Chapter 7
Dynamic Descriptors

It is common in belief revision theory to distinguish between *static* and *dynamic* information about a belief state. Static information refers to what the agent believes in that belief state. Dynamic information refers to what changes the agent's beliefs are disposed to undergo, in particular in response to various external inputs. For example, I know that my friend Sarah believes herself to be in excellent health. This is a piece of static information about her belief state. I am also convinced that she will give up that belief if her doctor tells her that she has the early signs of a serious autoimmune disease. This is a piece of dynamic information about her belief state. The following are other examples of such dynamic information:

"If the agent receives the input . . . , then she will believe that . . . ."
"If the agent comes to believe that . . . , then she will also believe that . . . ."
"The agent might in the future come to believe that . . . ."
"Nothing can bring the agent to believe that . . . ."

We can introduce expressions into the formal language that represent these and other patterns of change. Such expressions will be called *dynamic descriptors* in contrast to the *static descriptors*, formed with $\mathfrak{B}$, that we have been concerned with up until now.

In Section 7.1 we will introduce dynamic and autoepistemic descriptors and discuss whether they should preferably be parts of the object language and included in the belief sets, or only be parts of the metalanguage (which is how we have treated the predicate $\mathfrak{B}$). In Section 7.2 an important class of sentences carrying dynamic information, namely Ramsey test conditionals, are generalized to a class of dynamic descriptors called Ramsey descriptors. In Section 7.3 the logical properties of Ramsey descriptors are determined. Section 7.4 puts focus on standard (sentential) conditionals and introduces two alternatives to the Ramsey test. In Section 7.5 we turn to the logic of non-monotonic inference. It is usually considered to be a fragment of the logic of conditional sentences, but that turns out not to be quite true. Finally, in Section 7.6 we introduce modal expressions into our belief change framework.

S.O. Hansson, *Descriptor Revision*, Trends in Logic 46,
DOI 10.1007/978-3-319-53061-1_7

7.1 Representing Autoepistemic Beliefs

By an *autoepistemic belief* is meant a belief that an agent has about her or his own beliefs. Autoepistemic beliefs can be either static or dynamic. In the example above, Sarah is aware that she considers herself to be in excellent health. This is a belief that she has about her own beliefs at the same point in time, in other words a *static autoepistemic belief*. If I were inconsiderate enough to ask her whether she would retain that belief if her doctor told her she has the early signs of a serious autoimmune disease, then her answer would be in the negative. That answer would report a *dynamic autoepistemic belief*, a belief about how she would change her own beliefs in response to new information.

Should (static and dynamic) autoepistemic beliefs be included in belief sets, or does their special nature require that they be kept out? As indicated in Section 3.6, this has been a difficult and sometimes controversial issue in the belief change literature. In descriptor revision, autoepistemic beliefs already have a representation in the form of descriptors, so all we have to do is to move these descriptors from the metalanguage to the object language and extend the belief sets to contain some of them. Let $\mathbb{X}$ be an outcome set. We can form an augmented version $\overline{\mathbb{X}}$ of it, such that $\overline{\mathbb{X}} = \{\overline{X} \mid X \in \mathbb{X}\}$ where each $\overline{X}$ contains exactly the sentences that are introduced with (sequential) use of the following rules:

1. If $\alpha \in \mathrm{Cn}(X)$, then $\alpha \in \overline{X}$
2. If $\alpha \in \overline{X}$ then $\mathfrak{B}\alpha \in \overline{X}$
3. If $\alpha \notin \overline{X}$ then $\neg\mathfrak{B}\alpha \in \overline{X}$

This means that if $p \in X$ and X is consistent, then $\mathfrak{B}p$, $\mathfrak{B}\mathfrak{B}p$, $\mathfrak{B}\neg\mathfrak{B}\neg\mathfrak{B}p$ and an infinite number of other such composite autoepistemic sentences are all elements of X.

Next, let us turn to dynamic descriptors. There is a well-known recipe for including conditional sentences into (extended) belief sets, namely the Ramsey test that was introduced in Section 3.6. It prescribes that $p \rightarrowtail q$ holds at K if and only if $q \in K * p$. Just like the above recipe for $\mathfrak{B}$, this one can be used repeatedly:

1. If $\alpha \in \mathrm{Cn}(X)$ then $\alpha \in \overline{X}$
2. If $\beta \in \overline{X} * \alpha$ then $\alpha \rightarrowtail \beta \in \overline{X}$

In this way, nested conditionals of unlimited length can be included in extended belief sets, for instance[1]:

$p_1 \rightarrowtail (p_2 \rightarrowtail (p_3 \rightarrowtail p_4)) \in \overline{X}$
if and only if $p_2 \rightarrowtail (p_3 \rightarrowtail p_4) \in \overline{X} * p_1$
if and only if $p_3 \rightarrowtail p_4 \in \overline{X} * p_1 * p_2$
if and only if $p_4 \in \overline{X} * p_1 * p_2 * p_3$

[1] At this point we can set aside the problems with the Ramsey test referred to in Section 3.6. It will be shown in Section 7.2 how these problems can be overcome.

However, from the logical feasibility of these constructions it does not follow that they are philosophically plausible.[2] Some authors have claimed that it would be a conceptual mistake to include autoepistemic beliefs in belief sets. For instance, Alvaro del Val has argued that belief sets should be kept free from dynamic autoepistemic information since there is a "need to separate the specification of the agent's beliefs from the specification of the agent's revision policy, which are fully orthogonal, independent issues" [39, p. 223].[3] This view can be called the *thesis of autoepistemic ignorance*; it claims that the formal representation of actual (static) beliefs should contain no information pertaining to how the agent's beliefs will change in response to new inputs.

From a philosophical point of view this is not a plausible standpoint, for the simple reason that an agent's current beliefs and the ways in which she tends to change these beliefs are far from independent or "orthogonal" issues. For instance, there are strong connections between the justificatory structure of an agent's beliefs and how they will change in response to various inputs. If p represents my only justification for believing that q, and I am aware of this justificatory relationship, then we should expect q to be lost if p is given up.[4] Furthermore, and perhaps more importantly, some forms of modal and conditional beliefs seem to be strongly connected with how one would change one's (non-modal and non-conditional) beliefs upon receipt of certain inputs. For a simple example, suppose that you believe the bog bilberry to be acutely poisonous. If you revise your belief set to accommodate the information that I have just picked and eaten a considerable amount of bog bilberries you will, presumably, also believe that I will soon be sick. Other such examples are easily found; it may in fact be more difficult to find examples of "purely static" beliefs that have no impact on how our beliefs will be modified in response to any type of new information. Our static and dynamic beliefs are closely interwoven and in practice often inseparable. A belief set that adequately represents the beliefs held at a particular moment will of necessity contain an abundance of information pertaining to how it will change in response to different inputs. Therefore the thesis of autoepistemic ignorance is untenable.

According to the diametrically opposite standpoint, the agent is fully aware of her own belief state. This can be called the *thesis of autoepistemic omniscience*.[5] For static autoepistemic beliefs it means that the agent is assumed to have perfectly accurate beliefs about what she believes and does not believe at present (even with respect to complex sentences such as $\mathfrak{B}\neg\mathfrak{B}\neg\mathfrak{B}p$). This is not plausible if we take the belief set to represent her actual beliefs, but it may be plausible if we follow Isaac Levi in taking it to represent the beliefs that she is committed to hold. With the former

[2] On what it means to know one's own beliefs, see [231]. On logics employing an autoepistemic belief operation, see [158, 190, 241, 259].

