Chapter 4
Putting the Building-Blocks Together

The previous two chapters were devoted to negative work. We have inventoried problems and implausible properties that are connected with the traditional approach to belief change. But the purpose of all this negative work was positive. In Section 3.8 we summarized our findings in the form of a list of desiderata for an alternative approach. In this chapter, the outlines of such an approach will be constructed. The rest of the book is devoted to its further development and evaluation.

We will start from scratch. Section 4.1 introduces a very general model for belief change that is based on primitive belief states and inputs, neither of which has any sentential structure. This model has the advantage of making few controversial assumptions but also the disadvantage of low expressive power. It is used as a starting-point to which more structure will successively be added in a guarded fashion, allowing us to see what assumptions are needed to obtain the resulting increase in expressive power. In Section 4.2 sentences are associated with the belief states. In Section 4.3 we introduce descriptors, a versatile tool for expressing properties of belief states, and in Section 4.4 their properties are investigated. In Section 4.5 descriptors are used as a general means for expressing success conditions of operations of change. In Section 4.6 the main features of the resulting model of belief change are summarized. At this point we will have arrived at the fundamental framework for belief change, descriptor revision, that will be further investigated in the rest of the book.

4.1 Beginning Without Sentences

It is almost universally assumed in the belief change literature that beliefs are fully representable as sentences in some language. The totality of beliefs held by an agent is represented by a belief set that is a logically closed set of sentences. Inputs specify a sentence (or sometimes a set of sentences) that has to be either added to the belief set or removed from it. The use of sentences has the immense advantage of making

S.O. Hansson, *Descriptor Revision*, Trends in Logic 46,
DOI 10.1007/978-3-319-53061-1_4

logical treatments possible. Logic operates with sentences, and it is an astoundingly efficient and versatile tool for modelling a wide array of phenomena [100]. However, like other modelling tools it puts emphasis on some aspects of the objects it models at the expense of others. One of the major characteristics of logical models is the linguistic structure that they impose on their subject matter.

Some belief changes can be adequately described in terms of sentences. When I learned that Georg Friedrich Händel wrote the *Messiah* in 1741, the resulting effect on my belief state can be summarized by saying that I started to believe in the sentence "Georg Friedrich Händel wrote the *Messiah* in 1741".[1] However, there are many belief changes that cannot easily be expressed in sentential terms. For instance, when I first heard the *Messiah* I acquired a whole set of new beliefs based on my auditory impressions, namely beliefs about how the music sounds, but I was not able to express all these beliefs in sentences. Similarly, I have beliefs about how Barack Obama's voice sounds, what Picasso's *Guernica* looks like, how my favourite brand of cheese tastes, and how hydrogen sulphide smells. In all these cases my beliefs take the form of "mental pictures" or sensory impressions that can only partially be translated into words. Such perceptually based beliefs are typically adopted "automatically", without any decision or reflection. (See [119], [192, p. 62], and [197, p. 313].) They form a large part of our beliefs. This is one of the reasons why the police use identity parades, photo-lineups, and facial composites in addition to asking witnesses to verbally describe a suspect. A witness may know what a suspect looks like without being able to express this knowledge in words.

Belief change theory is usually assumed to represent changes in the beliefs of individual persons. With this interpretation the exclusion of non-sentential beliefs is a significant limitation. The theory can also refer to belief-holders other than individual persons. In some such cases the sentential format may be less problematic. We can for instance use the theory to model database management. In that case sentential representation is at least in principle fully adequate since the contents of databases are typically representable by sentences. Another example is changes in collectively created and maintained stocks of information or knowledge, such as the corpus of scientific beliefs. Collective information processes are usually based on sentential representations since these are needed for inter-individual communication [112, 125, 136]. However, in order to cover the central case of changes in the beliefs of individual human beings, it is useful to investigate a more general approach that does not require all beliefs to be expressible in sentences. For that purpose we can use a set of primitive belief states, i.e. belief states that are not assumed to have any particular internal structure. Such a belief state may comprise both sentential and non-sentential beliefs. Changes have the effect of taking us from one such belief state to another (or vacuously keeping us in the original one).

[1] And in other sentences containing the same information. The joint information content of sentences with the same meaning is called a proposition. All this could alternatively be expressed in terms of propositions.

Definition 4.1 *A* (deterministic) generic belief state model *is a triple* $\langle \mathbb{K}, \mathbb{I}, \circledcirc \rangle$*, where* $\mathbb{K} = \{\mathcal{K}_1, \mathcal{K}_2, \dots\}$ *is a set of* belief states*,* $\mathbb{I} = \{\imath_1, \imath_2, \dots\}$ *a set of* inputs*, and* $\circledcirc$ *an* input assimilation function[2] *from* $\mathbb{K} \times \mathbb{I}$ *to* $\mathbb{K}$*.*

In such a model all changes are brought about by inputs, and we use the universal operation $\circledcirc$ to express their impact. For each $\mathcal{K} \in \mathbb{K}$ and $\imath \in \mathbb{I}$, $\mathcal{K} \circledcirc \imath$ is the outcome of subjecting $\mathcal{K}$ to the input $\imath$. $\mathcal{K} \circledcirc \imath$ is a new belief state on which further operations can be performed. Therefore this framework allows for iterated change such as $\mathcal{K} \circledcirc \imath_1 \circledcirc \imath_2 \dots \circledcirc \imath_n$ for arbitrary inputs $\imath_1, \imath_2, \dots, \imath_n$.

Although this is a fairly general framework it relies on a couple of assumptions that should be stated. It is *input-assimilating*, by which is meant that all changes stem from an input. Input-assimilating models highlight the causes and mechanisms of change. The inputs are usually interpreted as externally generated, which means that these models contain no representation of internally generated changes such as the loss or deterioration of information or the drawing of new inferences from old information. This can be remedied by allowing for internally generated inputs.[3] Furthermore, inputs come *consecutively*, i.e., one at a time. However, this is not a serious limitation since the set $\mathbb{I}$ of inputs can contain "combined inputs" in the same manner as the inputs of multiple contraction in AGM-style models. (We can define an operation $\div$ such that $K \div \{p, q\}$ has the success condition that neither p nor q should be an element of the outcome.) A much more important limitation is the *lack of explicit representation of time*. It does not seem possible to include a representation of time in an input assimilation model of belief change without making it inordinately complex and unmanageable for most purposes.

