

Belief Change

RICHARD BOOTH* and THOMAS MEYER†

Abstract

In this paper we present a brief overview of *belief change*, a research area concerned with the question of how a rational agent ought to change its mind in the face of new, possibly conflicting, information. We limit ourselves to *logic-based* belief change, with a particular emphasis on classical propositional logic as the underlying logic in which beliefs are to be represented. Our intention is to provide the reader with a basic introduction to the work done in this area over the past 30 years. In doing so we hope to sketch the main historical results, provide appropriate pointers to further references, and discuss some current developments. We trust that this will spur on the interested reader to learn more about the topic, and perhaps to join us in the further development of this exciting field of research.

1 Introduction

Consider the situation in which an agent has just encountered a bird: let's call it Tweety. Part of the agent's beliefs about the world is that birds fly. Being a logical agent, it therefore believes that Tweety flies. On closer inspection, though, the agent learns that Tweety is an ostrich. Since the agent also believes that ostriches don't fly, it is now faced with a dilemma: Can Tweety fly, or can't it?

The simple scenario above aptly illustrates the central topic of this paper—that a rational intelligent agent is sometimes forced to adjust its current beliefs in some appropriate fashion when confronted with new information. The investigation of the reasoning patterns involved in such a task is known as the study of *belief change*.

The approach to the problem of belief change that we discuss in this paper is logic-based. Both the beliefs of an agent and new information presented to it will be represented in a logic language with a strong emphasis on the case where the underlying logic is a classical propositional

*Univ. of Luxembourg Luxembourg

†Meraka Institute, CSIR and School of Computer Science, Univ. of Kwazulu-Natal, South Africa

logic. Although much of the early work on belief change has a somewhat weaker assumption about the underlying logic, requiring simply that it be a logic equipped with a Tarskian consequence relation and satisfying Compactness, the usual assumption in practice was to use a propositional logic. However, as we shall soon see, propositional logic on its own is not enough to obtain unique answers to the problems of belief change. The primary principle we shall use to guide us is known as the *Principle of Minimal Change*. The idea is simple and intuitive. Information is hard to come by and if an agent has gone to the trouble of incorporating a piece of information into its set of beliefs, it has presumably done so for a good reason. It should therefore give up any beliefs it has only if it is forced to do so. That is, any changes to its current stock of beliefs should be minimal.

Traditionally, approaches to belief change have followed one of two trajectories, with the differences centred around the question of whether beliefs should be represented as *belief bases* (arbitrary sets of sentences) or logically closed *theories*. The AGM approach to belief change [1, 28] (named after its originators Alchourrón, Gärdenfors and Makinson), perhaps the most influential voice within this field of research, is based on the assumption that beliefs need to be represented as theories. The idea here is that we are interested in belief change on the *knowledge level* and that the particular syntactic formulation that we choose for representing the beliefs of an agent is largely irrelevant. On the other hand, the case made for the use of belief bases, which originated with the work of Sven Ove Hansson [36], is that the sentences chosen to represent the beliefs of an agent are somehow more basic than those that merely follow logically from these basic sentences. Although these two approaches start off with different, seemingly conflicting basic assumptions, we shall see that they actually have much in common. In fact, one of the assumptions underlying both approaches is the necessity of introducing additional structure to the representation of beliefs in order to obtain unique results to specific problems in belief change. In much of the work on this topic the additional structure is not represented in the underlying logic itself, but is viewed as meta-information of some kind, and our work here strongly emphasises that approach. Having said that, it is important to note that there is a growing body of work in which information such as this is incorporated into the logic itself—a topic which we touch on in Section 8.

Our intention in this paper is to provide the reader with a basic introduction to the work done in the area of belief change over the past 30 years. In doing so we hope to sketch the main historical results, provide appropriate pointers to further references, and discuss some current developments. Of course, it is impossible to present a truly comprehensive account of a

research area in a paper such as this and our perspective on matters will invariably be subjective, to some extent. The reader is urged to keep this in mind when going through the paper.

On to more concrete matters, then. We commence with a discussion of the formal preliminaries needed to digest the rest of the paper in Section 2. This is followed in Sections 3 and 4 by accounts of the two basic operators investigated in belief change: *belief contraction* and *belief revision*. In Section 5 we take a closer look at the semantic methods for constructing belief change operators before we use this approach to consider *iterated belief revision* in Section 6. Section 7 discusses the links between belief change and the area of *nonmonotonic reasoning*, while Section 8 considers approaches to belief change using *epistemic logics*. Finally, Section 9 takes a brief look at recent developments in belief change before we conclude in Section 10.

2 Preliminaries

First the logical framework. We start with a quite abstract formulation (L, Cn) , where we just have a set L whose elements are the *sentences*, together with a consequence operator Cn which takes sets of sentences $B \subseteq L$ to sets of sentences $Cn(B)$ which intuitively represents all the sentences which are *entailed* by B . Cn is assumed to be a compact *Tarskian* consequence operator (after Alfred Tarski), i.e., it satisfies the following four properties for all $B, B_1, B_2 \subseteq L$:

- $B \subseteq Cn(B)$ **(Reflexivity)**
- $B_1 \subseteq B_2$ implies $Cn(B_1) \subseteq Cn(B_2)$ **(Monotony)**
- $Cn(Cn(B)) = Cn(B)$ **(Idempotence)**
- If $\varphi \in Cn(B)$ then $\varphi \in Cn(B')$ for some finite $B' \subseteq B$ **(Compactness)**

We call any arbitrary set $B \subseteq L$ a *belief base*, but if $B = Cn(B)$, i.e., B is *closed* under Cn , then we call B a *theory*. Following the tradition of the AGM approach, we will use K rather than B to denote theories. $\alpha \in L$ is a tautology iff $\alpha \in Cn(\emptyset)$. From Cn we can define a notion of *consistency*. A set B of sentences is *Cn-consistent* iff $Cn(B) \neq L$. We will just say *consistent* if the consequence operator is clear from the context.

Definition 1. An abstract deduction system is a pair (L, Cn) as above.

This logical setup, although very general, is surprisingly already rich enough to explore many interesting issues in belief change. However, traditionally researchers (including AGM) have worked within more specific background logical systems. In particular the machinery of propositional logic is usually taken as minimum. We may take $L = L_P$, consisting of all sentences built up from some set of propositional variables using the connectives $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$. The classical propositional consequence operator is denoted by Cn_0 . We call a *supraclassical deduction system* a pair (L_P, Cn) where in addition to the four properties mentioned above, Cn is assumed to satisfy

- $Cn_0(B) \subseteq Cn(B)$ **(Supraclassicality)**
- $\varphi \in Cn(B \cup \{\theta\})$ iff $(\theta \rightarrow \varphi) \in Cn(B)$ **(Deduction)**

For supraclassical deduction systems we have $B \cup \{\alpha\}$ is consistent iff $\neg\alpha \notin Cn(B)$. In particular α is consistent iff $\neg\alpha \notin Cn(\emptyset)$.

2.1 The problem formalised

We are now ready to state formally the problem of belief revision:

Assume some fixed abstract deduction system (L, Cn) as background. Then given an initial belief base $B \subseteq L$ and some new information represented as a sentence $\alpha \in L$, find a new belief base $B * \alpha$ which includes α and is consistent.

The requirement that $B * \alpha$ be consistent is crucial here. Without it we might as well just add α set-theoretically to B and stop there. But if $B \cup \{\alpha\}$ is inconsistent then entailment is trivialised, since by definition of inconsistency then *all* sentences are entailed by $B \cup \{\alpha\}$, rendering it useless. So how should we approach this problem?

One influential idea, which comes from Isaac Levi [50] is to decompose the operation into two main steps. First, B is altered if necessary so as to “make room” for, i.e., become consistent with, the incoming sentence α . This is achieved by making B deductively weaker. This is known as *contraction*. Here we should adhere to the principle of minimal change, according to which this weakening should be made as “small” as possible. (See [64] for a discussion of this principle.) In the second, trivial step, the new formula is then simply joined on to the result (this is known as *expansion*). Clearly the difficult step is the first one. So in order to answer the problem of revision we first need to address the problem of contraction. We turn to this in the next section.