[3] Isaac Levi has expressed a similar view with respect to the dynamic information contained in conditional sentences. See [161] and [163, pp. 49–50]. See also [61, 65, 102].

[4] The effects of justificatory relationships on patterns of belief change has been investigated with models employing belief bases, see for instance [91, 104].

[5] This term was used with essentially the same meaning by Hans Rott [212].

interpretation it would be desirable for static autoepistemic beliefs to be includible in belief sets, but they should not be automatically included whenever they are true. Then an agent who believes in p may or may not believe in $\mathfrak{B}\neg\mathfrak{B}\neg\mathfrak{B}p$.

For dynamic autoepistemic beliefs, the thesis of autoepistemic omniscience is much more demanding. It requires that the agent has perfectly accurate beliefs not only about her present belief state but also about how her beliefs will change in response to any (arbitrarily long) series of inputs that she may be exposed to. Needless to say, this is an utterly unrealistic feature of a formal model.[6] Therefore, a realistic treatment of dynamic autoepistemic beliefs requires that we find a middle way between autoepistemic ignorance and autoepistemic omniscience. We can call this a *thesis of autoepistemic incompleteness*: agents should be modelled as having belief sets that answer some but not all questions about what they will believe after various (series of) operations of belief change.

We can expect the coverage of dynamic autoepistemic beliefs to be highest for single-step changes that do not require the retraction of highly entrenched beliefs. I have fairly well-developed (and probably accurate) beliefs about how my belief set will change if I receive some unsurprising piece of information such as that Real Madrid won their latest match against Granada CF with $5-1$. I am much less certain about how my belief set will change if I learn that Granada CF won over Real Madrid with $13-0$. And if you provide me with a longish list of statements, each of which contradicts a strong belief of mine, then I will be at a loss for what my belief set would look like after I had successively revised by all of them. In a distance-based model such as that developed in Section 6.2 this can be approximated by the assumption that the agent has true autoepistemic beliefs about the belief changes that only take her to belief sets within a certain (small) distance from the present belief set, but not in general for belief changes that take her further away from the present belief state.

Most treatments of autoepistemic beliefs have assumed that these beliefs are all truthful. However, there is no reason to take that for granted. We are no more infallible in these issues than in any others [119]. Therefore, false autoepistemic beliefs should be includible in belief sets, and their inclusion should not (or at least not always) make the belief set inconsistent.

7.2 Ramsey Descriptors

There are many varieties of conditional sentences ("if... then..."-sentences), and several ways to classify them in terms of their meanings have been put forward [6, 92]. For our purposes, a simple typology proposed by Lindström and Rabinowicz is particularly useful. They divided conditionals into two groups: ontic and epistemic (doxastic) conditionals. The crucial difference is that "ontic conditionals concern

[6]The same type of cognitive unrealism is inherent in standard probability theory. If an agent with a probability function $\mathfrak{p}$ learns that q, then (provided that $\mathfrak{p}(q) \neq 0$) her new probability function $\mathfrak{p}'$ is derivable from $\mathfrak{p}$ through the simple formula $\mathfrak{p}'(x) = \mathfrak{p}(x \mid q) = \mathfrak{p}(x\&q)/\mathfrak{p}(q)$.

hypothetical modifications of the *world*, but epistemic conditionals have to do with hypothetical modifications of our *beliefs* about the world" [170, p. 225]. To exemplify the difference, suppose that one late night you have just arrived in a small town that has only two snackbars.[7] You meet a man eating a hamburger. This makes you believe in the following conditional sentence:

(1) If snackbar A is closed, then snackbar B is open.

Soon afterwards, you see that bar A is in fact open. You would then probably not assent to the following conditional sentence:

(2) If snackbar A were closed, then snackbar B would be open.

(1) is most naturally interpreted as an epistemic conditional, i.e. it signalizes that belief in the antecedent would make you believe in the consequent. In contrast, (2) is an ontic conditional, expressing patterns in the world rather than in your beliefs about it. Grammatically, the antecedent of (1) is expressed with a verb in the indicative mood ("is") and that of (2) with a verb in the subjunctive mood ("were"). The grammatical difference reflects common usage in the English language: the antecedents of epistemic conditionals are typically expressed with an indicative verb form and those of ontic conditionals with a subjunctive verb form. However, this connection between meaning and mood only holds in some languages. The difference between ontic and epistemic conditionals is also present in languages that do not express it by shifting between indicative and subjunctive verb forms. Furthermore, as was pointed out by Michael R. Ayers and more recently by Hans Rott, the correlation between meaning and mood is far from perfect in English [9, 216]. Consider the following examples:

(3) If everyone in this room is legally married to someone else in the room, then there is an even number of persons in the room.
(4) If everyone in this room were legally married to someone else in the room, then there would have been an even number of persons in the room.

(3) and (4) differ in grammatical form but they do not differ in terms of the ontic−epistemic distinction. Instead, the indicative form in (3) ("is") imparts the impression that the antecedent is reasonably plausible, whereas the subjunctive form in (4) signals an assumption that it does not hold.

Due to its higher philosophical relevance, the distinction between epistemic and ontic conditionals should preferably replace that between conditionals expressed in the subjunctive respectively indicative mood in the English language.

The Ramsey test has its origin in a famous footnote by Frank Ramsey:

> If two people are arguing 'If p will q?' and are both in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis about q. [210, p. 247]

This rather sketchy proposal was developed by Robert Stalnaker into a general principle that is now commonly called the Ramsey test [240, pp. 101–105].[8] The test

[7] This example is an improvement by Hans Rott [216] of an example first published in [81].

[8] On the Ramsey test, see also Section 3.6 and [6, 85].

is intended for epistemic conditionals, and it can be expressed as an equivalence between on the one hand the epistemic agent's acceptance of the conditional "If p, then q" and on the other hand her propensity to believe in q after revising by p. In formal notation:

$p \rightarrowtail q$ holds at K *if and only if* $q \in K * p$.

In the framework of descriptor revision the right-hand part of this formula can be written $K \circ \mathfrak{B}p \Vdash \mathfrak{B}q$. This opens up for an obvious generalization: we can replace $\mathfrak{B}p$ and $\mathfrak{B}q$ by more general (i.e. not necessarily atomic) descriptors to serve as antecedent respectively consequent. This results in *Ramsey descriptors*, a generalization of (sentential) Ramsey test conditionals that will be denoted $\Psi \Rightarrow \Xi$ [132]. This formula means that if the belief set is revised by Ψ, then the outcome will satisfy Ξ. The Ramsey test can be straightforwardly generalized as follows:

$\Psi \Rightarrow \Xi$ holds at K *if and only if* $K \circ \Psi \Vdash \Xi$.