The belief state model in Definition 4.1 was called deterministic. That is because the income assimilation function determines for each input exactly what the new belief state will be. In other words, for each $\mathcal{K} \in \mathbb{K}$ and each $\imath \in \mathbb{I}$, we have $\mathcal{K} \circledcirc \imath \in \mathbb{K}$. Another option is to use an input assimilation function that takes us to a non-empty set of belief states. In formal terms we then have a function $\check{\circledcirc}$ such that $\varnothing \neq \mathcal{K} \check{\circledcirc} \imath \subseteq \mathbb{K}$.[4] Such an indeterministic function can be used to reflect that we do not (and perhaps cannot) know exactly what the outcome will be. Alternatively it can signify that the outcome is, in an ontological sense, undetermined. In this book, the focus will be on deterministic models of belief change, but we will return to indeterministic models in Section 5.3.

The structure introduced in Definition 4.1 can also be used to represent an agent's overall state of mind rather than the part of her state of mind that constitutes her belief state. We can for instance conceive the elements of $\mathbb{K}$ as incorporating value judgments, emotions, and desires. Such extensions will not be discussed here, but they can be useful tools for investigating the relationships among these different

---

[2] In the terminology of automata theory it is a transition function.

[3] Changes consisting of the drawing of new inferences from old information have been included in some belief change models; see [83], [91, pp. 20–21], [112, 204].

[4] The symbol ˇ above the symbol representing a (deterministic) belief change operation will be used to denote the indeterministic generalization of that operation.

components of mental states, for instance the effects of changes in belief on value judgments and vice versa.

In the subsequent sections we will add structure to the generic belief state models. But before doing so we will have a look at some interesting properties of these models that can be expressed already with the structure that we have. The following notation is useful:

Definition 4.2 *Let* $\langle \mathbb{K}, \mathbb{I}, \circledcirc \rangle$ *be a generic belief state model and let* $\mathcal{K} \in \mathbb{K}$. *Then:*

(1) $\mathbb{K}_{\mathcal{K}} = \{\mathcal{K} \circledcirc \imath \mid \imath \in \mathbb{I}\}$ *is the set of* directly reachable *belief sets from* $\mathcal{K}$.
(2) $\mathbb{K}^{+}_{\mathcal{K}} = \{\mathcal{K} \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n \mid \{\imath_1, \ldots, \imath_n\} \subseteq \mathbb{I}\}$ *is the set of* indirectly reachable *belief sets from* $\mathcal{K}$.

The following are interesting reachability-related properties of generic belief state models:

$\mathbb{K}_{\mathcal{K}} \neq \{\mathcal{K}\}$ for some $\mathcal{K} \in \mathbb{K}$. (changeability)
$\mathcal{K} \in \mathbb{K}_{\mathcal{K}}$ for all $\mathcal{K} \in \mathbb{K}$. (retainability)
$\mathbb{K}_{\mathcal{K}} = \mathbb{K}$ for all $\mathcal{K} \in \mathbb{K}$. (direct access)
$\mathbb{K}^{+}_{\mathcal{K}} = \mathbb{K}$ for all $\mathcal{K} \in \mathbb{K}$. (successive access)

Retainability can be seen as a technical property; it ensures that the option of changing nothing is represented in the input set. Direct access says that we can go directly (through one single input) from any belief state to any other belief state. This is a problematic property since there seem to be situations where several successive inputs are needed to reach a new belief state. 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 can take her there through a mechanism whereby the earlier of these inputs facilitate her assimilation of those coming later. If that is always possible, then successive access holds.

The following two properties express intuitions that run contrary to those expressed by direct access and successive access:

If $\mathcal{K} \neq \mathcal{K}'$, then $\mathcal{K} \circledcirc \imath \neq \mathcal{K}' \circledcirc \imath'$. (non-convergence)
If $\mathcal{K} \circledcirc \imath_1 \neq \mathcal{K}$, then $\mathcal{K} \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n \neq \mathcal{K}$. (non-reversion)

Observation 4.3 *(1) No belief state model satisfies changeability, successive access, and non-reversion.*

(2) No belief state model with at least two belief states satisfies retainability, direct access, and non-convergence.

(3) If a belief state model satisfies retainability and non-convergence, then it satisfies non-reversion.

Non-convergence and non-reversion are both plausible under the assumption that we "carry our history with us" in the sense that previous beliefs leave traces behind them, for instance in the form of beliefs about what one believed previously.

The following three finiteness properties all refer to the number of alternative belief states that are in some sense available. They are stated here in order of increasing strength.

$\mathbb{K}_{\mathcal{K}}$ is finite for all $\mathcal{K}$. (finite direct access set)
$\mathbb{K}^{+}_{\mathcal{K}}$ is finite for all $\mathcal{K}$. (finite successive access set)
$\mathbb{K}$ is finite. (finite outcome set)

The framework introduced in Definition 4.1 has the important advantage of being general enough to cover a wide range of more specified models of belief revision (and mental dynamics in general) within one and the same formal structure. It can therefore be used to compare different such models. However, no such general investigation of different frameworks will be pursued here. Instead the remainder of this book is primarily devoted to one particularly promising type of model that can be developed within this framework. A couple of comparisons with other models will be made, namely with the AGM model (Sections 8.1, 8.2, and 10.3) and dynamic epistemic logic (Section 7.6).

4.2 Support Functions

With the introduction of generic belief state models we have discarded in one fell swoop all the assumptions about relations between sentence structure and operations of change that were found to be problematic in Chapter 3. But we may have thrown out too much. Actual belief states sustain both beliefs that are expressible in sentences and beliefs that are not. By removing sentences altogether we have deprived ourselves of all means to say something interesting about the special characteristics of the former class of beliefs. In order to regain that capability we will now reintroduce sentences in a cautious manner, avoiding some of the more controversial assumptions of the traditional approach.

The first and crucial step is to assign to each belief state a set consisting of exactly those sentences (in a given language) that represent beliefs held in that state. Formally, this assignment is expressed with a *support function* $\mathfrak{s}$ that takes us from elements of $\mathbb{K}$ to sets of sentences in the object language $\mathcal{L}$.

Definition 4.4 ([85, p. 525]) *Let $\mathbb{K}$ be a set of belief states and $\mathcal{L}$ a language. A* support function *for $\mathbb{K}$ in $\mathcal{L}$ is a function $\mathfrak{s}$ such that $\mathfrak{s}(\mathcal{K}) \subseteq \mathcal{L}$ for all $\mathcal{K} \in \mathbb{K}$.*

In the intended interpretation, $\mathfrak{s}(\mathcal{K})$ is the set of sentences in $\mathcal{L}$ that are supported (believed by the epistemic agent) in the belief state $\mathcal{K}$. Importantly, a support function always refers to a specific language. One and the same belief state $\mathcal{K}$ may be associated with several support functions, $\mathfrak{s}_{\mathcal{L}_1}, \mathfrak{s}_{\mathcal{L}_2}, \ldots$, for different languages. There may also be different support functions referring to different epistemic attitudes that the agent may have to sentences in one and the same language, such as the epistemic attitudes of assuming something, taking it for granted, believing it, and being sure of it. We may for instance distinguish between the set $\mathfrak{s}^{s}_{\mathcal{L}}$ of sentences in $\mathcal{L}$ that the agent is sure of in the state $\mathcal{K}$ and the set $\mathfrak{s}^b_{\mathcal{L}}$ of sentences in the same language that she believes in. In studies comparing different epistemic attitudes it will be useful to have more than one support function. Here the focus will be on a single epistemic attitude, namely that of belief, and a single object language.