3 Belief contraction

Note that for the purposes of revision we just need to make B consistent with α . In the case where the background deduction system is supraclassical this is the same as ensuring $\neg\alpha \notin Cn(B)$. But in general (for instance if negation is not available in the language) these two things will be different. So in general there are two kinds of contraction operator: the first is *inconsistency-based* and the second is *entailment-based*.¹ We will focus on entailment-based contraction here. We denote the result of contracting B so that it no longer entails a given sentence α by $B - \alpha$. We will first deal with the case where B is an arbitrary belief base. Later we will look at the special case where it is a theory.

3.1 Partial meet base contraction

One of the best-known approaches to contraction is *partial meet contraction* [1, 39]. Here the idea is to calculate contraction in three steps:

1. Focus for the first step on those subsets of B which do not entail α and which are *maximal* with this property. We denote this set by $B \perp \alpha$.
2. Then, a certain number of the elements of this set are somehow selected as the “best” or “most preferred” by means of a selection function γ : $\gamma(B \perp \alpha) \subseteq B \perp \alpha$.
3. Finally, the intersection of these best elements is taken: $\bigcap \gamma(B \perp \alpha)$

Let’s formalise all this, starting with the set $B \perp \alpha$ in step 1.

Definition 2. Let $B \subseteq L$ and $\alpha \in L$. Then $B \perp \alpha$ is the set of subsets $X \subseteq L$ such that $X \in B \perp \alpha$ iff (i). $X \subseteq B$, (ii). $\alpha \notin Cn(X)$, (iii). For all $X' \subseteq B$, if $X \subset X'$ then $\alpha \in Cn(X')$. We call $B \perp \alpha$ the set of α -remainders of B

If α is a tautology then $B \perp \alpha = \emptyset$, but this is the only case for which $B \perp \alpha = \emptyset$. This is a result of the following fact, the proof of which requires **Monotony** and **Compactness** of Cn as well as Zorn’s Lemma.

Fact 3 ([2]). If $\alpha \notin Cn(Y)$ and $Y \subseteq B$ then there exists $X \in B \perp \alpha$ such that $Y \subseteq X$.

In other words, every non- α -implying subset of B may be extended to a maximal non- α -implying subset of B . Next comes the definition of selection function.

¹See also Section 9.

Definition 4. Let $B \subseteq L$. A selection function for B is a function γ such that for all $\alpha \in L$, (i) if $B \perp \alpha \neq \emptyset$ then $\emptyset \neq \gamma(B \perp \alpha) \subseteq B \perp \alpha$, and (ii) if $B \perp \alpha = \emptyset$ then $\gamma(B \perp \alpha) = \{B\}$.

Finally we can use a selection function for B to define a contraction operator $-_\gamma$ for B :

$$B -_\gamma \alpha = \bigcap \gamma(B \perp \alpha).$$

Definition 5. If $-$ can be defined via some selection function γ for B as above then $-$ is a partial meet base contraction operator (for B).

Two special cases of partial meet contraction deserve mention. If the selection function picks a single element of $B \perp \alpha$, it is called a *maxichoice contraction*. If it picks the whole of $B \perp \alpha$, it is called a *full meet contraction*. Observe that full meet contraction is unique, whereas there are many different maxichoice contractions: one for each element of $B \perp \alpha$.

Partial meet base contraction may be characterised as follows.

Theorem 6 ([38]). $-$ is a partial meet base contraction operator for B iff it satisfies the following properties:

- If $\alpha \notin \text{Cn}(\emptyset)$ then $\alpha \notin \text{Cn}(B - \alpha)$ **(Success)**
- $B - \alpha \subseteq B$ **(Inclusion)**
- If $\beta \in B \setminus B - \alpha$ then there exists B' such that $B - \alpha \subseteq B' \subseteq B$, $\alpha \notin \text{Cn}(B')$ and $\alpha \in \text{Cn}(B' \cup \{\beta\})$ **(Relevance)**
- If for all $B' \subseteq B$ we have $\alpha \in \text{Cn}(B')$ iff $\beta \in \text{Cn}(B')$ then $B - \alpha = B - \beta$ **(Uniformity)**

The above properties may be explained as follows. **Success** says the sentence to be removed is actually removed, i.e., is no longer a consequence of the base² and **Inclusion** states that no new beliefs may be added in the course of removing α .³ **Relevance** seeks to avoid unnecessary loss of information. It says that a sentence β should be given up only if it contributes to the fact that B , and not $B - \alpha$, entails α . **Uniformity** states that if two sentences are indistinguishable from the viewpoint of B , in that every subset of B which implies one also implies the other, then the results of contracting by them should be the same. As was noted in [42], the only properties of Cn which are actually used in the proof of Theorem 6 are **Monotony** and **Compactness**.

The following two reasonable properties can be shown to follow from those in Theorem 6 (see [36]), and thus are satisfied by any partial meet base contraction operator.

²But see [25] for a discussion on why this is not always desirable.

³See [10] for an argument against this postulate.

- If $\alpha \notin Cn(B)$ then $B - \alpha = B$ **(Vacuity)**
- If $Cn(\alpha) = Cn(\beta)$ then $B - \alpha = B - \beta$ **(Preservation)**

Vacuity says that if α is not entailed by B to begin with, then nothing needs to be changed. It can be shown to be a consequence of **Relevance** and **Inclusion**. **Preservation** says that if two sentences are equivalent under logical consequence then contracting by them should give the same results. It can be shown to follow from **Uniformity**, while in the special case when B is a theory, it is actually equivalent to **Uniformity** in the presence of **Vacuity**.

3.2 Partial meet theory contraction

The preceding construction works equally well when B is taken to be a theory K . But in this case, since the input to contraction is a theory, we should expect the output to be a theory too. This is ensured because in this case the elements of $K \perp \alpha$ are themselves theories, and the intersection of any family of theories is again a theory. When applied to a theory K we will refer to the above construction as *partial meet theory contraction*.

In this case we obtain a different representation theorem, which was one of the main results of AGM.⁴

Theorem 7 ([1]). *Assume we work with a supraclassical deduction system (L_P, Cn) , and let K be a theory. Then $-$ is a partial meet theory contraction operator for K iff it satisfies **Success**, **Inclusion**, **Vacuity**, **Preservation** and the following properties:*

- $K - \alpha = Cn(K - \alpha)$ **(Closure)**
- $K \subseteq Cn((K - \alpha) \cup \{\alpha\})$ **(Recovery)**

Note that this result requires the assumption of a supraclassical deduction system as background. It may not hold for general abstract deduction systems (see [26] for discussion on this).

The postulates listed in the above theorem are collectively known as the *basic AGM contraction postulates*. **Closure** says that the result of theory contraction is another theory, while **Recovery** says that if one removes α and then simply adds it again (and then closes under logical consequence) then one should get back all the initial beliefs K . **Recovery** has been by far the most controversial of the AGM contraction postulates, with many authors calling it into question (see [36, pp. 72-74]). It should be noted that

⁴Note that, historically, partial meet theory contraction actually pre-dates the more general version for arbitrary bases given above.

this postulate is specific to the *theory* version of partial meet contraction, i.e., it does not hold in general for partial meet *base* contraction, where B is allowed to be an arbitrary base. For supraclassical deduction systems, in the presence of **Closure**, **Inclusion** and **Vacuity** it is equivalent to **Relevance** [27].

3.3 The supplementary postulates

In partial meet contraction, when selecting the remainder sets via γ , what we have is an instance of a *choice situation*. We have a number of alternatives up for selection, namely $B \perp \alpha$, and some of them are singled out as being in some sense more preferred. So we make some crossover into the realm of *rational choice*. How can this choice be made? We can assume it is made on the basis of a binary preference relation \sqsubseteq over the set of all possible remainder sets for B , i.e., the set $\{X \mid X \in B \perp \alpha \text{ for some } \alpha \in L\}$. For any two possible remainder sets X, Y , we write $X \sqsubseteq Y$ to mean that Y is at least as preferred as X , and use \sqsubset to denote the strict part of \sqsubseteq , i.e., $X \sqsubset Y$ iff both $X \sqsubseteq Y$ and $Y \not\sqsubseteq X$. Then \sqsubseteq can be used as the basis for a selection function γ_{\sqsubseteq} by setting, in the principal case in which $B \perp \alpha \neq \emptyset$,

$$\gamma_{\sqsubseteq}(B \perp \alpha) = \{X \in B \perp \alpha \mid Y \sqsubseteq X \text{ for all } Y \in B \perp \alpha\}.$$

That is, $\gamma_{\sqsubseteq}(B \perp \alpha)$ consists of those elements of $B \perp \alpha$ which are at least as preferred as *all* other elements of $B \perp \alpha$. If a selection function γ for B is generated from some relation \sqsubseteq in this way then we say γ is a *relational* selection function, and a partial meet contraction operator $-\gamma$ which can be generated from some relational selection function γ will be called a *relational partial meet contraction* operator.