Standard (sentential) Ramsey test conditionals are of course a special case of Ramsey descriptors, obtainable by defining $p \rightarrowtail q$ as $\mathfrak{B}p \Rightarrow \mathfrak{B}q$. But more interestingly, other forms of conditional belief patterns can also be expressed, such as the following:

"If he gives up his belief that his wife is faithful to him, then he will also lose his belief that she loves him." ($\neg\mathfrak{B}p \Rightarrow \neg\mathfrak{B}q$)

"If she gives up her belief that the first chapter of Genesis is literally true, then she will still believe that God exists." ($\neg\mathfrak{B}p \Rightarrow \mathfrak{B}q$)

"If she makes up her mind on whether this painting is a genuine Picasso or not, then she will come to believe that it is genuine." ($\mathfrak{B}p \vee \mathfrak{B}\neg p \Rightarrow \mathfrak{B}p$)[9]

Dorothy Edgington has pointed out that the conventional form of conditional sentences (represented here as $p \rightarrowtail q$) is insufficient to cover the wide variety of conditionalities that are expressible in ordinary language:

> Any kind of propositional attitude can occur within the scope of a supposition... and hence... a theory of conditionals should be applicable to more than conditional statements. [45, p. 177]

Ramsey descriptors can hopefully facilitate investigations of the wide range of conditional expressions, in addition to standard "if p then q" conditionals, that are available in ordinary language.

7.3 The Logic of Ramsey Descriptors

In the above explication of $\Rightarrow$, $\Psi \Rightarrow \Xi$ was said to hold at K if and only if $K \circ \Psi \Vdash \Xi$. Importantly, whether $\Psi \Rightarrow \Xi$ holds at a belief set K is not a property of that belief

---

[9] It is an interesting issue whether a rational agent can have the autoepistemic belief $\mathfrak{B}p \vee \mathfrak{B}\neg p \Rightarrow \mathfrak{B}p$ without also having the (static) belief p. This relates to the discussion in Section 7.1 on the connection between static and dynamics beliefs.

set alone. There may be two different operations $\circ$ and $\circ'$ such that $K \circ \Psi \Vdash \Xi$ but $K \circ' \Psi \nVdash \Xi$. Then $\Psi \Rightarrow \Xi$ holds at K according to $\circ$ but not according to $\circ'$. Consequently, sentences formed with $\Rightarrow$ have to be evaluated both *at a specific belief set* and *in relation to a specific operation*. In this they differ from sentences with $\mathfrak{B}$ that represent static beliefs. For any static belief p, whether $\mathfrak{B}p$ holds at K is not influenced by the operation of revision we use. It is a property of K alone. This difference between $\Rightarrow$ and $\mathfrak{B}$ is, of course, the defining difference between static and dynamic descriptors.

It follows that whereas $K \Vdash \Psi$ is an adequate representation of what it means for a static descriptor Ψ to be "held" at K, it would be misleading to substitute $\Psi \Rightarrow \Xi$ for Ψ in that formula to express that $\Psi \Rightarrow \Xi$ holds at K. For dynamic descriptors we need to mention the operation of revision. The symbol $\Lleftarrow$ will be used to denote that a dynamic descriptor holds. The truth condition associated with $\Lleftarrow$ will have to refer both to the belief set and to the operation of revision:

Definition 7.1 *Let $\circ$ be a descriptor revision on K and let $\mathbb{X}$ be its outcome set. The* Ramsey descriptor *associated with $\circ$ is the relation $\Rightarrow$ on descriptors such that for all descriptors Ψ and Ξ:*

$\langle K, \circ\rangle \Lleftarrow \Psi \Rightarrow \Xi$ *if and only if* $K \circ \Psi \Vdash \Xi$.

Obviously, it does not follow from $\langle K, \circ\rangle \Lleftarrow \Psi \Rightarrow \Xi$ that the agent is aware that $\Psi \Rightarrow \Xi$ holds at K (or, more precisely, at $\langle K, \circ\rangle$). Using Definition 7.1 we will develop a logic of Ramsey descriptors that does *not* assume that sentences containing $\Rightarrow$ are believed by the agent or included in belief sets. Whether they should be so is a separate issue that we will return to at the end of this section.

If $\circ$ satisfies confirmation then Definition 7.1 has the following special case:

Observation 7.2 *Let $\circ$ be a descriptor revision on K that satisfies confirmation ($K \circ \Psi = K$ whenever $K \Vdash \Psi$), and let $\Rightarrow$ be the Ramsey descriptor associated with $\circ$. Then:*

$\langle K, \circ\rangle \Lleftarrow \mathfrak{B}\top \Rightarrow \Xi$ *if and only if* $K \Vdash \Xi$.

In studies of Ramsey descriptors it is useful to assume that the underlying operation of revision satisfies confirmation. This makes it possible to regain the belief set from the set of satisfied Ramsey descriptors. It follows from Observation 7.2 that $p \in K$ holds if and only if $\langle K, \circ\rangle \Lleftarrow \mathfrak{B}\top \Rightarrow \mathfrak{B}p$.

Definition 7.1 also has the following implications:

Observation 7.3 *Let $\circ$ be a centrolinear revision on K and let $\mathbb{X}$ be its outcome set. Let $\Rightarrow$ be the Ramsey descriptor and $\succeq$ the relation of epistemic proximity that are associated with $\circ$.*

If $\Psi \cup \Xi$ is satisfiable within $\mathbb{X}$, then:

(1) $\langle K, \circ\rangle \Lleftarrow \Psi \Rightarrow \Xi$ *if and only if* $K \circ \Psi = K \circ (\Psi \cup \Xi)$.

(2) $\langle K, \circ\rangle \Lleftarrow \Psi \Rightarrow \Xi$ *if and only if* $\Psi \cup \Xi \succeq \Psi$.

If Ψ is satisfiable within $\mathbb{X}$, then:

(3) $\Psi \succeq \Xi$ *if and only if* $\langle K, \circ\rangle \twoheadleftarrow (\Psi \veebar \Xi \Rightarrow \Psi)$.

Based on Definition 7.1 one would expect $\circ$ and $\Rightarrow$ to be interdefinable so that if we have one of them, then the other can be derived from it. Such a one-to-one relationship between two areas of logic, belief revision and conditional logic, is of course highly interesting.[10] However, there is a limiting case that creates problems for the interdefinability, namely the case of unsatisfiable inputs respectively antecedents. If Ψ is unsatisfiable, then $\Psi \Rightarrow \Xi$ cannot be evaluated with reference to belief sets in which Ψ is satisfied. Belief revision and the logic of conditionals tend to treat this case in different ways. In belief revision, when the input cannot be satisfied, the standard solution is to let the outcome be equal to the original belief set.[11] In conditional logic, the tradition is instead to follow the *ex falso quodlibet* principle according to which a false sentence implies all other sentences.[12] To exemplify this practice, let p be a sentence that is not included in any element of $\mathbb{X}$. Then $\mathfrak{B}p$ is not satisfied at any element of $\mathbb{X}$. It follows from Definition 7.1 that $\mathfrak{B}p \Rightarrow \mathfrak{B}p$ does not hold at K, and the same applies to its translation into sentential conditional logic, $p \rightarrowtail p$. This is contrary to a well-established tradition in conditional logic, where $p \rightarrowtail p$ is almost universally assumed to hold. (See for instance [217, pp. 33, 112–114] and [257, p. 294].)

In order to avoid the convention-bound translation problems in this rather uninteresting limiting case, it is preferable to relate the two frameworks to each other only in the main case. As the following theorem shows, a one-to-one correspondence can then be obtained with plausible postulates for the Ramsey descriptors:

Theorem 7.4 ([132], modified) *Let $\Rightarrow$ be a Ramsey descriptor. Then the following three conditions are equivalent:*

(I) There is a coextensive centrolinear *revision $\circ$ such that the restriction of $\circ$ to inputs that are satisfiable within its outcome set has an associated Ramsey descriptor that coincides with $\Rightarrow$.*

(II) There is a coextensive linear *revision $\circ$ such that the restriction of $\circ$ to inputs that are satisfiable within its outcome set has an associated Ramsey descriptor that coincides with $\Rightarrow$.*

(III) $\Rightarrow$ satisfies:

[10] In the AGM framework such a connection was introduced in [184]. See also [13, 176, 214, 217, 228].