In this framework, operations of change are primarily performed on belief states, not on the sets of supported sentences that are associated with them. In other words, we do not apply the input assimilation function $\circledcirc$ to the set $\mathfrak{s}(\mathcal{K})$ of sentences. Instead we apply it to the (non-sentential) belief state $\mathcal{K}$, and then we apply $\mathfrak{s}$ to the outcome $\mathcal{K} \circledcirc \imath$ to obtain the new set of supported sentences, $\mathfrak{s}(\mathcal{K} \circledcirc \imath)$.

The introduction of support functions makes it possible to express a series of important properties of belief change models, such as:

$$\mathfrak{s}(\mathcal{K}) = \text{Cn}(\mathfrak{s}(\mathcal{K})) \qquad \text{(closure)}$$
$$\perp \notin \mathfrak{s}(\mathcal{K}) \qquad \text{(consistency)}$$

As before, Cn is a consequence operation that includes classical truth-functional consequence. In what follows we will assume that closure holds, i.e. that the support function assigns a belief set to each belief state. We will also mostly assume that the assigned belief sets are consistent. However, the presence of inconsistent belief sets may not be as devastating here as it is in frameworks such as the original AGM model where belief changes take place directly on belief sets. In classical truth-functional logic, there is only one logically closed inconsistent set, namely the whole language. Therefore, if K_1 and K_2 are inconsistent belief sets, i.e. $\perp \in K_1$ and $\perp \in K_2$, then $K_1 = K_2$. Since further changes are performed on the belief sets that are now identical, no posterior change can reintroduce the lost distinction.[5] In contrast, the present framework can accommodate distinct inconsistent belief states,[6] i.e. belief states $\mathcal{K}_1$ and $\mathcal{K}_2$ such that $\perp \in \mathfrak{s}(\mathcal{K}_1)$, $\perp \in \mathfrak{s}(\mathcal{K}_2)$, $\mathfrak{s}(\mathcal{K}_1) = \mathfrak{s}(\mathcal{K}_2)$, and $\mathcal{K}_1 \neq \mathcal{K}_2$.[7] Since further changes are performed on $\mathcal{K}_1$ and $\mathcal{K}_2$, not on $\mathfrak{s}(\mathcal{K}_1)$ and $\mathfrak{s}(\mathcal{K}_2)$, distinctions can be reintroduced at a later stage, for instance through revision by some input $\imath$ such that $\mathfrak{s}(\mathcal{K}_1 \circ \imath) \neq \mathfrak{s}(\mathcal{K}_2 \circ \imath)$. This is a property that corresponds to an important feature of actual belief systems, namely that inconsistencies are reparable in a way that does not blur all distinctions.[8]

---

[5] As noted by Hans Rott [223], this problem is not present in extended versions of the AGM model where the outcome of a contraction or revision is not just a belief set but a larger object that contains information about how additional changes will be performed.

[6] More precisely: different belief states that generate inconsistencies in the language of the support function.

[7] The same is true of belief base models in which the belief state is represented by a set of sentences that is not logically closed. Different such belief bases may have the same logical closure and therefore represent belief states with the same belief set [84, 88, 89, 94].

[8] In addition, actual belief systems are capable of containing local inconsistencies that do not corrupt the entire belief system. It is "quite feasible to believe both that Jesus was a human being and that Jesus was not a human being, without believing that the moon is made of cheese" [139, p. 49]. To represent this feature we can employ a support function $\mathfrak{s}$ that does not satisfy closure under classical consequence (but possibly some weaker, paraconsistent closure condition). On local inconsistencies, see [139].

The following properties are related to direct access and successive access that were introduced in Section 4.1.[9]

If $\perp \notin \mathrm{Cn}(\{p\})$, then there is some input $\imath$ with $p \in \mathfrak{s}(\mathcal{K} \circledcirc \imath)$. (direct believability)

If $\perp \notin \mathrm{Cn}(\{p\})$, then there is a series $\imath_1, \ldots, \imath_n$ of inputs with $p \in \mathfrak{s}(\mathcal{K} \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n)$. (successive believability)

If $p \notin \mathrm{Cn}(\varnothing)$, then there is some input $\imath$ with $p \notin \mathfrak{s}(\mathcal{K} \circledcirc \imath)$. (direct removability)

If $p \notin \mathrm{Cn}(\varnothing)$, then there is a series $\imath_1, \ldots, \imath_n$ of inputs with $p \notin \mathfrak{s}(\mathcal{K} \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n)$. (successive removability)

There is some input $\imath$ with $s(\mathcal{K} \circledcirc \imath) = \mathrm{Cn}(\varnothing)$. (direct depletability)

There is a series $\imath_1, \ldots, \imath_n$ of inputs with $s(\mathcal{K} \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n) = \mathrm{Cn}(\varnothing)$. (successive depletability)

These are all fairly strong and arguably problematic properties. As mentioned above in connection with direct access, some persons may have beliefs (such as articles of religious faith) that nothing can make them give up. There may also be potential beliefs that they will never adopt, come whatever may. However, it can be questioned whether such stubbornness is compatible with rationality. The present framework allows for different answers to that question, expressible in terms of whether or not the believability and removability postulates hold. (By way of comparison, the equivalents of direct believability and direct removability hold in the AGM framework.[10])

The following are two finiteness properties that refer to the properties of individual belief states.

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

If $\mathfrak{s}(\mathcal{K})$ is finite-based, then so is $\mathfrak{s}(\mathcal{K} \circledcirc \imath)$. (finite-based outcome)

Finite-based outcome, the weakest of the two, was discussed in Section 3.1, where we found its absence in the AGM framework to be problematic. Interestingly, it is prone to conflict with the finiteness properties introduced in Section 4.1.

Observation 4.5 *Let $\mathfrak{s}$ be a support function for the belief states of some generic belief state model, and let the language $\mathcal{L}$ to which it refers be logically infinite. Then:*

(1) Direct believability, finite-based outcome, and finite direct access set are not all satisfied.