By putting some mild constraints on the relation \sqsubseteq , we can constrain the behaviour of the resulting relational partial meet contraction operator in interesting ways. Consider the following two properties:

- $X \sqsubseteq Y$ and $Y \sqsubseteq Z$ implies $X \sqsubseteq Z$ **(Transitivity)**
- If $X \subset Y$ then $X \sqsubset Y$ **(Maximising)**

The first property is a standard requirement for a relation of preference. The second is motivated by minimal change considerations: when contracting B it is always preferable to retain as much of B as possible, so a given subset Y of B should always be strictly more preferred to any of its strict subsets. If γ is generated from some \sqsubseteq satisfying **Transitivity** then we say $-\gamma$ is a *transitively relational* partial meet contraction operator, while if it generated from some \sqsubseteq satisfying, in addition, **Maximising** then $-\gamma$ is a *transitively, maximisingly relational* partial meet contraction.

Note that, if B is a theory, then these two collapse into the same thing, in the sense that \neg_γ is transitively relational iff it is transitively, maximisingly relational [35].

We obtain the following results, both of which assume we work in a supraclassical deduction system:

Theorem 8 ([39]). *Assume a supraclassical deduction system as background. Let B be a belief base and suppose $-$ is a transitively, maximisingly relational partial meet base contraction operator for B . Then $-$ satisfies the following property:*

- $(B - \alpha) \cap (B - \beta) \subseteq B - (\alpha \wedge \beta)$ (**Conjunctive Overlap**)

For the case of relational partial meet theory contraction we can say more:

Theorem 9 ([1]). *Assume a supraclassical deduction system as back-ground. Let K be a theory and $-$ an operator for K . Then $-$ is a transitively relational partial meet theory contraction operator for K iff it satisfies all the basic AGM contraction postulates (see Theorem 7) plus **Conjunctive Overlap** and the following property:*

- If $\alpha \notin K - (\alpha \wedge \beta)$ then $K - (\alpha \wedge \beta) \subseteq K - \alpha$ (**Conjunctive Inclusion**)

The postulates **Conjunctive Overlap** and **Conjunctive Inclusion** are known as the *AGM supplementary contraction postulates*. They go a step beyond the basic postulates, in that they relate the results of contracting by conjunctions $\alpha \wedge \beta$ with the result of contracting by the individual conjuncts. We refer the interested reader to [1, 36] for discussion on these postulates.

3.4 Kernel contraction

The partial meet approach to contraction focusses on *maximal* subsets of B which do not imply the sentence α to be removed. Another approach is instead to single out the *minimal* subsets which *do* entail α , and then to make sure at least one sentence is removed from each. This is the idea behind Hansson's operation of *kernel contraction* [40], which is a generalisation of the *safe contraction* of Alchourrón and Makinson [3].

Definition 10. *Let B be a belief base and $\alpha \in L$. Then $B \perp \alpha$ is the set of sets X such that (i). $X \subseteq B$, (ii). $\alpha \in Cn(X)$, (iii). If $X' \subset X$ then $\alpha \notin Cn(X')$. We call the elements of $B \perp \alpha$ the α -kernels of B .*

To remove α , it is necessary and sufficient to remove at least one sentence from every α -kernel. To this end, we assume the existence of a function which makes an "incision" into every such set, returning the sentences which must be discarded.

Definition 11. σ is an incision function for B if (i). $\sigma(B \perp \alpha) \subseteq \bigcup B \perp \alpha$, and (ii). $\emptyset \neq X \in B \perp \alpha$ implies $\sigma(B \perp \alpha) \cap X \neq \emptyset$.

Every incision function σ then yields a contraction operator by setting $B -_{\sigma} \alpha = B \setminus \sigma(B \perp \alpha)$.

Definition 12. Let $-$ be an operator for B . If $-$ equals $-_{\sigma}$ for some incision function σ for B then it is called a kernel base contraction operator (for B).

Kernel base contraction may be characterised in the following way:

Theorem 13 ([40]). $-$ is a kernel base contraction operator for B iff it satisfies **Success**, **Inclusion**, **Uniformity** and the following property:

- If $\beta \in B \setminus B - \alpha$ then there exists B' such that $B' \subseteq B$, $\alpha \notin Cn(B')$ and $\alpha \in Cn(B' \cup \{\beta\})$ **(Core retainment)**

As with Theorem 6, it was noted in [42] that the only properties required to prove this result are **Monotony** and **Compactness**. Note that **Core retainment** is weaker than **Relevance** and so we see that every partial meet base contraction operator is a kernel base contraction operator. The converse, however, does not hold, i.e., there exist kernel base contraction operators which are not partial meet base contraction operators (see [36, p91]) for a counterexample, and [24] for more on the relation between partial meet and kernel base contraction).

The above discussion was about kernel *base* contraction. It is also possible to employ the construction in the case when B is a theory except, since $K -_{\sigma} \alpha$ is not guaranteed to be a theory (even when K is), it is necessary to add a post-processing step of closing under Cn . That is, a kernel theory contraction operator for K is any operator of the form $K \approx_{\sigma} \alpha = Cn(K -_{\sigma} \alpha)$, where σ is an incision function for K and $-_{\sigma}$ is defined from σ as for kernel base contraction. However, in this case (at least for supraclassical deduction systems), the distinction between kernel theory contraction and partial meet theory contraction disappears, in that every kernel theory contraction operator is a partial meet theory contraction operator, and vice versa [40].

4 Belief Revision

As stated earlier, once we have a contraction operation for a belief base B , we can use it to define a revision operator via the Levi Identity. The Levi Identity comes in two flavours, according to whether we want the result of revision to be a theory or not. In the former case we take $B * \alpha =$

$(B - \neg\alpha) \cup \{\alpha\}$, in the latter case we take $K * \alpha = Cn((K - \neg\alpha) \cup \{\alpha\})$. We call the former the non-closing Levi Identity, and the latter the closing Levi Identity. Whenever we talk of the Levi Identity in connection with an arbitrary belief base B we shall implicitly assume it is the non-closing version we are using, while if we use it in connection with a theory K , we shall assume the closing version.⁵ Throughout this section we shall assume, for simplicity, that we work in a supraclassical deduction system.

First we deal with arbitrary belief bases, where the result is not expected to be logically closed.

Definition 14. *Let B be a belief base. If $*$ can be defined from some partial meet base contraction operator for B using the (non-closing) Levi Identity then it is a partial meet base revision operator for B .*

Partial meet base revision may be characterised as follows:

Theorem 15 ([39]). *$*$ is an operation of partial meet base revision for B iff it satisfies the following properties:*

- $\alpha \in B * \alpha$ **(Success)**
- If α is consistent then $B * \alpha$ is consistent **(Consistency)**
- $B * \alpha \subseteq B \cup \{\alpha\}$ **(Inclusion)**
- If $\beta \in B \setminus B * \alpha$ then there exists B' such that $B * \alpha \subseteq B' \subseteq B \cup \{\alpha\}$, B' is consistent and $B' \cup \{\beta\}$ is inconsistent. **(Relevance)**
- If, for all $B' \subseteq B$, we have $B' \cup \{\alpha\}$ is consistent iff $B' \cup \{\beta\}$ is consistent, then $B \cap (B * \alpha) = B \cap (B * \beta)$ **(Uniformity)**

Success and **Consistency** are taken as fundamental requirements here.⁶ **Inclusion** places an upper-bound on the result of revision. It says the result should not contain any sentence not included in B , apart from the new information α . **Relevance** and **Uniformity** are similar to their namesakes in the list of base contraction postulates.