[11] This is the solution commonly chosen for contraction by a tautology [1], for shielded contraction in which some non-tautologous sentences are not contractible [51], and for non-prioritized revision in which some sentences cannot be incorporated into the belief set [137, 179]. In our presentation of descriptor revision, we have followed this tradition. (See for instance Definitions 5.2 and 5.9.)

[12] The *ex falso quodlibet* principle is seldom mentioned in presentations of conditional logic, but it follows from the common principle that if p logically implies q, then $p \rightarrowtail q$ holds in all belief states. See e.g. [26].

If $\Psi \dashVdash \Psi'$*, then* $\langle K, \circ\rangle \twoheadleftarrow \Psi \Rightarrow \Xi$ *if and only if* $\langle K, \circ\rangle \twoheadleftarrow \Psi' \Rightarrow \Xi$*. (left logical equivalence)*

For all Ψ *there is some belief set* $Y \subseteq \mathcal{L}$ *such that for all* Ξ*:* $\langle K, \circ\rangle \twoheadleftarrow \Psi \Rightarrow \Xi$ *if and only if* $Y \Vdash \Xi$*. (unitarity*[13]*)*

$\langle K, \circ\rangle \twoheadleftarrow \Psi \Rightarrow \Psi$ *(reflexivity), and*

If $\langle K, \circ\rangle \twoheadleftarrow \Psi \Rightarrow \Xi$*, then* $\langle K, \circ\rangle \twoheadleftarrow \Psi \Rightarrow \Phi$ *if and only if* $\langle K, \circ\rangle \twoheadleftarrow \Psi \cup \Xi \Rightarrow \Phi$*. (cumulativity)*

The postulates used in the theorem are all generalizations of properties commonly referred to in the logic of (sentential) conditionals.[14]

Reflexivity and left logical equivalence have been given the same names as properties of sentential conditionals that they generalize, namely:

$p \rightarrowtail p$ (reflexivity)
If $\vdash p \leftrightarrow p'$, then $p \rightarrowtail q$ if and only if $p' \rightarrowtail q$. (left logical equivalence)

Cumulativity also has a direct analogue in sentential conditional logic:

If $p \rightarrowtail q$, then $p \rightarrowtail r$ if and only if $p\&q \rightarrowtail r$. (cumulativity)

In sentential conditional logic, cumulativity is usually split into the following two conditions ([177, p. 43], [12]):

If $p \rightarrowtail q$ and $p\&q \rightarrowtail r$, then $p \rightarrowtail r$. (cut)
If $p \rightarrowtail q$ and $p \rightarrowtail r$, then $p\&q \rightarrowtail r$. (cumulative monotony)[15]

The connection between unitarity and well-known properties of sentential connectives is somewhat less obvious but can easily be brought to light. The following is a sentential variant of unitarity:

For all p there is some $Y \subseteq \mathcal{L}$ such that for all q: $p \rightarrowtail q$ if and only if $Y \vdash q$.

It is equivalent with the following property:

$\{q \mid p \rightarrowtail q\} = \mathrm{Cn}(\{q \mid p \rightarrowtail q\})$ (left absorption, left logical absorption)[16]

In a compact logic, left absorption it is equivalent with the combination of the following two, well-known properties of conditionals:

If $p \rightarrowtail q$ and $q \vdash r$, then $p \rightarrowtail r$. (right weakening)

---

[13]Technically, in the logic of descriptors a belief set X is interchangeable with a descriptor Π_X that is satisfied by X but not by any other belief set. (See Definition 4.14.) Therefore unitarity can equivalently be expressed by a requirement that the descriptor $\bigcup\{\Xi \mid \Psi \Rightarrow \Xi\}$ is satisfied by exactly one belief set; this is also why the name "unitarity" was chosen for this postulate.

[14]See [177] or [217, pp. 111–119] for useful overviews of properties of sentential conditionals.

[15]This postulate is called "cautious monotonicity" in [153, p. 178].

[16]See [177, p. 45] and [57, pp. 164–165].

If $p \rightarrowtail q_1$ and $p \rightarrowtail q_2$, then $p \rightarrowtail q_1 \& q_2$. (*And*)

Thus, in summary, the logical properties of Ramsey descriptors ($\Rightarrow$) used in Theorem 7.4 generalize the following properties of sentential conditionals ($\rightarrowtail$): reflexivity, left logical equivalence, cut, cumulative monotony, right weakening, and *And*. (*And* can be omitted since it follows from the other postulates.[17]) These are exactly the logical principles for sentential conditionals that characterize the system C (cumulative reasoning) proposed by Kraus, Lehmann, and Magidor [153, p. 176].[18] However, the parallel between the two analogous systems of postulates is not complete. In particular, the following result has no counterpart for sentential conditionals:

Observation 7.5 ([132]) *If a Ramsey descriptor $\Rightarrow$ satisfies left logical equivalence, unitarity, reflexivity, and cumulativity, then it satisfies:*

If $\langle K, \circ\rangle \twoheadleftarrow \Psi_1 \Rightarrow \Xi$ and $\langle K, \circ\rangle \twoheadleftarrow \Psi_2 \Rightarrow \Xi$, then $\langle K, \circ\rangle \twoheadleftarrow \Psi_1 \veebar \Psi_2 \Rightarrow \Xi$. (Or)

This is analogous to a well-known postulate for sentential conditionals:

If $p \rightarrowtail r$ and $q \rightarrowtail r$, then $p \vee q \rightarrowtail r$. (*Or*)

However, the sentential *Or* does not hold in system C. To the contrary, its addition to C gives rise to the stronger system P (preferential reasoning) [153, p. 190]. This confirms again that the logic of descriptors is distinctly different from that of sentences.

We have constructed Ramsey descriptors as metalinguistic objects. They are not included in the object language from which belief sets are formed, and therefore they are not elements of the belief sets at which they are supported. The reason for this was given in Section 7.1: Including all Ramsey descriptors as elements of the belief sets at which they hold (relative to $\circ$) is tantamount to assuming that the agent is completely and correctly informed about how her beliefs will develop in response to any chain of inputs that she may receive in the future. However, it should be mentioned that in spite of the philosophical counterarguments to such an assumption, it can easily be implemented in the formal system. The same construction can be used that was mentioned in Section 7.1 for representations of autoepistemic beliefs. We can extend each belief set $X \in \mathbb{X}$ into a set $\overset{\Rightarrow}{X}$ that contains, in addition to X, all Ramsey descriptors $\Psi \Rightarrow \Xi$ such that $X \circ \Psi \Vdash \Xi$. We can also generalize $\circ$ to take such Ramsey descriptors as inputs. For instance, if $\overset{\Rightarrow}{Y}$ is the $\overset{\Rightarrow}{X}$-closest extended belief set containing $\Psi \Rightarrow \Xi$, then $\overset{\Rightarrow}{X} \circ (\Psi \Rightarrow \Xi) = \overset{\Rightarrow}{Y}$. This construction

[17] Let $p \rightarrowtail q_1$ and $p \rightarrowtail q_2$. Cumulative monotony yields $p \& q_1 \rightarrowtail q_2$. Reflexivity yields $p \& q_1 \& q_2 \rightarrowtail p \& q_1 \& q_2$, and with right weakening we obtain $p \& q_1 \& q_2 \rightarrowtail q_1 \& q_2$. Applying cut to $p \& q_1 \rightarrowtail q_2$ and $p \& q_1 \& q_2 \rightarrowtail q_1 \& q_2$ we obtain $p \& q_1 \rightarrowtail q_1 \& q_2$. Finally,we apply cut to $p \rightarrowtail q_1$ and $p \& q_1 \rightarrowtail q_1 \& q_2$, and obtain $p \rightarrowtail q_1 \& q_2$ [153, p. 179].