(2) Successive believability, finite-based outcome, and finite successive access set are not all satisfied.

[9] The direct versions of these properties are discussed in [121].

[10] Since $p \in K * p$ and $p \notin (K \div p) \setminus \mathrm{Cn}(\varnothing)$.

Finally, let us introduce properties indicating how much information about the belief state $\mathcal{K}$ is contained in the supported set $\mathfrak{s}(\mathcal{K})$.

$$\text{If } \mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}'), \text{ then } \mathcal{K} = \mathcal{K}'. \qquad \text{(injectivity)}$$

$$\text{If } \mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}'), \text{ then } \mathfrak{s}(\mathcal{K} \circledcirc \imath) = \mathfrak{s}(\mathcal{K}' \circledcirc \imath) \text{ for all } \imath \in \mathbb{I}. \qquad \text{(sententiality)}$$

According to injectivity, any difference between belief states is manifested on the sentential level. For instance, suppose that the only difference between $\mathcal{K}$ and $\mathcal{K} \circledcirc \imath$ is that in the latter you have looked somewhat more closely at your neighbour's hedge, and your mental picture of it has changed in consequence. Injectivity requires that there is some sentence (presumably about the hedge) that you could utter to express your beliefs in one of $\mathcal{K}$ and $\mathcal{K} \circledcirc \imath$ but not in the other. Notably, it does not require that all the differences between the two belief states can be expressed linguistically, only that at least one of them can.

Sententiality is the weaker of the two properties. It says that if two belief states are indistinguishable in terms of what sentences they support, then no series of changes will make their successors distinguishable in that respect.[11] This excludes the existence of belief states that are statically but not dynamically equivalent on the linguistic level, i.e. such that they cannot be distinguished in terms of the beliefs they support, but their successors after some operation(s) of change can be distinguished.[12] It also excludes belief changes that weaken or strengthen beliefs without moving any of them across the belief/non-belief border.[13] But contrary to injectivity, sententiality allows for the existence of essentially non-linguistic properties of belief states that will never show up when the beliefs are expressed linguistically [93].

The plausibility of these properties depends on the language $\mathcal{L}$ that the support function $\mathfrak{s}_{\mathcal{L}}$ operates with. The more expressive power the language has, the less problematic is the assumption that two distinct belief states must have some difference that is expressible in the language.[14] However, this assumption will never be entirely unproblematic since it deprives us of the possibility of distinguishing on the linguistic level between different inconsistent belief states (given that $\mathfrak{s}_{\mathcal{L}}$ satisfies the closure property, $\mathfrak{s}_{\mathcal{L}}(\mathcal{K}) = \text{Cn}(\mathfrak{s}_{\mathcal{L}}(\mathcal{K}))$.)

[11] It can be applied repeatedly, and can therefore equivalently be expressed as follows: If $\mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}')$, then $\mathfrak{s}(\mathcal{K} \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n) = \mathfrak{s}(\mathcal{K}' \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n)$ for all series $\imath_1, \ldots, \imath_n$ of elements of $\mathbb{I}$.

[12] On the difference between static and dynamic equivalence of belief states, see [83].

[13] Suppose that an input $\imath$ (1) strengthens p in $\mathcal{K}$, but (2) does not move any sentence across the belief/non-belief border. It would seem to follow from (1) that there is some series $\imath_1, \ldots, \imath_n$ of inputs such that $p \notin \mathfrak{s}(\mathcal{K} \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n)$ and $p \in \mathfrak{s}(\mathcal{K} \circledcirc \imath \circledcirc \imath_1 \circledcirc \ldots \circledcirc \imath_n)$, but it follows from (2) that $\mathfrak{s}(\mathcal{K} \circledcirc \imath) = \mathfrak{s}(\mathcal{K})$. This contradicts sententiality. On operations that strengthen or weaken beliefs, see [28].

[14] This is particularly pertinent if autoepistemic or conditional beliefs are included in the belief set. See Chapter 7.

4.3 Belief Descriptors

In order to reconstruct a sentential framework we need to represent not only belief states but also inputs in sentential terms. Standard operations of belief change are defined in terms of their success conditions, such as $p \in K * p$ for revision and $p \notin (K \div p) \setminus \text{Cn}(\varnothing)$ for contraction. These are statements about what is believed in the belief state that the operation results in. We need a versatile way to express such properties of belief states.

For that purpose, we will introduce a metalinguistic belief predicate $\mathfrak{B}$. As arguments it takes sentences in the object language in which beliefs are expressed. For any sentence $p \in \mathcal{L}$, the expression $\mathfrak{B}p$ denotes that p is believed in the belief state under consideration. A belief set satisfies $\mathfrak{B}p$ if and only if it has p as an element. An expression like this, consisting of $\mathfrak{B}$ followed by a sentence in the object language, will be called an *atomic belief descriptor*. The term "atomic" signals that these sentences are the smallest building-blocks in the language of belief descriptors that we are now building. However, atomic belief descriptors are not atomic in the sense of being logically independent. For instance, from $\mathfrak{B}p$ and $\mathfrak{B}q$ we can conclude that $\mathfrak{B}(p\&q)$.[15]

Atomic belief descriptors can be combined with the usual truth-functional connectives, classically interpreted. Hence, $\mathfrak{B}p \vee \mathfrak{B}q$ denotes that either p or q is believed, and $\neg\mathfrak{B}r$ that r is not believed. The truth conditions of these expressions follow the standard pattern: $\neg\mathfrak{B}p$ is satisfied whenever $\mathfrak{B}p$ is not satisfied, $\mathfrak{B}p \;\&\; \mathfrak{B}q$ whenever both $\mathfrak{B}p$ and $\mathfrak{B}q$ are satisfied, $\mathfrak{B}p \vee \mathfrak{B}q$ whenever either $\mathfrak{B}p$ or $\mathfrak{B}q$ is satisfied. These composite expressions are called *molecular belief descriptors*.

Finally, we can form sets of (molecular) belief descriptors, such as $\{\mathfrak{B}(p \vee q), \neg\mathfrak{B}p, \neg\mathfrak{B}q\}$. Sets of molecular belief descriptors will be called *composite belief descriptors* or in short just *descriptors*. A composite belief descriptor is satisfied if and only if all its elements are satisfied. Hence the descriptor $\{\mathfrak{B}(p \vee q), \neg\mathfrak{B}p, \neg\mathfrak{B}q\}$ is satisfied by the belief set $\text{Cn}(\{p \vee q\})$ but not by the belief set $\text{Cn}(\{q\})$.[16]

A descriptor is (obviously) called finite if it has a finite number of elements. Strictly speaking, finite descriptors are superfluous since they can be replaced by the conjunction of their elements. For instance, $\{\mathfrak{B}(p \vee q), \neg\mathfrak{B}p, \neg\mathfrak{B}q\}$ is satisfied by exactly the same belief sets that are satisfied by the molecular descriptor $\mathfrak{B}(p \vee q)$ $\& \neg\mathfrak{B}p \;\&\; \neg\mathfrak{B}q$. However, the set-theoretical notation is often more convenient, and it will be used freely in what follows.