Let us move on to theory revision.

Definition 16. *Let K be a theory. If $*$ can be defined from some partial meet theory contraction operator for K via the (closing) Levi Identity, then $*$ is a partial meet theory revision operator for K .*

⁵The Levi Identity breaks revision by α into two steps: contraction (by $\neg\alpha$) and expansion (by α), in **that** order. Another possibility is to reverse this order and do the expansion (by α) first, followed by the contraction (by $\neg\alpha$). This possibility is explored in [39].

⁶Although **Success** is not beyond controversy, since one can certainly imagine situations in which new information is not accepted. See [41, 37] for studies of revision operators which don't satisfy it.

The following result is the AGM characterisation of partial meet theory revision.

Theorem 17 ([1]). ** is a partial meet theory revision operator for a theory K iff it satisfies **Success**, **Consistency** and the following basic AGM postulates for theory revision:*

- $K * \alpha = Cn(K * \alpha)$ **(Closure)**
- $K * \alpha \subseteq Cn(K \cup \{\alpha\})$ **(Inclusion)**
- *If $K \cup \{\alpha\}$ is consistent then $K * \alpha = Cn(K \cup \{\alpha\})$* **(Vacuity)**
- *If $Cn(\alpha) = Cn(\beta)$ then $K * \alpha = K * \beta$* **(Preservation)**

Furthermore, ** is a transitively relational partial meet theory revision operator (i.e. is defined via the Levi Identity from some transitively relational partial meet contraction operator) iff it satisfies, in addition, the following two supplementary AGM revision postulates:*

- $K * (\alpha \wedge \beta) \subseteq Cn((K * \alpha) \cup \{\beta\})$ **(Subexpansion)**
- *If $(K * \alpha) \cup \{\beta\}$ is consistent then $Cn((K * \alpha) \cup \{\beta\}) \subseteq K * (\alpha \wedge \beta)$* **(Superexpansion)**

Observe that, for the remainder of this paper we will take *AGM revision* to mean transitively relational partial meet revision.

It is also possible to use the Levi Identity to define revision from kernel contraction, leading to *kernel revision* operators. Axiomatic characterisations are given in [42], and we refer the reader to that paper for details.

Finally in this section, while the Levi Identity deals with how to define revision in terms of contraction, it is also possible to go the other way and define contraction in terms of revision by using the *Harper Identity* [43]:

$$B - \alpha = B \cap (B * \neg\alpha).$$

The Levi and Harper identities can be thought of as inverses to each other. They ensure a very tight connection between contraction and revision.

5 On the semantic side

The previous sections have been developed against the background of some given, fixed abstract (sometimes supraclassical) deduction system (L, Cn) , which represents the background logic we work in. These systems can be said to be *syntactical*, in the sense that they simply declare (via Cn) which sentences are entailed by which sentences. There is, of course, usually

another side to logic which is the *semantical* side. It is the semantics of a logic which tells us what are the objects, or *possible worlds*, or *models* which the sentences in L are actually *talking about*. In this section we investigate belief change from a more semantical viewpoint. The ideas behind this approach originate in a famous paper by Adam Grove [33]. For this and the next section we make a number of simplifying assumptions: (i) we assume that we are working in a supraclassical deduction system (L_P, Cn) , (ii) we furthermore assume L_P is generated by only finitely many propositional atoms, and (iii) we will focus only on *theory* revision and contraction.

What are the models in a supraclassical deduction system? One way to define them is as the set of *maximally consistent theories* of L_P .

$$\mathcal{W} \stackrel{\text{def}}{=} \{M \subseteq L_P \mid M \text{ is a } Cn\text{-consistent theory and for no } Cn\text{-consistent theory } M' \subseteq L \text{ do we have } M \subset M'\}.$$

Given $M \in \mathcal{W}$ and $B \subseteq L_P$, we say M is a *model* of B iff $B \subseteq M$. Then the set of models of B is denoted by $[B]$.

The set \mathcal{W} defines a consequence relation $Cn_{\mathcal{W}}$ by setting, for any $B \subseteq L_P$, $Cn_{\mathcal{W}}(B) = \bigcap [B]$, i.e., a sentence is entailed by B iff it is contained in all models of B . Then \mathcal{W} provides a semantics which is sound and complete with respect to (L_P, Cn) , in the sense that, for any $B \subseteq L_P$, the identity $Cn_{\mathcal{W}}(B) = Cn(B)$ holds.

Now suppose we have a theory K representing our initial beliefs. This corresponds to the belief that the actual “true” world is one of the worlds in $[K]$. It turns out that performing transitively relational partial meet theory contraction on K is equivalent to choosing, on the basis of some *total preorder* over the set \mathcal{W} , some countermodels of (i.e, models of the negation of) the sentence to be contracted, and adding them to $[K]$. To be more precise, let \leq be a total preorder⁷, or *tpo* for short, over \mathcal{W} . For $M_1, M_2 \in \mathcal{W}$, $M_1 \leq M_2$ may be informally read as “ M_1 is at least as plausible (as a candidate to be the real world) as M_2 ”.⁸ Given any subset $T \subseteq \mathcal{W}$, we denote by $\min_{\leq}(T)$ the minimal elements of T under \leq , i.e., $\min_{\leq}(T) = \{t \in T \mid t \leq t' \text{ for all } t' \in T\}$. We assume \leq is *anchored on* $[K]$, i.e., $[K] = \min_{\leq}(\mathcal{W})$. Then we may use \leq to define a contraction operator

⁷A binary relation \leq over a set S is a *total preorder* iff it is (i) reflexive, i.e., $s \leq s$ for all $s \in S$, (ii) transitive, and (iii) connected, i.e., either $s \leq t$ or $t \leq s$ for all $s, t \in S$.

⁸Note the direction of the inequality. Here, “lower” means “more plausible”. Although somewhat counterintuitive, this is the convention in the area.

for K as follows:

$$K \text{ } \text{-}_{\leq} \alpha = \begin{cases} K & \text{if } \alpha \in \text{Cn}(\emptyset) \\ K \cap \bigcap \min_{\leq}([\neg\alpha]) & \text{otherwise.} \end{cases}$$

In other words, the models of the new theory are obtained by taking the minimal models of $\neg\alpha$ and adding them to the models of K .

Theorem 18 ([33]). *– is a transitively relational partial meet theory contraction operator for a theory K iff – equals -_{\leq} for some tpo over \mathcal{W} which is anchored on $[K]$.*

This is not the only way we could use a plausibility order to define a contraction operator. Rott and Pagnucco introduced and axiomatically characterised the operation of *severe withdrawal* [66].

$$K \text{-}\alpha = \begin{cases} K & \text{if } \alpha \in \text{Cn}(\emptyset) \\ \bigcap \{M \in \mathcal{W} \mid M \leq M' \text{ for some } M' \in \min_{\leq}([\neg\alpha])\} & \text{otherwise.} \end{cases}$$

Here, the models of the new theory are obtained by taking *all* models which are at least as plausible as the \leq -minimal $\neg\alpha$ -models. This operation was independently proposed, using a different construction, by Isaac Levi under the name *mild contraction* [51]. Unlike partial meet theory contraction, severe withdrawal does not satisfy **Recovery**. Yet another possibility was explored by Meyer et al. [52]. *Systematic withdrawal* is just like severe withdrawal except we add to $[K]$ not only the most plausible $\neg\alpha$ -models, but all models which are *strictly* more plausible than them.

$$K \text{-}\alpha = \begin{cases} K & \text{if } \alpha \in \text{Cn}(\emptyset) \\ K \cap \bigcap \{M \in \mathcal{W} \mid M < M' \text{ for some } M' \in \min_{\leq}(\neg\alpha)\} \cap \bigcap \min_{\leq}([\neg\alpha]) & \text{otherwise.} \end{cases}$$

What about defining revision from a plausibility order \leq ? We may just apply the Levi Identity to each of the three families of contraction operator above. It turns out we get the same revision operator in each case, viz.

$$K \text{ } *_{\leq} \alpha = \begin{cases} L & \text{if } \neg\alpha \in \text{Cn}(\emptyset) \\ \bigcap \min_{\leq}([\alpha]) & \text{otherwise} \end{cases}$$

In other words the models of the new theory which results from revising by α are exactly \leq -minimal α -models.