[18] This was pointed out to me by John Cantwell.

has the formal advantage of allowing conditional sentences satisfying the Ramsey test into belief sets, without being affected by the Gärdenfors impossibility theorem that prevents the inclusion of sentential Ramsey test conditionals into the belief sets of AGM [68].[19] But as already indicated, it would be philosophically much more interesting to investigate constructions in which the extended belief set contains a smaller collection of conditional sentences that has at least some likeness to the set of autoepistemic beliefs that an agent can actually hold.

7.4 Alternative Approaches to Conditionals

Although the Ramsey test provides a highly useful account of conditional sentences, at least for some purposes it should only be seen as a first approximation. Even if we restrict our attention to epistemic conditionals, natural language contains several types of such sentences, and we should not expect a single formal account to cover them all. One important source of this complexity is that we are often reluctant to either approve or disapprove of a conditional sentence. Even if we have no difficulty in understanding the two sentences p and q, we may have great difficulties in taking a stand on the conditional sentence "If p then q". In this section I will sketch out two formal approaches to such reluctance and to the mechanisms by which it is sometimes overcome. In the first of these approaches it is overcome with additional deliberation and in the second with additional information.

For the first approach, consider the following example:

> THE NEW COACH: If we replace Susan by Dorothy as a central defender, will the team as a whole play better?
> THE RECENTLY RETIRED COACH: That is very difficult to say, I do not really know.
> THE NEW COACH: Yes, I know this is difficult, but I really need your opinion. Can you think it over?
> THE RECENTLY RETIRED COACH (after thinking for a while): Well, yes. The team as a whole will play better if you replace Susan by Dorothy. [129]

Let p denote that Susan is replaced by Dorothy and q that the team as a whole improves its play. One way to interpret this dialogue is that for her first answer, the retired coach hypothetically revised her beliefs by $\mathfrak{B}p$. She arrived at a belief set $K \circ \mathfrak{B}p$ that satisfied neither $\mathfrak{B}q$ nor $\mathfrak{B}\neg q$. Then she reconsidered the issue, but now aiming to arrive at a belief set satisfying either $\mathfrak{B}q$ or $\mathfrak{B}\neg q$. We can express this

[19] The Gärdenfors theorem is based on the combination of two properties of a belief revision framework: (1) If a sentence p is logically compatible with a belief set K, i.e. $\neg p \notin K$, then the revision $K * p$ does not remove anything from K, i.e. $K \subseteq K * p$. (2) All Ramsey test conditionals are included in the belief sets at which they hold, i.e. $p \rightarrowtail q \in K$ if and only if $q \in K * p$. The combination of (1) and (2) implies that if $q \in K * r$ and $\neg p \notin K$, then $r \rightarrowtail q \in K \subseteq K * p$, thus $q \in K * p * r$. Counterexamples to this pattern are easily found; see for instance the taxi driver example in Section 3.5. Gärdenfors showed that the combination of (1) and (2) is incompatible with a set of plausible formal properties of a belief revision framework [68]. Descriptor revision avoids these problems since it does not satisfy (1). For arguments against (1), see Section 3.5 and [85, 212].

as an extended success condition. She was no longer searching for the most credible belief set satisfying $\mathfrak{B}p$ but for the most credible belief set satisfying both $\mathfrak{B}p$ and $\mathfrak{B}q \vee \mathfrak{B}\neg q$. We can generalize this pattern by defining the following conditional:

$p \rightarrowtail\!\!\!\!\twoheadrightarrow q$ is an abbreviation of $\{\mathfrak{B}p, \mathfrak{B}q \vee \mathfrak{B}\neg q\} \Rightarrow \mathfrak{B}q$.

We can call this an *elicited conditional* [129]. If $\Rightarrow$ is based on centrolinear revision, then $\rightarrowtail\!\!\!\!\twoheadrightarrow$ will be weaker than the standard sentential conditional $\rightarrowtail$, i.e. it holds that

If $p \rightarrowtail q$, then $p \rightarrowtail\!\!\!\!\twoheadrightarrow q$,

but the reverse implication does not hold.

There are interesting logical differences between the standard sentential Ramsey conditional $\rightarrowtail$ and the elicited conditional $\rightarrowtail\!\!\!\!\twoheadrightarrow$. Within the framework of centrolinear revision, the former satisfies the following postulate:

If $p \rightarrowtail q_1$ and $p \rightarrowtail q_2$, then $p \rightarrowtail (q_1 \& q_2)$. (*And*)

However, the corresponding principle for the elicited conditional,

If $p \rightarrowtail\!\!\!\!\twoheadrightarrow q_1$ and $p \rightarrowtail\!\!\!\!\twoheadrightarrow q_2$, then $p \rightarrowtail\!\!\!\!\twoheadrightarrow (q_1 \& q_2)$.

does not hold in general.

To introduce the second approach we can use the following example that was put forward by Lewis [167, p. 1] to illustrate the well-known observation that many conditionals are context dependent[20]:

If kangaroos had no tails, they would topple over.

In a discussion on the principles of mechanics we would have good reasons to assent to this statement. However, in a discussion on evolutionary biology we would probably say to the contrary that if kangaroos had no tails, then their bodies would have had a different weight distribution, so that they would not topple over.

The underlying reason for this variability in interpretation seems to be that the antecedent:

Kangaroos have no tails. (p)

is so indeterminate that epistemic agents cannot be expected to know how to revise by it.[21] In consequence, (sentential) revision by p will be unsuccessful. Provided

[20] The context dependence of conditionals has been referred to as the shiftability problem [79]. Other early discussions can be found in [166, p. 465] and [202, pp.134–135]. Several other examples have been given in the literature: “If frogs were mammals, they would have mammae.” – “If frogs were mammals, they would be the only ones not to have mammae.” [256]. “If I had been John Keats, I should not have been able to write the *Ode to a Nightingale*.” – “If I had been John Keats, then I should have been the man who wrote the *Ode to a Nightingale*.” [79, pp. 5–6].