Upper-case Greek letters such as $\Psi, \Xi, \ldots$ will be used to denote (composite) descriptors. Occasionally, when a notation is needed for molecular descriptors, lower-case Greek letters such as $\alpha, \beta, \ldots$ will be used for that purpose.

[15] Frank Ramsey noted in 1925 that "A believes p" is not a truth function of p but can instead be treated as "one of other atomic propositions". [210, p. 9n].

[16] Composite descriptors with one element will be used interchangeably with the molecular descriptor that they contain. For instance, $\{\mathfrak{B}p\}$ and $\mathfrak{B}p$ will be used interchangeably.

All this is important enough to be summarized in a formal definition:

Definition 4.6 ([124]) *An* atomic belief descriptor *is a sentence $\mathfrak{B}p$ with $p \in \mathcal{L}$. It is* satisfied *by a belief state $\mathcal{K}$ according to a support function $\mathfrak{s}$ in $\mathcal{L}$ if and only if $p \in \mathfrak{s}(\mathcal{K})$.*

A molecular belief descriptor *(denoted by lower-case Greek letters α, $\beta, \ldots$) is a truth-functional combination of atomic descriptors. Conditions of satisfaction are defined inductively, such that $\mathcal{K}$ satisfies $\neg\alpha$ according to $\mathfrak{s}$ 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 Ψ, $\Xi, \ldots$) is a non-empty set of molecular descriptors. A belief state $\mathcal{K}$ satisfies a composite descriptor Ψ according to $\mathfrak{s}$ if and only if it satisfies all its elements.*

A descriptor is satisfiable within *a set of belief states if and only if it is satisfied by at least one of its elements.*

As defined here, the symbol $\mathfrak{B}$ is not part of the object language, and therefore it cannot be used to express an agent's beliefs about her own beliefs. (It is possible to include an autoepistemic belief predicate into the language. It may or may not coincide with $\mathfrak{B}$, depending on whether the agent's autoepistemic beliefs accord with her epistemic conduct. See Section 7.1.) It should also be noted that our definition does not allow $\mathfrak{B}$ to be iterated.[17] Therefore expressions such as $\mathfrak{B}\mathfrak{B}p$ or $\mathfrak{B}(\mathfrak{B}p \rightarrow \mathfrak{B}q)$ are not well-formed. The reason for this is that it is very unclear what such expressions could possibly mean, given the metalinguistic interpretation of $\mathfrak{B}$.

Descriptors are well suited to express the success conditions of different types of belief change operations. In revision, a specified sentence p should be included in the outcome, in other words the success condition has the characteristic form $\mathfrak{B}p$. In contraction, a specified sentence p is instead required not to be present in the outcome, thus a success condition of the form $\neg\mathfrak{B}p$ has to be satisfied. The success conditions of many other, less common, types of operations can be expressed analogously. Multiple revision by a set $\{p_1, \ldots, p_n\}$ of sentences has two variants, package revision that requires all of them to be believed in the new belief state, and choice revision that only requires that at least one of them be believed.[18] Package revision has the success condition $\{\mathfrak{B}p_1, \ldots, \mathfrak{B}p_n\}$, and choice revision the success condition $\mathfrak{B}p_1 \vee \ldots \vee \mathfrak{B}p_n$. Similarly, multiple contraction by a set $\{p_1, \ldots, p_n\}$ of sentences has two variants, package contraction that requires all of them to be removed and choice contraction that only requires that at least one of them be removed [64]. Package contraction has to satisfy the success condition $\{\neg\mathfrak{B}p_1, \ldots, \neg\mathfrak{B}p_n\}$, and choice contraction the success condition $\neg\mathfrak{B}p_1 \vee \ldots \vee \neg\mathfrak{B}p_n$. The operation of replacement is constructed to remove one specified sentence and incorporate another [110].

[17] More precisely: It does not allow the formation of expressions in which an instance of $\mathfrak{B}$ appears within the scope of another instance of $\mathfrak{B}$.

[18] This terminology is used in [107] and [239, p. 280]. It is based on the terminology for two types of multiple contraction used in [64]. Hans Rott uses the terms "bunch revision" and "pick revision" for the same concepts [217, p. 65].

Its success condition has the form $\{\neg\mathfrak{B}p, \mathfrak{B}q\}$. Finally, the operation of "making up one's mind" aims at either belief or disbelief in a specified sentence p. Its success condition is $\mathfrak{B}p \vee \mathfrak{B}\neg p$ [264]. In summary, descriptors can be used to express a wide range of success conditions in a precise and unified way. We will use this locution to construct a uniform type of belief change that covers all operations whose success conditions are expressible with descriptors. But before that we need to have a brief look at some of the formal properties of descriptors.

4.4 Properties of Descriptors

Descriptors refer to what sentences a belief state supports, i.e. to the contents of the belief set $\mathfrak{s}(\mathcal{K})$ supported by a belief state $\mathcal{K}$. We can therefore assume that if $\mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}')$, then $\mathcal{K}$ and $\mathcal{K}'$ satisfy the same descriptors. For simplicity, we can then refer to descriptors as satisfied by belief sets rather than by belief states. The symbol $\Vdash$ will be used for that relation of satisfaction:

Definition 4.7 ([124]) *Let K be a belief set and let Ψ and Ξ be descriptors.*

$K \Vdash \Psi$ means that K satisfies Ψ, and $\Psi \Vdash \Xi$ that all belief sets satisfying Ψ also satisfy Ξ.

The corresponding equivalence relation is written $\dashv\Vdash$; hence $\Psi \dashv\Vdash \Xi$ holds if and only if both $\Psi \Vdash \Xi$ and $\Xi \Vdash \Psi$ hold.

As can be seen from the definition, $\Vdash$ is (for simplicity) used to denote two binary relations. First, it stands for a relation between belief sets and descriptors, such that $K \Vdash \Psi$ holds if and only if K satisfies Ψ (in the sense of satisfaction specified in Definition 4.6, which means that it has to satisfy all elements of Ψ). Secondly, $\Vdash$ also represents a relation between descriptors, such that $\Psi \Vdash \Xi$ holds if and only if it holds for all belief sets K that if $K \Vdash \Psi$ then $K \Vdash \Xi$.