Theorem 19 ([33]). ** is a transitively relational partial meet theory revision operator for a theory K iff * equals $*_{\leq}$ for some tpo over \mathcal{W} which is anchored on $[K]$.*

In the theory case, we have that the AGM postulates are equivalent to total preorders over the set of models. There is a third way of characterising AGM theory contraction, namely as an ordering of entrenchment among the sentences of the language [56, 57]. The best-known version of such entrenchment orderings is the *epistemic entrenchment orderings* put forward by Gärdenfors and Makinson [28, 29]. We do not consider these in details here, but rather refer the interested reader to the references provided.

5.1 A general framework for semantic theory contraction

Although AGM theory contraction is acknowledged as central to the endeavour of belief contraction, we have seen that there are other important versions of theory contraction as well. In Section 5 above we discussed, in addition to AGM transitively relational partial meet theory contraction, also the severe withdrawal operators of Rott and Pagnucco [66] and the systematic withdrawal operators of Meyer et al. [52]. In a recent paper Booth et al. [11] provide a general framework from which these, and a number of other well-known theory contraction operators, can be generated. In addition to the $\text{tpo } \leq$ over \mathcal{W} , they assume a second binary relation \trianglelefteq on \mathcal{W} . Together, \leq and \trianglelefteq provide the *context* in which an agent makes changes to its current beliefs. Intuitively, \trianglelefteq is intended to serve as an aid to \leq in the provision of context. Its role is to relate relevant worlds to one another.

Formally, given a theory K , (\leq, \trianglelefteq) is defined as a K -context iff \leq is a tpo (on \mathcal{W}) anchored on $[K]$ and \trianglelefteq is a reflexive sub-relation of \leq (i.e. \trianglelefteq is reflexive and if $M_1 \trianglelefteq M_2$ then $M_1 \leq M_2$). The K -context (\leq, \trianglelefteq) can then be used to define a contraction operator $_{-(\leq, \trianglelefteq)}$ as follows:⁹

$$K_{-(\leq, \trianglelefteq)} \alpha = \bigcap \{M_1 \mid M_1 \trianglelefteq M_2 \text{ for some } M_2 \in \min_{\leq}[\neg\alpha]\}.$$

That is, the models of the theory resulting from a contraction by α are obtained by locating all the \leq -best models of $\neg\alpha$ and adding to those all worlds that are at least as \trianglelefteq -plausible (i.e. at least as low down in \trianglelefteq).

Booth et al. proceed to show that imposing a variety of constraints on \trianglelefteq allows us to recover a number of different contraction (and related) operators proposed in the literature. For our purposes here, the most important results are the ones repeated below.

Consider the following two properties on \trianglelefteq :

(a) If $M_1 \in \min_{\leq}(\mathcal{W})$ then $M_1 \trianglelefteq M_2$ for all $M_2 \in \mathcal{W}$

⁹Booth et al. [11] assume that the pathological case of trying to remove a tautology does not occur.

(b) If $M_1 \trianglelefteq M_2$ then $M_1 = M_2$ or $M_1 \in \min_{\leq}(\mathcal{W})$

Property (a) states that the \leq -minimal worlds are also the \trianglelefteq -minimal worlds, while (b) requires that, apart from itself, nothing but \leq -minimal worlds may be below any world in \trianglelefteq .

Theorem 20 ([11]). *– is a transitively relational partial meet theory contraction operator for a theory K iff – is equal to $-(\leq, \trianglelefteq)$ for some K -context (\leq, \trianglelefteq) which satisfies (a) and (b).*

So, by insisting on properties (a) and (b) we recover exactly the AGM transitively relational partial meet contraction operators.

Now, consider the following properties on \trianglelefteq :

(c) If $M_1 \leq M_2$, $M_2 \leq M_1$, and $M_3 \trianglelefteq M_1$ then $M_3 \trianglelefteq M_2$

(d) If $M_1 < M_2$ then $M_1 \trianglelefteq M_2$

Property (c) says that whether or not a world M_3 is below M_1 according to \trianglelefteq depends only on the \leq -plausibility rank of M_1 . In the case of property (d), two worlds M_1 and M_2 are required to be \trianglelefteq -connected if M_1 is strictly more \leq -plausible than M_2 . In such a case, since \trianglelefteq is a sub-relation of \leq , it has to be the case that $M_1 \trianglelefteq M_2$.

Theorem 21 ([11]). *– is a severe withdrawal operator for a theory K iff – is equal to $-(\leq, \trianglelefteq)$ for some K -context (\leq, \trianglelefteq) which satisfies (c) and (d).*

In this case, by requiring that properties (c) and (d) hold, we are able to recover exactly the severe withdrawal operators.

Next, consider the following properties:

(e) If $M_1, M_2 \in \min_{\leq}(\mathcal{W})$ then $M_1 \trianglelefteq M_2$

(f) If $M_1 \trianglelefteq M_2$ then $M_1 < M_2$, $M_1 = M_2$, or $M_1 \in \min_{\leq}(\mathcal{W})$

Property (e) forces all \leq -minimal worlds (i.e., all elements of $[K]$) to be \trianglelefteq -connected, while (f) allows the models of K to be connected according to \trianglelefteq , although it does not force them to be.

Theorem 22 ([11]). *– is a systematic withdrawal operator for a theory K iff – is equal to $-(\leq, \trianglelefteq)$ for some K -context (\leq, \trianglelefteq) which satisfies (d), (e) and (f).*

This result shows that the insistence on properties (d), (e) and (f) allows us to recover exactly the systematic withdrawal operators.

To conclude this section we note that the framework of Booth et al. can also be used to recover AGM transitively relational partial meet theory revision.

Theorem 23 ([11]). ** is an AGM transitively relational partial meet theory revision for a theory K iff there is a K -context (\leq, \sqsubseteq) where \sqsubseteq is the identity relation, such that $K * \alpha = K \text{ }_{-(\leq, \sqsubseteq)} \neg\alpha$ for every $\alpha \in L_{\mathcal{P}}$.*

6 Iterated theory revision as revising tpos

Everything in the preceding sections has been about “one-shot” belief change. There is an initial theory, there is some new input and then there is a new theory. However, in realistic settings, a rational agent does not simply “shut down” after incorporating this input, but must be ready to receive the *next* input, followed by further inputs after that. That is to say, belief change is an iterative process, and any theory of belief change worthy of the name should be able to account for this. The question is, then, does the theory sketched in the previous sections adequately handle iterated changes? The answer, as researchers began to realise in the mid 1990s, is “no”.

What does the theory described until now have to say about iterated belief change? Notice that the extra structure required to carry out revision, be it incision functions, selection functions, or total preorders over models is always defined *relative to the theory which is being changed*. Thus, for example, when using the tpo construction, there is a fixed total preorder \leq_K associated to each different theory K . So, to revise a theory K by a sentence φ , we can use the total preorder \leq_K associated to K to compute the result $K * \varphi$. If we then want to further revise this new theory by ψ , then we use the tpo $\leq_{K*\varphi}$ associated to it to compute $(K * \varphi) * \psi$. There are three, interrelated problems with this:

1. There need be hardly any relation between the successive tpos \leq_K and $\leq_{K*\varphi}$, where intuitively we might expect some.
2. Some intuitively plausible properties of iterated revision may be violated (see below).
3. This method totally disregards the role that “revision history” may play in determining results of belief change.

What researchers realised in the mid 1990s is that, to address these shortcomings, the theory of belief change should be widened so that it deals not only with change on the level of theories, but that it should address change in the very structure used to change those theories. A contraction or revision operator should tell us not only what the new theory should be, but should also provide us with a new selection function/incision function/tpo over models which is then the target for the next input. In fact most of the best-known approaches to iterated change deal with tpos rather

than the other ways of modelling the extra-structure. Furthermore the focus in this area tends to be more on revision than contraction (but see [21, 19, 20, 59, 44]) so in the following we focus on iterated theory revision as a problem of revising tpos.