[21] This is also an illustration of the difficulties involved in representing an actual or hypothetical input (element of $\mathbb{I}$) by a single sentence. (Cf. Section 4.1.) Serious considerations of what would happen if kangaroos had no tails do not come out of the blue, but would typically take place in some context that makes it clear whether physical or biological principles are under scrutiny.

that relative success is satisfied, we will then have $K * p = K$ and consequently $p \notin K * p$. Such epistemic behaviour is in conflict with the AGM theory due to its exceptionless success postulate ($p \in K * p$ for all p), but as we saw in Section 3.2 that postulate does not express a realistic general feature of belief revision. There are input sentences that a rational agent may well reject, either because they are too unrealistic or because they are too vague. Our sentence p belongs to the latter category.[22]

The Ramsey test for conditionals requires revision to be successful. It does not make sense to evaluate $p \rightarrowtail q$ based on whether q holds in $K * p$ unless the latter set actually contains p. But now consider the two statements:

Kangaroos have suddenly lost their tails. (s)
Kangaroos have lost their tails in an evolutionary process. (e)

If a stranger at a party suddenly asks me: "Would kangaroos topple over if they had no tails?", then I will not be able to answer the question since I do not know how to revise by the sentence p. However, if I am asked the same question in a physics class, then I will assume that revision by $p\&s$ is intended. In a biology classroom I would instead interpret it as referring to revision by $p\&e$. Arguably, both $p\&s$ and $p\&e$ are specified enough to allow for successful (hypothetical) revision, i.e. $p\&s \in K * (p\&s)$ and $p\&e \in K * (p\&e)$.[23] Both these revisions can be expected to provide us with a belief set that has a clear answer to the question whether kangaroos will topple over (q). Consequently, in both these contexts the kangaroo conditional can be unambiguously evaluated [129].

This solution requires an adjustment of the underlying operation of revision. Most of the operations introduced in Chapters 4–6 satisfy the postulate of regularity, according to which it follows from $p \in K * (p\&s)$ that $p \in K * p$. This would of course block the solution just described. A precise formal development of this approach to the context-dependence of conditionals will have to be based on an operation of descriptor revision that does not satisfy regularity, such as blockage revision. (Cf. Observation 5.18.)

7.5 Non-Monotonic Inference

Inference and conditionality are both expressed with "If… then…". Not surprisingly, it has often been assumed that the logic of non-monotonic inference can be based on that of conditional sentences [19, 40, 76]. There is, however, an important difference that was well expressed by Kraus, Lehmann, and Magidor in their seminal 1990 paper on non-monotonic inference:

---

[22] On inputs that cannot be processed due to vagueness, see also [117, pp. 1021–1025].

[23] In their respective contexts, $p\&s$ and $p\&e$ are more adequate representations than p of the hypothetical input whose effect on the belief state, specifically with respect to q, is under consideration.

> [C]onditional logic considers a binary intensional connective that can be embedded inside other connectives and even itself, whereas we [in non-monotonic reasoning] consider a binary relation symbol that is part of the metalanguage. [153, p. 170]

Therefore, the logic of non-monotonic inference cannot be exactly the same as that of conditionals, but it can be constructed as the "flat (i.e. nonnested) fragment of a conditional logic" [153, p. 171]. In the same vein, Makinson and Gärdenfors suggested that "q follows non-monotonically from p" holds if and only if q is an element of the outcome of revising an "arbitrary but fixed background theory" by p [184, p. 189]. Denoting that theory by K and non-monotonic consequence by $\mid\!\sim$ we obtain:

The Ramsey test for non-monotonic inference
$p \mid\!\sim q$ holds if and only if $q \in K * p$.

This connection between non-monotonic inference and belief revision has been subject to much further study and refinement [39, 72, 149, 217, 258]. It can now be described as the standard view that the logic of non-monotonic inference coincides with a logic of non-nested conditional sentences and that it is connected to belief revision via the Ramsey test.

However, although conditionality and inferribility are related concepts, it is far from obvious that inferribility is nothing else than (non-nested) conditionality. For instance, let p denote that it rains in London today and q that Flamengo wins the match they are playing tonight in the Maracanã Stadium. If I become convinced that both p and q are true, then I may arguably conclude that "if p then q". However, it would be absurd to also conclude that "from p it can be inferred that q". More generally speaking, inferribility seems to imply conditionality, but not the other way around.

To account for the difference I propose that we retain the Ramsey test for conditionals (with the reservations made in the previous section) but apply another test to non-monotonic inference. The Ramsey test is based on the following criterion:

Ramsey's criterion
If the agent revises her beliefs by p, then she will believe that q.

For non-monotonic inference, the following criterion is proposed:

The co-occurrence criterion
If the agent comes to believe that p, then she will believe that q.

The two criteria differ since an agent can come to believe in p not only as the result of revising her beliefs by p but also as the result of revising them by some other input. According to the co-occurrence criterion, q has to be an element not only of $K * p$ but also of other belief sets containing p. The criterion concerns whether we will *in general* (given our present epistemic commitments) believe in q if we come to believe in p, not only whether we will do so in one single case. This seems to make the criterion better aligned with the notion of inferribility than the Ramsey criterion.

The co-occurrence test needs to be specified with respect to which of the belief sets containing p we should include in the analysis. A simple answer would be to include all potential belief change outcomes that contain p, i.e. all belief sets $K \circ \Psi$ such that $p \in K \circ \Psi$, or at least all belief sets $K * r$ such that $p \in K * r$. However, such an approach would be inadequate since it fails to reflect an essential feature of non-monotonic reasoning, namely that comparatively remote possibilities are left out of consideration. When you conclude from "Tweety is a bird" that "Tweety can fly", then that is precisely because you do not take remote possibilities into account. Importantly, the degree of remoteness referred to here is relative to the antecedent. Some of the possibilities that are too far-fetched to be taken into account when considering "Tweety is a bird" would be quite close at hand when considering "Tweety is a bird who was born in Antarctica".

Based on this, we arrive at the following test of inferribility:

The co-occurrence test for non-monotonic inference[24]
$p \mid\sim q$ holds if and only if q holds in all the p-satisfying belief change outcomes that are reasonably plausible as compared to other p-satisfying belief change outcomes.

We will develop this approach in a centrolinear model that is exhaustive in the sense that $\bigcup \mathbb{X} = \mathcal{L}$. This is what is required for the success postulate for sentential revision ($p \in K * p$) to be satisfied. Although it is not a realistic feature, this condition is adopted here as a simple way to get rid of the rather uninteresting limiting cases of conditionality and inference with non-satisfiable antecedents ($p \rightarrowtail q$ and $p \mid\sim q$ when the epistemic agent cannot be brought to believe that p is true).

In a centrolinear model, when evaluating non-monotonic inferences with p as the antecedent, we have to consider not only $K * p$ that is the most plausible ($\leqq$-minimal) p-containing belief set, but also a band of other p-containing belief sets that are less plausible than $K * p$ but still reasonably plausible.[25] That band has $K * p$ as its inner limit, and since it does not extend indefinitely we must assign an outer limit to it. In formal terms, for each potential outcome X there will be another potential outcome $\ell(X)$ that is the outer limit of the plausibility band that has X as its inner limit.[26] Intuitively, the plausibility band consists, in addition to X, of all the belief sets that are less plausible than X but only moderately so. In formal terms:

Definition 7.6 *The triple* $\langle \mathbb{X}, \leqq, \ell \rangle$ *is a* dilated centrolinear model *if and only if* $\langle \mathbb{X}, \leqq \rangle$ *is a centrolinear model and* ℓ *(the* delimiter*) is a function from and to* $\mathbb{X}$ *such that* $X \leqq \ell(X)$ *for all* $X \in \mathbb{X}$.