The following observation summarizes some elementary properties of descriptors:

Observation 4.8 *(1) Let K be a belief set and α a molecular descriptor. Then either $K \Vdash \alpha$ or $K \Vdash \neg\alpha$.*

(2) Let $\alpha_1, \ldots, \alpha_n$ be molecular descriptors. Then $\{\alpha_1, \ldots, \alpha_n\} \dashv\Vdash \{\alpha_1 \& \ldots \& \alpha_n\}$.

(3) For any descriptors Ψ and Ξ: $\Psi \Vdash \Xi$ if and only if there is some Ψ' such that $\Xi \subseteq \Psi'$ and $\Psi' \dashv\Vdash \Psi$.

Part (1) of the observation cannot be extended to formulas in which the belief set K has been replaced by a descriptor. It does not hold in general that if Ψ is a descriptor and α a molecular descriptor, then either $\Psi \Vdash \alpha$ or $\Psi \Vdash \neg\alpha$.[19]

A descriptor can be inconsistent in the sense that no belief set can satisfy it. The following notation is introduced to express such inconsistency:

[19] To see that, let p and q be logically independent elements of $\mathcal{L}$, and let $\Psi = \{\mathfrak{B}p\}$ and $\alpha = \mathfrak{B}q$.

Definition 4.9 ([124]) $\underline{\perp}$ *(descriptor falsum) denotes* $\{\mathfrak{B}p, \neg\mathfrak{B}p\}$ *for an arbitrary* p.

It is important to distinguish $\underline{\perp}$ from the falsum $\perp$ of the object language (that is introducable as $p\&\neg p$ for an arbitrary p). The inconsistent belief set $K = \text{Cn}(\{\perp\})$ satisfies the condition $K \vdash \perp$, but no belief set satisfies the condition $K \Vdash \underline{\perp}$.

We can apply ordinary conjunction and disjunction to molecular descriptors, forming sentences such as $\alpha\&\beta$ and $\alpha \vee \beta$. For composite descriptors, we can use set union with essentially the same effect as conjunction. The parallel is obvious: a belief set satisfies $\alpha\&\beta$ if and only if it satisfies α and it also satisfies β. Similarly, it satisfies $\Psi \cup \Xi$ if and only if it satisfies Ψ and it also satisfies Ξ. For disjunction, the following construction can be used:

Definition 4.10 ([126]) *The* descriptor disjunction $\underline{\vee}$ *is defined by the relationship* $\Psi \underline{\vee} \Xi = \{\alpha \vee \beta \mid \alpha \in \Psi$ *and* $\beta \in \Xi\}$.

Observation 4.11 *Let* K *be a belief set and let* Ψ *and* Ξ *be descriptors. Then:*
$K \Vdash \Psi \underline{\vee} \Xi$ *if and only if either* $K \Vdash \Psi$ *or* $K \Vdash \Xi$.

It follows from Definition 4.6 that the negation of a molecular descriptor α is a descriptor $\neg\alpha$ such that for any belief set X: $X \Vdash \neg\alpha$ if and only if $X \nVdash \alpha$. A generalization of negation to composite descriptors should have the same property, in other words the negation of a composite descriptor Ψ would have to be another descriptor $\underline{\neg}\,\Psi$ such that for any belief set X: $X \Vdash \underline{\neg}\,\Psi$ if and only if $X \nVdash \Psi$. For any finite descriptor $\{\alpha_1, \ldots, \alpha_n\}$ we can use the set:

$$\underline{\neg}\{\alpha_1, \ldots, \alpha_n\} = \{\neg\alpha_1 \vee \ldots \vee \neg\alpha_n\}$$

as its negation. However, as the following observation shows, there are infinite descriptors for which no construction with the desired property is possible. In other words, there are non-negatable descriptors.

Observation 4.12 *Let the object language* $\mathcal{L}$ *have infinitely many logically independent atoms. Then there are non-negatable descriptors, i.e. descriptors* Ψ *such that there is no descriptor* $\underline{\neg}\Psi$ *satisfying the condition that for any belief set* X*:* $X \Vdash \underline{\neg}\,\Psi$ *if and only if* $X \nVdash \Psi$.

To each descriptor Ψ we can assign a *characteristic set* of belief sets, namely the set consisting of those belief sets that satisfy the descriptor. The descriptors that characterize a single belief set are worth special attention since they are very useful in formal proofs.

Definition 4.13 ([126]) *A descriptor* Ψ *is* maxispecified *(maximally specified) if and only if there is exactly one belief set* Y *in* $\wp(\mathcal{L})$ *such that* $Y \Vdash \Psi$. *It is then a maxispecified descriptor* for Y.

There are many (equivalent) maxispecified descriptors for each belief set. For instance, if $X = \text{Cn}(\{q\})$ then both $\{\mathfrak{B}q\} \cup \{\neg\mathfrak{B}x \mid x \notin X\}$ and $\{\mathfrak{B}x \mid x \in X\} \cup \{\neg\mathfrak{B}x \mid x \notin X\}$ are maxispecified descriptors for X. For convenience, one of the maxispecified descriptors for a belief set X will be denoted as follows:

Definition 4.14 ([126]) *Let X be a belief set. Then Π_X is the maxispecified descriptor for X such that:*

$\Pi_X = \{\mathfrak{B}x \mid x \in X\} \cup \{\neg\mathfrak{B}x \mid x \notin X\}$

Whereas all (single) belief sets can be characterized by a descriptor, there are sets of belief sets that cannot:

Definition 4.15 *A set $\mathbb{Y}$ of belief sets is* descriptor-definable *if and only if there is some descriptor Ψ such that for all belief sets Y:*

$Y \in \mathbb{Y}$ if and only if $Y \Vdash \Psi$.

Observation 4.16 ([124]) *(1) Let $\mathbb{Y}$ be a finite set of belief sets. Then $\mathbb{Y}$ is descriptor-definable.*

(2) If $\mathcal{L}$ is logically infinite[20] *then there are sets of belief sets that are not descriptor-definable.*

We now have the formal means to analyze an issue that was brought up informally in Section 2.4, namely which success conditions are preserved under intersection. We noted that if each element of a set of belief sets satisfies the success condition for revision by the sentence p (i.e. they all contain p), then their intersection also satisfies that condition (i.e. the intersection contains p). Similarly, if all belief sets in a collection satisfy the success condition for contraction by the sentence p (i.e. none of them contains p), then their intersection does the same (i.e. it does not contain p). This is what makes the select-and-intersect method viable for both revision and contraction. But not all success conditions are preserved under intersection. Since success conditions can be represented by descriptors we can now express this condition in a fully formalized way:

Definition 4.17 *A descriptor Ψ is* preserved under intersection *if and only if it holds for all sets $\mathbb{Y}$ of belief sets that if $Y \Vdash \Psi$ for all $Y \in \mathbb{Y}$, then $\bigcap \mathbb{Y} \Vdash \Psi$.*

The following observation identifies an important class of descriptors that are preserved under intersection.