6.1 Revising total preorders

So given K and a total preorder \leq associated to K , the result of revision should be a new theory $K * \varphi$ together with a new associated tpo $\leq_{K*\varphi}$. However we can simplify a bit, since the tpo associated to any theory contains enough information to recapture the theory anyway (since $[K] = \min_{\leq}(\mathcal{W})$). So, our new revision problem may be formulated as follows:

Given an initial tpo \leq over \mathcal{W} , and revision input α , determine a new tpo \leq_{α}^* over \mathcal{W} .

The theory should extend the foregoing theory of single-step revision, which means the new belief set $K(\leq_{\alpha}^*)$ should be derived from the initial tpo and α using the partial meet revision recipe from Theorem 19. This means that the new lowest level $\min_{\leq_{\alpha}^*}(\mathcal{W})$ in the new tpo is determined already - it is equal to $\min_{\leq}([\alpha])$. But what about the rest of the ordering? The most obvious thing to do, if we want to be motivated by the principle of minimal change, is to simply leave the rest of the ordering untouched, and sure enough, this was one of the first proposals for tpo revision. Boutilier called it *Natural Revision* [17, 18], though the idea dates back to [71]. Formally it is defined as follows:

$$M_1 \leq_{\alpha}^{*B} M_2 \text{ iff } \begin{cases} \text{either} & M_1 \in \min_{\leq}([\alpha]) \\ \text{or} & M_1, M_2 \notin \min_{\leq}([\alpha]) \text{ and } M_1 \leq M_2. \end{cases}$$

The problem with natural revision is that it makes *too few* changes. This was recognised by Darwiche and Pearl, who proposed four postulates for regulating tpo revision [22]:

(CR1) If $M_1, M_2 \in [\alpha]$ then $M_1 \leq_{\alpha}^* M_2$ iff $M_1 \leq M_2$

(CR2) If $M_1, M_2 \in [-\alpha]$ then $M_1 \leq_{\alpha}^* M_2$ iff $M_1 \leq M_2$

(CR3) If $M_1 \in [\alpha]$ and $M_2 \in [-\alpha]$ and $M_1 \leq M_2$ then $M_1 \leq_{\alpha}^* M_2$

(CR4) If $M_1 \in [\alpha]$ and $M_2 \in [-\alpha]$ and $M_1 < M_2$ then $M_1 <_{\alpha}^* M_2$

(CR1) and **(CR2)** say that, when revising \leq by α , the relative ordering of models within $[\alpha]$, respectively within $[-\alpha]$, should remain unchanged.

(CR3) and (CR4) say that if a given α -model was judged to be at least as (respectively strictly more) plausible as a given $\neg\alpha$ -model before revising by α , then that relation should be preserved after the revision. Essentially revising by α should not cause any degradation in plausibility of any α -model with respect to the $\neg\alpha$ -models.

As noted by Darwiche and Pearl themselves, the above postulates do not rule out natural revision as a sensible approach to tpo revision, because $*_B$ satisfies all these postulates. However $*_B$ does not satisfy the following strengthening of (CR3) and (CR4), which was suggested independently in [12, 47]:

(CR5) If $M_1 \in [\alpha]$ and $M_2 \in [\neg\alpha]$ and $M_1 \leq M_2$ then $M_1 <_{\alpha}^* M_2$

(CR5) forces there to be a *strict* increase in plausibility of the α -models in relation to the $\neg\alpha$ -models which were not deemed more plausible to begin with.

The above postulates can be repackaged as postulates constraining the theory following a double revision:

(C1) If $\alpha \in Cn(\beta)$ then $K((\leq_{\alpha}^*)_{\beta}^*) = K(\leq_{\beta}^*)$

(C2) If $\neg\alpha \in Cn(\beta)$ then $K((\leq_{\alpha}^*)_{\beta}^*) = K(\leq_{\beta}^*)$

(C3) If $\alpha \in K(\leq_{\beta}^*)$ then $\alpha \in K((\leq_{\alpha}^*)_{\beta}^*)$

(C4) If $\neg\alpha \notin K(\leq_{\beta}^*)$ then $\neg\alpha \notin K((\leq_{\alpha}^*)_{\beta}^*)$

(C5) If $\neg\alpha \notin K(\leq_{\beta}^*)$ then $\alpha \in K((\leq_{\alpha}^*)_{\beta}^*)$

(C1) says if two inputs arrive, the second entailing the first, then the first can be ignored when calculating the resulting theory. (C2) says if two contradictory inputs arrive, then the effects of the first are cancelled out. (C3) and (C4) say that if α would be believed, resp. not rejected, after receiving β alone, then this should not change if β were to be preceded by an input α . Finally (C5) postulates a condition under which belief in an input α is guaranteed to survive the arrival of a subsequent input β .

Theorem 24 ([12, 22, 47]). *Let $*$ be a tpo revision operator such that always $K(\leq_{\alpha}^*) = \bigcap \min_{\leq}([\alpha])$. Then, for each $i = 1, 2, 3, 4, 5$, $*$ satisfies (CRi) iff it satisfies (Ci).*

A few concrete tpo revision operators have been proposed which satisfy all of the above postulates. For example in *lexicographic revision* [58, 71] the new tpo following input α is determined by placing *all* α -models strictly below all $\neg\alpha$ -models while leaving the relative ordering within the sets $[\alpha]$

and $[\neg\alpha]$ unchanged. This is a most radical form of tpo revision, where the new information α is given total priority over the initial ordering \leq . At the opposite end of the spectrum is *restrained revision* [12], in which the strict part of the initial ordering is preserved (apart from the minimal α -models, which become strictly more plausible than all the other models), with α -models being promoted only ahead of the $\neg\alpha$ -models which were on the same plausibility “level” (see also [65]).

6.2 Supplementary postulates for tpo revision

A closer inspection of the postulates **(CR1)**-**(CR5)** reveals that they all deal only with revising a tpo by a single sentence α . In this respect they can be viewed analogously to the set of *basic* AGM postulates for one-step theory revision (see Theorem 17). As we have seen in Section 3.3, a richer and more compelling account of theory change emerges once we move to consider postulates such as the *supplementary* ones, which relate the outcomes of revising by *different* sentences. Can any such rules be conceived of in the case of tpo revision? In [13] the following were considered:

(CR6) If $M_1 \in [\alpha]$, $M_2 \in [\neg\alpha]$ and $M_2 \leq_\alpha^* M_1$ then $M_2 \leq_\gamma^* M_1$

(CR7) If $M_1 \in [\alpha]$, $M_2 \in [\neg\alpha]$ and $M_2 <_\alpha^* M_1$ then $M_2 <_\gamma^* M_1$

(CR6) says that if a model M_2 is considered at least as plausible as another model M_1 *even after* receiving evidence α which clearly points more to M_1 being the actual “true” state of the world than it does to M_2 , then there can be *no* evidence γ which will lead M_1 to become strictly more plausible than M_2 . Postulate **(CR7)** is similar.

The class of tpo revision operators satisfying **(CR1)**-**(CR7)** (along with the postulate stating that the result \leq_α^* is again a tpo) was characterised in [13]. The characterisation assumes that to each model M we can associate two distinct objects M^+ and M^- which can be thought of intuitively as representing M in positive and negative circumstances, respectively. It is then assumed that the set \mathcal{W}^\pm collecting all of these objects for each possible model *itself* comes equipped with a total preorder $\underline{\subseteq}$ satisfying the following conditions for all $M_1, M_2 \in \mathcal{W}$, where \leq is the given tpo over \mathcal{W} which we wish to revise (and where $\underline{\subseteq}$ is the strict part of $\underline{\subseteq}$):

($\underline{\subseteq}$ 1) $M_1^+ \underline{\subseteq} M_2^+$ iff $M_1 \leq M_2$

($\underline{\subseteq}$ 2) $M_1^- \underline{\subseteq} M_2^-$ iff $M_1 \leq M_2$

($\underline{\subseteq}$ 3) $M_1^+ \underline{\subseteq} M_1^-$

Each possible tpo $\underline{\subseteq}$ over \mathcal{W}^\pm satisfying the above conditions determines a tpo revision operator for \leq as follows. Given revision input α we just define the new tpo \leq_α^* by setting, for each $M_1, M_2 \in \mathcal{W}$,

$$M_1 \leq_\alpha^* M_2 \text{ iff } r_\alpha(M_1) \underline{\subseteq} r_\alpha(M_2),$$

where for any γ and M , $r_\gamma(M) = M^+$ if $M \in [\gamma]$, and $r_\gamma(M) = M^-$ if $M \in [\neg\gamma]$. That is, the new tpo \leq_α^* is read off from $\underline{\subseteq}$, with each model M assuming the mantle of its positive instance M^+ if it satisfies the new evidence, and that of its negative instance M^- if it falsifies it.