[24]This criterion does not preclude the existence of belief change outcomes in which $p\&\neg q$ holds. There can be some sentence r, less plausible than p, such that $p\&\neg q$ holds in some or all of the r-satisfying belief set outcomes. For an example, let p denote that Bitsy is a female mammal, q that Bitsy can give birth to live young, and r that Bitsy is a platypus.

[25]From a formal point of view, this proposal is related to the proposals by Nute [201] and Schlossberger [230, p. 80] that in possible world semantics, the assessment of a conditional sentence should refer not only to the antecedent-satisfying possible worlds that are most similar to the actual world but to all those that are sufficiently similar.

[26]This construction has the property that if $K * p_1 = K * p_2$ then p_1 and p_2 are evaluated with the same set of belief sets. Another plausible property of ℓ is: If $X \leqq Y$ then $\ell(X) \leqq \ell(Y)$. It will not be needed here.

A belief set $X \in \mathbb{X}$ *is* self-limited *according to* ℓ *if and only if* $X = \ell(X)$.

We can now express the co-occurrence test in more precise terms:

The co-occurrence test in an exhaustive and dilated centrolinear model
$p \mathrel{|\!\sim} q$ holds if and only if it holds for all $Y \in \mathbb{X}$ that if $K * p \leqq Y \leqq \ell(K * p)$ and $p \in Y$, then $q \in Y$ [131].

In the limiting case when all elements of $\mathbb{X}$ are self-limited according to ℓ, the co-occurrence test coincides with the Ramsey test (and thus $\mathrel{|\!\sim}$ coincides with $\rightarrowtail$).

This recipe can be straightforwardly extended to take sets as antecedents:

$A \mathrel{|\!\sim} q$ if and only if it holds for all $Y \in \mathbb{X}$ that if
$K \circ \{\mathfrak{B}p \mid p \in A\} \leqq Y \leqq \ell(K \circ \{\mathfrak{B}p \mid p \in A\})$
and $A \subseteq Y$, then $q \in Y$.

With this reformulation we can define a non-monotonic inference operation C such that that $q \in \mathrm{C}(A)$ if and only if $A \mathrel{|\!\sim} q$. Such an operation is an important tool for studying non-monotonic inference and its relationship to classical consequence (as expressed by the consequence operation Cn) [177, 181]. The need for this extension to sets of sentences is, by the way, another reason why non-monotonic inference should not be assumed to coincide with the non-nested fragment of conditional logic.

The following theorem provides us with a close connection between the Ramsey test and the co-occurrence test for single-sentence antecedents:

Theorem 7.7 ([131]) *Let* $\langle \mathbb{X}, \leqq, \ell \rangle$ *be an exhaustive and dilated centrolinear model such that the strict part of* $\leqq$ *is a well-ordering with an order type that is either finite or* ω.[27] *Furthermore, let* $\mathrel{|\!\sim}$ *be the non-monotonic inference relation that is based on* $\langle \mathbb{X}, \leqq, \ell \rangle$ *via the co-occurrence test. Then there is a centrolinear model* $\langle \mathbb{X}', \leqq' \rangle$ *such that* $\mathrel{|\!\sim}$ *coincides with the conditional* $\rightarrowtail$ *that is based on* $\langle \mathbb{X}', \leqq' \rangle$ *via the Ramsey test.*

Furthermore, if K *is the* $\leqq$*-minimal element of* $\mathbb{X}$ *and* K' *is the* $\leqq'$*-minimal element of* $\mathbb{X}'$*, then* $K' \subseteq K$.

At first glance one might be tempted to see this theorem as an argument against the distinction made above between Ramsey test conditionals and non-monotonic inference based on the co-occurence test. However, that would be too rash a conclusion. Two distinct concepts may have the same logical properties.[28] Furthermore,

[27] A well-ordering is a linear ordering such that every non-empty subset of its domain has at least one minimal element. That the strict part $<$ of $\leqq$ has an order type that is either finite or ω means that $<$ is either isomorphic with a finite string $\langle 0, 1, \dots n \rangle$ of natural numbers or with the full infinite series $\langle 0, 1, 2 \dots \rangle$ of natural numbers. This is a stronger requirement than wellfoundedness. For instance, let $\mathbb{X}$ consist of all sets X_k where k is a natural number, and let $X_k < X_m$ hold if and only if either (a) X_k is even and X_m is odd, or (b) X_k and X_m are either both even or both odd, and $k < m$. (This is the sequence $X_0, X_2, X_4 \dots X_1, X_3, X_5 \dots$.) This relation is well-founded since every subset of $\mathbb{X}$ has a $<$-minimal element. However, it does not satisfy the criterion of Theorem 7.7.

[28] Logical necessity and physical necessity may both have the same (S5) logic, but that is no reason to conflate them. ([74, pp. 104–105], cf. [27, 59].) In social choice theory, we usually assume that

although the theorem provides us with a reconstruction of any co-occurrence test as a Ramsey test, this derived Ramsey test is based on another initial belief set and another operation of belief revision than those employed in the co-occurrence test that we started with. Therefore, although the logical properties of $\mid\!\sim$ alone coincide with those of $\rightarrowtail$ alone, the same cannot be said of the logical properties that connect $\mid\!\sim$ respectively $\rightarrowtail$ to the original beliefs or to the operation of belief revision. That is exemplified by the following property of conditonals:

Property CS
If p and q both hold, then so does $p \rightarrowtail q$.

In our centrolinear model, $\rightarrowtail$ satisfies CS, but $\mid\!\sim$ does not do so in general. This is a highly plausible difference, as shown in the above Flamengo example.[29]

7.6 Modalities of Belief

A common way to construct a modal epistemic logic is to let a belief operation $\mathfrak{B}$ take the role that the necessity operation has in alethic logic (the logic of necessity and possibility). The corresponding possibility operation $\diamondplus p$ will then be defined as follows:

$\diamondplus p$ if and only if $\neg\mathfrak{B}\neg p$.

Isaac Levi introduced the term "serious possibility" for this operation. ([159], cf. [61, 163, 164].) It can be translated "p is compatible with what the agent believes". It has also often been interpreted as "the agent considers p possible" [248, p. 23]. However, the latter interpretation is less convincing. If the agent's beliefs are consistent then $\diamondplus$ satisfies the property $\diamondplus p \vee \diamondplus \neg p$ [61, p.120]. This contradicts the common experience of having no opinion on what is possible in some particular matter. For instance, until recently I had no opinion on how high pitches it might be possible for a baboon to hear. Therefore, I did not hold it to be possible that a baboon hears a tone of the highest pitch a human can hear, and neither did I hold the opposite to be possible.[30]

But this is not the only way to introduce modal notions into the logic of belief. Much more expressive power can be obtained by introducing separate representations of necessity and possibility. We can then represent statements about what the agent believes to be possible or necessary, and also about what is possible or necessary for

(Footnote 28 continued)
the preferences of different persons satisfy the same logical rules, but in all non-trivial cases they differ in substance.

[29] $\rightarrowtail$ satisfies CS in any model such that $*$ satisfies confirmation. CS holds in many systems of conditional logic, see for instance [167, pp. 26–31], [207, p. 249], and [203]. However, it has also been criticized, for instance by Bennett [11, pp. 386–388] and Nozick [200, p. 176].

