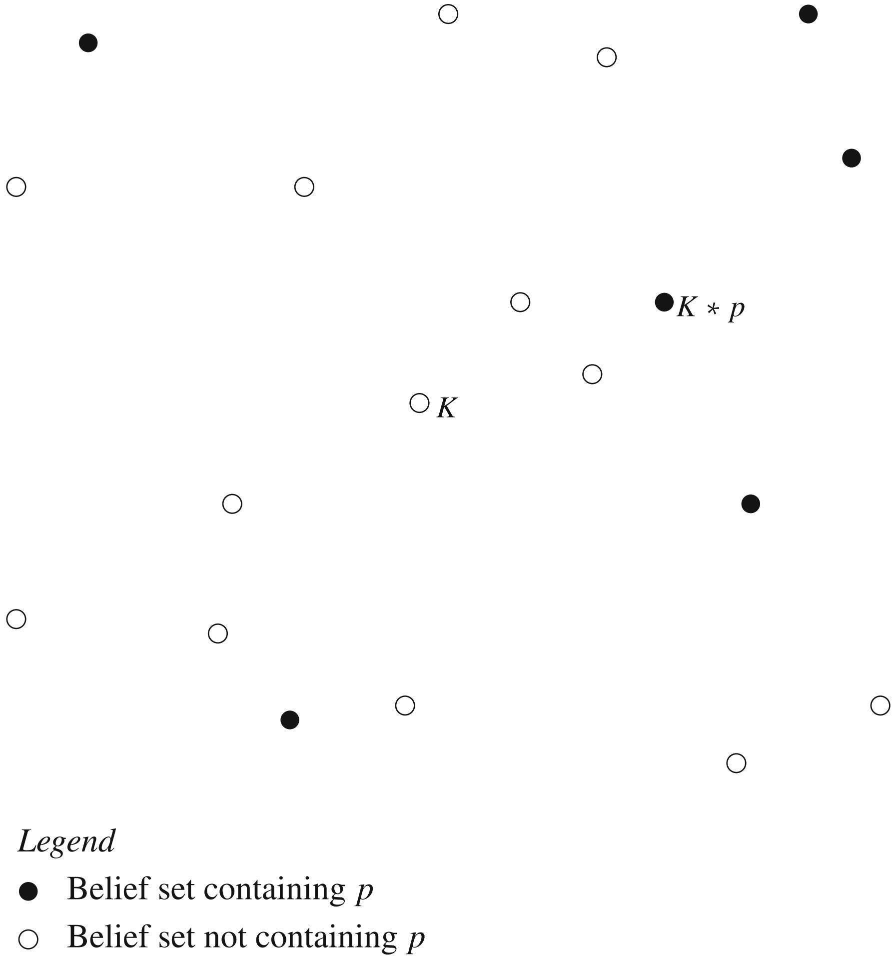

Fig. 20.4 Descriptor revision. $K * p$ is the belief set closest to K among those that contain p. Note that the circles denote belief sets, not possible worlds

Another alternative is to apply the selection mechanism directly to the set of possible outcomes of change. Presumably, not all logically closed sets are suitable contraction outcomes. We can therefore assume that there is an *outcome set* ($\mathbb{X}$) consisting of all the belief sets that can be reached by an operation of change. Intuitively, we can think of $\mathbb{X}$ as consisting of all those belief sets that are coherent, stable, and/or plausible enough to be suitable as outcomes of an operation of belief change. In a cognitively realistic model, all elements of $\mathbb{X}$ should be finite-based.

In addition to this we need a selection mechanism that always singles out exactly one element of the outcome set. One plausible such mechanism is a distance measure. We can think of the elements of the outcome set $\mathbb{X}$ as dispersed in some kind of metric space, as illustrated in Fig. 20.4. $K * p$ is then found as the belief set closest to K in which p holds. (We have to assume that distances are always dissimilar, or that there is some other mechanism to arbitrate ties.)

Contraction can be obtained in the same way: We can identify $K \div p$ as the belief set that is closest to K among those elements of $\mathbb{X}$ that are subsets of K and do not contain p. Furthermore, this model opens up for a more general approach to belief change called *descriptor revision* [13, 14]. To introduce it we need a belief operator $\mathfrak{B}$ such that $\mathfrak{B}p$ denotes that p is believed. We can use this notation to express a wide variety of success conditions for belief change. For instance, $\mathfrak{B}p \vee \mathfrak{B}q$ means that either p or q is believed and $\mathfrak{B}p \vee \neg\mathfrak{B}q$ means that either p is believed or q disbelieved. (The expressions formed with $\mathfrak{B}$ in this way are called belief descriptors.)