Theorem 25 ([13]). *Let $*$ be an operator for \leq which, given any sentence α as input, returns a new tpo \leq_α^* over \mathcal{W} . Then $*$ satisfies (CR1)-(CR7) iff it may be generated as above from some tpo $\underline{\subseteq}$ over \mathcal{W}^\pm satisfying $(\underline{\subseteq} 1)$ - $(\underline{\subseteq} 3)$.*

It should be noted that the family of tpo revision operators characterised in the above theorem departs from traditional AGM-style theory revision in one major respect, namely - recalling that each tpo has an associated theory given by $K(\leq) = \min_{\leq}(\mathcal{W})$ - we are *not* always guaranteed to get $\alpha \in K(\leq_\alpha^*)$. That is, the new input might not actually be *believed* after receiving it as input. For a simple counterexample suppose L_p is generated from a single propositional variable p , yielding $\mathcal{W} = \{M_p, M_{\neg p}\}$, where $M_p = Cn_0(p)$ and $M_{\neg p} = Cn_0(\neg p)$. Suppose the initial tpo \leq is such that $M_p < M_{\neg p}$ and that we revise \leq using the order over \mathcal{W}^\pm specified by

$$M_p^+ \underline{\subseteq} M_p^- \underline{\subseteq} M_{\neg p}^+ \underline{\subseteq} M_{\neg p}^-.$$

Then we can see in fact that *for any* sentence α we will have $\leq_\alpha^* = \leq$, i.e., the original tpo \leq will be left undisturbed. In particular we get $\neg p \notin K(\leq_{\neg p}^*) = K(\leq) = Cn_0(p)$.

In [13] it is shown how this general scheme for revising tpos is able to capture as instances several revision operators previously described in the literature. For example the lexicographic revision mentioned at the end of the previous subsection forms one such instance (although the other operator mentioned there - restrained revision - does not). For a detailed treatment of this topic the reader is referred to [13].

7 Belief revision and nonmonotonic reasoning

In this section we discuss the connections between belief revision and the work done in the nonmonotonic reasoning community. A logic is said to be

nonmonotonic if its associated entailment relation \vdash need not satisfy the following monotonicity property: if $A \vdash \beta$ then $A \cup \{\alpha\} \vdash \beta$. With \vdash seen as a relation of plausible consequence, there are many examples to show that monotonicity is an undesirable property. Perhaps the one most deeply entrenched in the nonmonotonic reasoning literature is the Tweety example (the example we used in the introduction). Given that Tweety is a bird, it seems plausible to infer that Tweety can fly. But given the additional evidence that Tweety is an ostrich, we should abandon our conclusion about Tweety's flying capabilities.

While there are many approaches to nonmonotonic reasoning (see e.g., [63, 55]), we consider here the influential framework proposed by Kraus, Lehmann, and Magidor [49] and show that it has a strong connection with AGM belief revision. Formally, Kraus et al. take \vdash to be a binary relation on sentences of a propositional logic where $\alpha \vdash \beta$ is to be read as “ β follows plausibly from α ”. For example, if we represent the information that Tweety is a bird by the atom b , and that Tweety can fly by the atom f , the statement $b \vdash f$ is to be read as “from the fact that Tweety is a bird it follows plausibly that Tweety can fly”. Kraus et al. define \vdash as a *rational consequence relation* iff it satisfies the properties Ref, LLE, RW, And, Or, CM, and RM given below.

- (Ref) $\alpha \vdash \alpha$ (Reflexivity)
 (LLE) If $Cn(\alpha) = Cn(\beta)$ and $\alpha \vdash \gamma$ then $\beta \vdash \gamma$ **(Left Logical Equivalence)**

 (RW) If $\gamma \in Cn(\beta)$ and $\alpha \vdash \beta$ then $\alpha \vdash \gamma$ (Right Weakening)
 (And) If $\alpha \vdash \beta$ and $\alpha \vdash \gamma$ then $\alpha \vdash \beta \wedge \gamma$
 (Or) If $\alpha \vdash \gamma$ and $\beta \vdash \gamma$ then $\alpha \vee \beta \vdash \gamma$
 (CM) If $\alpha \vdash \beta$ and $\alpha \vdash \gamma$ then $\alpha \wedge \beta \vdash \gamma$ (Cautious Monotonicity)
 (RM) If $\alpha \vdash \gamma$ then either $\alpha \wedge \beta \vdash \gamma$ or $\alpha \vdash \neg \beta$ (Rational Monotonicity)

We do not discuss these properties in detail here. Instead, the interested reader is referred to the paper of Kraus et al. [49]. To make the connection with AGM belief revision, we need to go one step further. Gärdenfors and Makinson [30] define \vdash as an *expectation based consequence relation* iff it is a rational consequence relation which also satisfies the property CP given below.

- (CP) If $\alpha \vdash \perp$ then α is *Cn*-inconsistent (Consistency Preservation)

(where \perp is any truth-functional contradiction, e.g., $p \wedge \neg p$). The underlying intuition provided by Gärdenfors and Makinson is that the reasoning

of an agent is guided by its *expectations*. Every expectation based consequence relation \vdash is based on a set of expectations E , playing a role that is analogous to that of a belief set K in theory change. Intuitively, E is the “current” set of expectations of the agent, and the plausible consequences of a sentence α are those sentences β for which $\alpha \vdash \beta$ holds. The set of expectations E is not explicitly mentioned in the definition of an expectation based consequence relation \vdash , but a suitable E can be recovered from \vdash as follows: $E = \{\alpha \mid \top \vdash \alpha\}$. That is, E is taken as the set of plausible consequences of a tautology.

This places us in a position to define a method for translating between belief revision and expectation based consequence relations. Given a consequence relation \vdash , we take the set of expectations E associated with \vdash as the theory K to be revised, and we define $K * \alpha$ as $\{\beta \mid \alpha \vdash \beta\}$. Conversely, given a theory K and a revision operator $*$, we define a nonmonotonic consequence relation \vdash as follows: $\alpha \vdash \beta$ iff $\beta \in K * \alpha$. The main result, linking belief revision to nonmonotonic reasoning is the following theorem by Gärdenfors and Makinson [30] proving that these definitions allow us to show that AGM revision and expectation based nonmonotonic consequence coincide:

Theorem 26. *Let \vdash be an expectation based consequence relation and let $E = \{\beta \mid \top \vdash \beta\}$. Then $E = \text{Cn}(E)$ (i.e. E is a theory). Furthermore, the revision operator $*$ for E , defined in terms of \vdash as follows: $E * \alpha = \{\beta \mid \alpha \vdash \beta\}$, is an AGM revision operator. Conversely, consider a theory K , and let $*$ be an AGM revision operator for K . Then the consequence relation \vdash defined as follows: $\alpha \vdash \beta$ iff $\beta \in K * \alpha$, is an expectation based consequence relation.*

8 Information change in epistemic logic

In the work on belief change we have described thus far we have avoided discussion on its most fundamental notion: that of belief itself. Since the seminal book of Hintikka [45] this question is traditionally explored in the context of modal logic within *doxastic and epistemic logics*.¹⁰ In these logics belief is studied *within* the object language as a modal operator. Although the initial work in this area was concerned purely with *static* notions of belief and knowledge, in more recent times there has been much

¹⁰Strictly speaking, since we deal here with belief rather than knowledge, the adjective *doxastic* (rather than *epistemic*) is the appropriate one. However, since the use of the latter is widespread we shall use it in the rest of this section.

interest in bringing dynamics into the picture, and studying how beliefs change in response to various learning events.