[30] Actually, baboons can even hear tones that are an octave above the upper limit of what a human can hear [243].

the agent to believe. In what follows I will sketch out how the latter development can be achieved in descriptor revision. We can unproblematically introduce necessity and possibility operations that refer to what the agent can come to believe as a result of belief change[31]:

Direct possibility:
$\Diamond\Psi$ holds at K if and only if there is some Ξ such that $K \circ \Xi \Vdash \Psi$.
Direct necessity:
$\Box\Psi$ holds at K if and only if $K \circ \Xi \Vdash \Psi$ for all Ξ.

To exemplify this, $\Diamond(\neg\mathfrak{B}p \;\&\; \neg\mathfrak{B}\neg p)$ means that there is some (single-step) belief change that will make the agent open-minded about p, and $\Box\neg\mathfrak{B}q$ that no (single) change can make the agent believe in q.

Sometimes it takes a whole series of changes to arrive at a new pattern of belief. To express what is necessary or possible through iterated belief change we need another pair of modal operations:

Iterative possibility:
$\diamondplus\Psi$ holds at K if and only if there is some series $\Xi_1, \ldots, \Xi_n$ of descriptors such that $K \circ \Xi_1 \circ \ldots \circ \Xi_n \Vdash \Psi$.
Iterative necessity:
$\boxdot\Psi$ holds at K if and only if $K \circ \Xi_1 \circ \ldots \circ \Xi_n \Vdash \Psi$ for all series $\Xi_1, \ldots, \Xi_n$ of descriptors.

A semantic model for these modal operations can be based on the same type of accessibility relation that is used in possible world models. However, that relation will have to operate on the entities that can be accessed through belief change, and these are belief sets rather than possible worlds. We will therefore have use for the following construction:

Definition 7.8 *A* possible theories model[32] *is a pair* $\langle \mathbb{X}, a \rangle$ *where* $\mathbb{X}$ *is a set of logically closed sets and* a *is a binary relation on* $\mathbb{X}$. *A* pointed possible theories model *is a triple* $\langle \mathbb{X}, a, K \rangle$, *where* $K \in \mathbb{X}$.

The evaluation of modal sentences follows standardly:

$\Box\Psi$ holds at X if and only if it holds for all Y that if XaY then $Y \Vdash \Psi$.
$\Diamond\Psi$ holds at X if and only if there is some Y such that XaY and $Y \Vdash \Psi$.

Obviously, possible worlds models are special cases of possible theories models.

The intended epistemic interpretation is that XaY holds if and only if there is some Ξ such that $X \circ \Xi = Y$. Not surprisingly there is a close connection between possible theories models and global monoselective revision as introduced in Definition 6.2.

[31] The introduction of modal notions with similar definitions into the AGM framework is less promising. Due to the success property ($p \in K * p$), $\Diamond\mathfrak{B}p$ would hold in AGM for all sentences p.

[32] Due to the general nature of this definition, the term "theory" for a logically closed set of sentences is used rather than "belief set" that is limited to epistemological interpretations.

Observation 7.9 *Let $\langle \mathbb{X}, a\rangle$ be a possible theories model. Then the following two conditions are equivalent:*

(1) a is reflexive.

(2) There is a global monoselective revision $\circ$ on $\mathbb{X}$ such that a is the accessibility relation on which $\circ$ is based.

The operations $\Box$, $\Diamond$, $\boxdot$, and $\diamondsuit$ all refer to what belief patterns can at all be reached with our operation of change.[33] Alternatively, we can restrict our deliberations to changes that satisfy some particular pattern. For instance, we may ask whether there is, among the belief changes that retain a person's belief in p, some change that would make her give up q. Generalizing this pattern we can write $\langle\Phi\rangle\Psi$ to denote that there is some Ξ such that $X \circ \Xi \Vdash \Phi$ and $X \circ \Xi \Vdash \Psi$, and correspondingly $[\Phi]\Psi$ to denote that $X \circ \Xi \Vdash \Psi$ for all Ξ such that $X \circ \Xi \Vdash \Phi$. However these operations are only of limited interest since $\langle\Phi\rangle\Psi$ is satisfied in X if and only if $\Diamond(\Phi \cup \Psi)$ is satisfied in X, and if Ψ is negatable then $[\Phi]\Psi$ is satisfied in X if and only if $\neg\Diamond(\Phi \cup \neg\!\!\!\neg\, \Psi)$ is satisfied in X.

But we can go further in this type of restriction on modalities. We can focus on what is necessary or possible after revision by a specific input:

$$[\Xi]\Psi \text{ holds if and only if } \Psi \text{ is satisfied in all outcomes of revision by } \Xi. \tag{1}$$

In deterministic descriptor revision there is exactly one outcome that can result from revision by Ξ, and therefore (1) is equivalent to:

$$[\Xi]\Psi \text{ holds if and only if } K \circ \Xi \Vdash \Psi. \tag{2}$$

Furthermore, in deterministic descriptor revision the corresponding possibility operation $\langle\ \rangle$ will coincide with the necessity operation, i.e.:

$$\text{For all } \Xi \text{ and } \Psi\text{: } [\Xi]\Psi \text{ is satisfied if and only if } \langle\Xi\rangle\Psi \text{ is satisfied.} \tag{3}$$

It follows that in deterministic descriptor revision, $[\Xi]\Psi$ and $\langle\Xi\rangle\Psi$ are nothing else than alternative notations for $K \circ \Xi \Vdash \Psi$. However, in *in*deterministic descriptor revision this trivialization of the modal operations does not take place, since neither (2) nor (3) will hold. In that case, the following definitions will be adequate:

$$[\Xi]\Psi \text{ holds if and only if } X \Vdash \Psi \text{ for all } X \in K \check{\circ} \Xi. \tag{4}$$
$$\langle\Xi\rangle\Psi \text{ holds if and only if } X \Vdash \Psi \text{ for some } X \in K \check{\circ} \Xi. \tag{5}$$

and clearly $\langle\Xi\rangle\Psi$ can be true while $[\Xi]\Psi$ is false. This notation is useful, not least since it brings us into direct contact with other approaches that use similar notations for belief change, such as the update logics developed by van Benthem, Fuhrmann, and de Rijke [37, 62, 246, 247], and Krister Segerberg's Dynamic Doxastic Logic (DDL) [28, 157, 172, 234, 236]. In Segerberg's notation, the standard operations of belief revision come out as follows:

---

[33] See [91] for a study of corresponding modal notions in a belief base framework.

$[*p]\mathfrak{B}q$ (q is believed after revision by p)
$[\div p]\mathfrak{B}q$ (q is believed after contraction by p)
$[+p]\mathfrak{B}q$ (q is believed after expansion by p)

These can all be obtained as special cases of the modal descriptor notation $[\Xi]\Psi$. Segerberg was right in pointing out that the modal notation provides us with an account of belief change that is "a generalization of ordinary Hintikka type doxastic logic", whereas strictly speaking, "AGM is not really logic; it is a theory about theories" [235, p. 136].

The descriptor-based modal operations $[\Xi]$ and $\langle\Xi\rangle$ have the advantage of being easily combinable with other modal operations such as $\Box$, $\Diamond$, $\boxdot$, and $\diamonddot$, as introduced above. We can for instance write $[\Xi]\boxdot\neg\mathfrak{B}p$ to express that after revision by Ξ, any series of belief changes will result in a belief set in which p is not believed. The logic of all these modal operations can be explored in the framework of possible theories semantics. This type of semantics may also be useful for studies of other concepts, such as intentions, goals, and various types of inference. However, these are topics that will have to be left for later investigations.