Descriptor revision has a single operation of change, denoted $\circ$. It can be applied to all types of descriptors. For instance, in a distance-based framework, the operation $K \circ \neg\mathfrak{B}p$ takes us from K to the belief set closest to K that does not contain p.[5] Similarly, the operation $K \circ (\mathfrak{B}p \vee \mathfrak{B}q)$ produces as outcome the closest belief set in which $\mathfrak{B}p \vee \mathfrak{B}q$ is satisfied, i.e. the closest belief set containing either p or q. Common revision (by sentences) is a special case of descriptor revision, since we can identify $K * p$ with $K \circ \mathfrak{B}p$. Both general descriptor revision ($\circ$) and its restriction to sentential revision ($*$) have been axiomatically characterized. One of the advantages of descriptor revision is that it can easily be extended to iterated change. Hence, in a distance-based framework 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 is a rather straight-forward exercise to show that neither Recovery nor Deductivism holds in this framework. However, provided that the elements of $\mathbb{X}$ are finite-based, both Finite-based contraction and Finite-based revision hold for descriptor revision.

20.7 Conclusion

The outcome of this investigation is summarized in Table 20.1. We have found that the properties of operations of change depend to a large degree on what formal structures we use as objects of choice. Two important philosophical questions are involved here: When we choose rationally what to believe, what are the objects that our choices should be applied to? And what structural properties should a rational agent's choice patterns comply with? Both these are questions that can

[5] In the limiting case when the outcome set contains no element that satisfies the descriptor, nothing is changed, i.e. the original belief set is the outcome of the operation.

Table 20.1 Summary of how the application of the choice mechanism to different objects impacts the satisfaction of four postulates of belief change

Objects of choice	Satisfaction of postulates			
	Recovery	Deductivism	Finite-based contraction	Finite-based revision
Remainders of belief sets	+	+	−	−
Possible worlds	+	+	−	−
Remainders of belief bases	−	+	+	+
Potential outcomes	−	−	+	+

be asked in an informal language, but we need a formal language to perform the logical analysis that shows how closely the two questions are connected with each other.

References and Proposed Readings

1. Alchourrón, C. E., & Makinson, D. (1981). Hierarchies of regulation and their logic. In R. Hilpinen (Ed.). *New studies in deontic logic* (pp. 125–148). Dordrecht: D. Reidel Publishing Company.
2. Alchourrón, C. E., & Makinson, D. (1985). On the logic of theory change: Safe contraction. *Studia Logica, 44*, 405–422
3. *Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic, 50*, 510–530. [The major starting-point of the whole research area.].
4. *Fermé, E., & Hansson, S. O. (2011). AGM 25 years. Twenty-five years of research in belief change. *Journal of Philosophical Logic, 40*, 295–331. [A comparatively brief and non-technical overview of the research area.].
5. *Fermé, E., & Hansson, S. O. (In press). *Belief change. Introduction and overview*. Springer 2018. [A comparatively non-technical overview of the area.].
6. *Gärdenfors, P. (1988). *Knowledge in flux. Modeling the dynamics of epistemic states*. Cambridge: The MIT Press. [A philosophical account of the AGM model by one of its originators.].
7. Gärdenfors, P., & Makinson, D. (1988). Revisions of knowledge systems using epistemic entrenchment. In M. Y. Vardi (Ed.). *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge* (pp. 83–95). Los Altos: Morgan Kaufmann.
8. Glaister, S. M. (2000). Recovery recovered. *Journal of Philosophical Logic, 29*, 171–206.
9. Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic, 17*, 157–170.
10. Hansson, S. O. (1991). Belief contraction without recovery. *Studia Logica, 50*, 251–260.
11. *Hansson, S. O. (1999). *A textbook of belief dynamics. Theory change and database updating*. Dordrecht: Kluwer. [Detailed introduction with a strong focus on AGM and belief base operations. Contains proofs of all the major theorems.].
12. Hansson, S. O. (2008). Specified meet contraction. *Erkenntnis, 69*, 31–54.
13. Hansson, S. O. (2014). Descriptor revision. *Studia Logica, 102*, 955–980.
14. *Hansson, S. O. (2017). *Descriptor revision. Belief change through direct choice.* Cham: Springer. [A full treatment of descriptor revision.].

15. Makinson, D. (1997). On the force of some apparent counterexamples to recovery. In E. G. Valdés, et al. (Eds.). *Normative systems in legal and moral theory*, *Festschrift for Carlos Alchourrón and Eugenio Bulygin* (pp. 475–481). Berlin: Duncker and Humblot.
16. *Rott, H. (2001). *Change, choice and inference: A study of belief revision and nonmonotonic reasoning*. Oxford: Oxford University Press. [An exposition of AGM theory with a strong emphasis on its connections with the study of rational choice.].
17. Rott, H. (2003). Basic entrenchment. *Studia Logica, 73*, 257–280.
18. Rott, H., & Hansson, S. O. (2014). Safe contraction revisited. In S. O. Hansson (Ed.). *David Makinson on classical methods for non-classical problems* (pp. 35–70). Dordrecht: Springer.
19. Sandqvist, T. (2000). On why the best should always meet. *Economics and Philosophy, 16*, 287–313.