Observation 4.18 ([135][21]) *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) $\mathfrak{B}p_1 \vee \ldots \vee \mathfrak{B}p_n \vee \neg\mathfrak{B}q$, *with* $q \vdash p_1 \vee \ldots \vee p_n \rightarrow p_k$ *for some* p_k.

---

[20] A set of sentences is logically infinite if and only if it has infinitely many equivalence classes in terms of logical equivalence. Cf. Section 2.5.

[21] This observation is related to the well-known theorem that a theory is equivalent to a Horn theory if and only if the set of its models is closed under intersection. This was proved (in a generalized form) in [187]. A more accessible proof can be found in [38, pp. 254–257], and an excellent introduction to Horn clauses in [148].

4.5 Descriptor Revision Introduced

We are now ready for the final step in the construction of a new framework for belief change: the transition from primitive inputs, i.e. elements of $\mathbb{I}$, to changes based on success conditions, expressed with descriptors. We are looking for a way to revise a belief state $\mathcal{K}$ by a descriptor Ψ rather than by an element of $\mathbb{I}$. This means that we need an operation $\circ$ of belief change that takes descriptors as inputs. Such an operation will supersede and unify the traditional operations, thus $\mathcal{K} \circ \mathfrak{B}p$ takes the role of revision, $\mathcal{K} \circ \neg\mathfrak{B}p$ that of contraction, $\mathcal{K} \circ \{\neg\mathfrak{B}p_1, \dots, \neg\mathfrak{B}p_n\}$ that of multiple (package) contraction, etc.

Importantly, the use of descriptors instead of elements of $\mathbb{I}$ as inputs does not require the introduction of new outputs. We can assume that $\mathbb{I}$ is exhaustive in the sense that every new belief state that is directly reachable from $\mathcal{K}$ can be reached through revision by one of the inputs in $\mathbb{I}$. This means that for every descriptor Ψ we can identify $\mathcal{K} \circ \Psi$ with a belief state $\mathcal{K} \circledcirc \imath$ for some $\imath \in \mathbb{I}$. In other words, we should have $\mathcal{K} \circ \Psi \in \mathbb{K}_{\mathcal{K}}$. Furthermore, since the operation $\circ$ should be successful, the outcome $\mathcal{K} \circ \Psi$ should satisfy Ψ, i.e. we should have $\mathfrak{s}_{\mathcal{L}}(\mathcal{K} \circledcirc \imath) \Vdash \Psi$ (unless, of course, there is no set $\mathcal{K} \circledcirc \imath$ with this property). Combining the two conditions, we obtain:

$$\mathcal{K} \circ \Psi \in \{\mathcal{K}' \in \mathbb{K}_{\mathcal{K}} \mid \mathfrak{s}_{\mathcal{L}}(\mathcal{K}') \Vdash \Psi\} \text{ if } \Psi \text{ is satisfiable within } \mathbb{K}_{\mathcal{K}}.$$

In order to construct such an operation we need to select an element of $\{\mathcal{K}' \in \mathbb{K}_{\mathcal{K}} \mid \mathfrak{s}_{\mathcal{L}}(\mathcal{K}') \Vdash \Psi\}$. Typically that set will have more than one element, i.e. there will be more than one element of $\mathbb{K}_{\mathcal{K}}$ that satisfies 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. 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 expect $K \circ \Psi$ to have as few features as possible that are not shared by all the reasonably credible revision outcomes that satisfy Ψ. The crucial assumption that we have to make when modelling deterministic belief change is that one among the various potential outcomes satisfying Ψ is singled out to be *the* outcome of revision *by* Ψ. In the formal language, this singling out is most conveniently represented by a choice function that extracts *only one* element from the set it is applied to. (In indeterministic belief change, we instead have an operation $\check{\circ}$ such that $K \check{\circ} \Psi$ is a non-empty set of belief sets, each of which is equal to $\mathfrak{s}(\mathcal{K}')$ for some $\mathcal{K}' \in \mathbb{K}_{\mathcal{K}}$ with $\mathfrak{s}(\mathcal{K}') \Vdash \Psi$.)

A function that singles out a single element can be constructed as a special case of the definition of a choice function. Then the formal object that we obtain will be a set with the chosen belief state as its only element. Alternatively we can construct a function that directly delivers this belief state (instead of a set in which it is the only element). It does not make much of a difference which of these two formal constructions we employ. They are both introduced in the following definition:

Definition 4.19 ([120]) *Let $\mathbb{Y}$ be a set. A* monoselective choice function *for $\mathbb{Y}$ is a choice function C for $\mathbb{Y}$ such that if $\varnothing \subset \mathbb{Y}' \subseteq \mathbb{Y}$ then $C(\mathbb{Y}')$ has exactly one element. Alternatively it can be represented by a function $\widehat{C}$ such that $\widehat{C}(\mathbb{Y}') \in \mathbb{Y}'$ whenever $\varnothing \subset \mathbb{Y}' \subseteq \mathbb{Y}$, and otherwise $\widehat{C}(\mathbb{Y}')$ is undefined.*

We will apply monoselective choice functions to a predetermined set of potential outcomes, namely the set $\mathbb{K}_{\mathcal{K}}$ of belief states that are directly reachable from $\mathcal{K}$. In this way the select-and-intersect method is replaced by a direct choice among the potential outcomes. We can use this method to construct our first version of descriptor revision:

Definition 4.20 *Let $\mathbb{K}$ be a set of belief states, $\mathbb{I}$ a set of inputs, $\circledcirc$ an input assimilation function on $\mathbb{K} \times \mathbb{I}$, $\mathfrak{s}$ a support function for $\mathbb{K}$ in a language $\mathcal{L}$, and $\widehat{C}$ a monoselective choice function for $\mathbb{K}$. The* (deterministic) descriptor revision[22] *based on $\langle \mathbb{K}, \mathbb{I}, \circledcirc, \mathfrak{s}, \widehat{C} \rangle$ is the operation $\circ$ such that for all $\mathcal{K} \in \mathbb{K}$ and all descriptors Ψ for the language $\mathcal{L}$:*

(i) If Ψ is satisfiable within $\mathbb{K}_{\mathcal{K}}$, then $\mathcal{K} \circ \Psi = \widehat{C}(\{\mathcal{K}' \in \mathbb{K}_{\mathcal{K}} \mid \mathfrak{s}(\mathcal{K}') \Vdash \Psi\})$, and

(ii) otherwise $\mathcal{K} \circ \Psi = \mathcal{K}$.

This definition introduces a uniformity property for descriptor revision. If it holds for a belief state $\mathcal{K}$ and two descriptors Ψ_1 and Ψ_2 that

$$\{\mathcal{K}' \in \mathbb{K}_{\mathcal{K}} \mid \mathfrak{s}(\mathcal{K}') \Vdash \Psi_1\} = \{\mathcal{K}' \in \mathbb{K}_{\mathcal{K}} \mid \mathfrak{s}(\mathcal{K}') \Vdash \Psi_2\},$$

then $\mathcal{K} \circ \Psi_1 = \mathcal{K} \circ \Psi_2$. To see why this is a plausible principle, it may be helpful to consider the special case when there are sentences p and q such that $\Psi_1 = \{\mathfrak{B}p\}$ and $\Psi_2 = \{\mathfrak{B}q\}$. It then follows from

$$\{\mathcal{K}' \in \mathbb{K}_{\mathcal{K}} \mid \mathfrak{s}(\mathcal{K}') \Vdash \mathfrak{B}p\} = \{\mathcal{K}' \in \mathbb{K}_{\mathcal{K}} \mid \mathfrak{s}(\mathcal{K}') \Vdash \mathfrak{B}q\}$$

that exactly those belief changes that make the agent believe in p will also make her believe in q, and vice versa. Therefore, making her believe in p and making her believe in q seems to be essentially the same thing.

Definition 4.20 provides the most general form of (deterministic) descriptor revision. We will now introduce two useful simplifications of the model, both of which were anticipated in the previous sections of this chapter. First, we can assume that the set of reachable belief states is the same irrespective of what belief state we begin with, i.e. that $\mathbb{K}_{\mathcal{K}} = \mathbb{K}_{\mathcal{K}'}$ for all $\mathcal{K}, \mathcal{K}' \in \mathbb{K}$. Since we do not need to consider belief states that are not reachable from anywhere, this is equivalent to adopting the postulate of Direct access, i.e. $\mathbb{K}_{\mathcal{K}} = \mathbb{K}$. This allows us to make a small but important modification of clause (i) in Definition 4.20:

(i_L) If Ψ is satisfiable within $\mathbb{K}$, then $\mathcal{K} \circ \Psi = \widehat{C}(\{\mathcal{K}' \in \mathbb{K} \mid \mathfrak{s}(\mathcal{K}') \Vdash \Psi\})$.

---

[22] The term “descriptor revision” refers to operations that take descriptors as inputs. For clarity, the operations called “revision” in the traditional approach will be called “sentential revision”.

The index of *(i_L)* stands for "local". Obviously, replacing *(i)* by *(i_L)* makes no difference in studies of local change, i.e. one-step changes that all start with the same belief state.

The second simplification is somewhat more far-reaching. It consists in adopting the principle of injectivity from Section 4.2. (If $\mathfrak{s}(\mathcal{K}) = \mathfrak{s}(\mathcal{K}')$ then $\mathcal{K} = \mathcal{K}'$.) There will then be a one-to-one correspondence between the set $\mathbb{K}$ of belief states and the set $\{\mathfrak{s}(\mathcal{K}') \mid \mathcal{K}' \in \mathbb{K}\}$ of the support sets of its elements. The following observation shows that the belief states that are reachable with $\circledcirc$ will then coincide with those that are reachable with $\circ$:

Observation 4.21 *Let $\circ$ be the descriptor revision based on $\langle \mathbb{K}, \mathbb{I}, \circledcirc, \mathfrak{s}, \widehat{C} \rangle$. If injectivity holds, then for each $\mathcal{K} \in \mathbb{K}$ and $\imath \in \mathbb{I}$ there is a descriptor Ψ with $\mathcal{K} \circledcirc \imath = \mathcal{K} \circ \Psi$.*

We can use these correspondences to construct a version of descriptor revision that refers directly to belief sets and descriptors, without mentioning the primitive belief states and inputs that we started with. For that purpose, let $\mathbb{X} = \{\mathfrak{s}(\mathcal{K}') \mid \mathcal{K}' \in \mathbb{K}\}$ and $\mathbb{X}_{\mathfrak{s}(\mathcal{K})} = \{\mathfrak{s}(\mathcal{K}') \mid \mathcal{K}' \in \mathbb{K}_{\mathcal{K}}\}$:

Definition 4.22 *Let $\mathcal{L}$ be a language, $\mathbb{X}$ a set of belief sets in $\mathcal{L}$, a an* accessibility function *that assigns to each K in $\mathbb{X}$ a set $\mathbb{X}_K$ with $K \in \mathbb{X}_K \subseteq \mathbb{X}$, and $\widehat{C}$ a monoselective choice function for $\mathbb{X}$. The* descriptor revision $\circ$ *based on $\langle \mathcal{L}, \mathbb{X}, a, \widehat{C} \rangle$ is the operation $\circ$ such that for all $K \in \mathbb{X}$ and all descriptors Ψ for the language $\mathcal{L}$:*

(i_S) *If Ψ is satisfiable within $\mathbb{X}_K$, then $K \circ \Psi = \widehat{C}(\{X \in \mathbb{X}_K \mid X \Vdash \Psi\})$, and*

(ii_S) *otherwise $K \circ \Psi = K$.*

The index of (i_S) and (ii_S) stands for "sentential". We can of course combine the two simplifications. This amounts to letting a in Definition 4.22 have the property $a(K) = \mathbb{X}$ for all K. We can then replace *(i_S)* by the following:

(i_{SL}) If Ψ is satisfiable within $\mathbb{X}$, then $K \circ \Psi = \widehat{C}(\{X \in \mathbb{X} \mid X \Vdash \Psi\})$.

4.6 Conclusion

In this chapter we have done two things in parallel. First, we have removed all references to sentences in the belief state model, and then reintroduced them in a step-by-step fashion, identifying the assumptions required at each stage. At the completion of this process, we have a fully sentential model. However, we have not reintroduced the more problematic assumptions related to possible worlds and remainders that are usually associated with sentential models. In particular, we now have the means to perform belief change through a choice among a finite set of logically finite potential outcomes rather than among an infinite set of logically infinite entities that are not themselves potential outcomes. The expansion property does not hold, and (as will be shown in detail later on) neither does the recovery property.

Secondly, we have introduced the two major formal elements of descriptor revision, namely: (1) the use of belief descriptors as a general representation of the success conditions of belief change, and (2) the direct application of a choice function to the set of potential outcomes of the operation. The rest of this book is devoted to the further development of belief change models employing these two principles.