In a language for epistemic logic the sentences are built up from propositional variables using the usual propositional connectives, but now also an explicit modality for belief, so that whenever α is a sentence then so is $B\alpha$. The latter sentence has the intuitive reading that the agent believes α . Semantics is provided by a (single-agent) *epistemic model* $\mathcal{M} = (S, R, v)$, where S is a set of *states*, R is a binary *accessibility relation* over S , and v is a function which assigns a truth-value to every propositional variable at every state. For each $s \in S$, the set $R(s)$ of all states t which are accessible from s , i.e., such that $R(s, t)$ holds, intuitively represents those states which are consistent with the agent's information at state s . Evaluation of sentences is made with respect to a model-state pair (\mathcal{M}, s) , where $s \in S$, with the crucial clause for $B\alpha$ being as follows:

$$(\mathcal{M}, s) \models B\alpha \text{ iff } (\mathcal{M}, t) \models \alpha \text{ for all } t \in S \text{ s.t. } R(s, t).$$

By putting various restrictions on the accessibility relation R we can obtain different properties for the B -operator. For example by assuming that R is serial, transitive and Euclidean¹¹ we obtain the most common modal logic of belief, known as **KD45**, which is the axiomatic system with inference rules Modus Ponens (if α and $\alpha \rightarrow \beta$ are theorems then so is β) and Necessitation (if α is a theorem then so is $B\alpha$) and which has as axioms all instances of propositional tautologies together with the following:

- | | |
|---|---------------------------------|
| K $B(\alpha \rightarrow \beta) \rightarrow (B\alpha \rightarrow B\beta)$ | (Distribution) |
| D $B\alpha \rightarrow \neg B\neg\alpha$ | (Consistency) |
| 4 $B\alpha \rightarrow BB\alpha$ | (Positive Introspection) |
| 5 $\neg B\alpha \rightarrow B\neg B\alpha$ | (Negative Introspection) |

As can be seen in axioms **4** and **5** above, epistemic logics afford us the possibility to express *higher-order* beliefs, or *beliefs about beliefs*, directly in the object language. Moreover, in *multi-agent* extensions of epistemic logic, in which we have a number of different agents $\{1, \dots, n\}$, to each of which we assign its own accessibility relation R_i , we can also have sentences of the form $B_i B_j \alpha$, expressing what one agent i believes about the beliefs of another agent j .

So far our description of epistemic logic only deals with static belief. In order to build dynamics into this framework one may introduce *dynamic*

¹¹ R is serial iff for every s there exists some t such that $R(s, t)$. It is Euclidean iff $R(t, u)$ whenever both $R(s, t)$ and $R(s, u)$.

modalities. This is what is done in *Dynamic Epistemic Logic* (DEL) [75]. Full DEL comes with a rich typology of different belief-changing events, up to (different varieties of) private announcements and complex forms of epistemic action in which different agents can have different perspectives on the learning event. We will just mention here the prototypical kind of event, namely *public announcement* [31, 60] in which all agents in the scenario being modelled simultaneously learn that φ is true. In public announcement logic (PAL) we introduce new dynamic modalities $[\!|\varphi]$, where φ can be any sentence (including one containing belief modalities). The crucial semantic clause is as follows:

$$(\mathcal{M}, s) \models [\!|\varphi]\alpha \text{ iff if } (\mathcal{M}, s) \models \varphi \text{ then } (\mathcal{M}|\varphi, t) \models \alpha,$$

where $\mathcal{M}|\varphi$ is the epistemic model obtained from \mathcal{M} by eliminating all states s' for which $(\mathcal{M}, s') \not\models \varphi$, with the accessibility relations and valuation functions restricted accordingly. Note that a public announcement of φ represents *hard information* that φ is true. As such it is more in the spirit of belief *expansion* than revision (but see [4, 74]). Some interesting things happen if we try to reformulate the AGM revision postulates in terms of PAL. For one thing, the natural translation of the **Success** postulate does not hold, i.e., the sentence $[\!|\varphi]B\varphi$ need not be valid for all choices of φ . The best known counterexample is if we take φ to be a Moore-type sentence such as $p \wedge \neg Bp$, where p is a propositional variable (“ p is true, but I don’t believe it”). While it may well be the case that $(\mathcal{M}, s) \models p \wedge \neg Bp$, the sentence $B(p \wedge \neg Bp)$ is inconsistent (i.e., true in no state). Thus although a Moore sentence may be truthfully announced, it can *never* be believed *after* the announcement.

The above considerations give rise to a distinction in the epistemic logic literature between *static* and *dynamic* belief revision. In static belief revision the result of revision expresses what an agent comes to believe about what *was* the case *before* the actual learning event took place. Dynamic revision deals with what the agent now comes to believe *after* the learning. The distinction only comes into effect when revising by higher-order beliefs such as with the Moore sentence above. When dealing with *factual* beliefs as with AGM the two notions coincide.

One way to model static belief revision which has been explored is via *doxastic conditionals*. This involves allowing sentences of the form $B^\alpha\varphi$ into the language, expressing the hypothetical belief that if the agent learned α then he would believe that φ was true before the learning. Roughly-speaking, the semantical structures for this language are obtained by replacing an agent’s set of accessible states from each state by binary plausibility relations (total preorders) over states [7, 9]. Then the doxastic con-

ditional $B^\alpha\varphi$ evaluates to true iff $(\mathcal{M}, s') \models \varphi$ for all the minimal states s' in the ordering such that $(\mathcal{M}, s') \models \alpha$. Within this framework one can define dynamic modalities for learning events of *soft* information, unlike the hard information of public announcement. These events have the effect of modifying the agents' plausibility relations rather than eliminating states from the picture completely. Van Benthem [73] studies modalities for two such events: *lexicographic upgrade* $[\uparrow\varphi]$ and *conservative upgrade* $[\uparrow\varphi]$, which essentially correspond respectively to the lexicographic tpo revision method and to Natural revision described in Section 6.1 of the present paper.

For more details on belief change in epistemic logics we refer the interested reader to the survey articles [32] and [8]. Section 7 of the latter also includes a detailed comparison with AGM belief revision.

9 Current developments: belief change for other logics

From the work discussed so far it is clear that belief change has come a long way in the past 30 years. However, a look back at the work done over this period reveals an interesting tendency. Although the original aims were phrased in terms of a broad class of logic—all those with Tarskian consequence relation and satisfying **Compactness**—most of the work done in the area is actually based on the assumption of an underlying *propositional logic*, whether finitely or infinitely generated. In this section we consider a departure from this trend, and discuss recent developments in belief change expressed in two logics other than full propositional logic: *propositional Horn logic* and *description logics*.

9.1 Propositional Horn contraction

One of the main reasons for considering belief change for Horn logic is that it has found extensive use in artificial intelligence and database theory, in areas where belief change is an issue to consider, such as logic programming, truth maintenance systems, and deductive databases. Delgrande [23] was the first to point this out and to investigate the contraction of theories for propositional Horn logic.

A *Horn clause* is a sentence of the form $p_1 \wedge p_2 \wedge \dots \wedge p_n \rightarrow p_{n+1}$ where $n \geq 0$, and where the p_i s are propositional atoms or one of \perp or \top . A *Horn sentence* is a conjunction of Horn clauses. A *Horn set* is a set of Horn sentences. Given a propositional language L_P , the Horn language L_H

generated from L_P is simply the Horn sentences occurring in L_P . The Horn logic obtained from L_H has the same semantics as the propositional logic obtained from L_P , but just restricted to Horn sentences. A *Horn theory* is a Horn set closed under logical consequence, but containing only Horn sentences. We denote Horn consequence by $Cn_H(\cdot)$.

Delgrande's main contributions were threefold. Firstly, he showed that the move to Horn logic leads to two different types of contraction which coincide in the full propositional case. Given a Horn theory H , the *entailment-based* contraction, or *e-contraction*, of a sentence α should result in a new Horn belief $H -_e \alpha$ of which α is not a logical consequence: $H -_e \alpha \not\models \alpha$. On the other hand, the *inconsistency-based* contraction, or *i-contraction*, of a sentence α should result in a new Horn belief $H -_i \alpha$ which is such that adding α to it does not result in an inconsistency: $H -_i \alpha \cup \{\alpha\} \not\models \perp$. In full propositional logic, a way to express *i-contraction* in terms of *e-contraction* would be to require that $H -_e \neg\alpha \cup \{\alpha\} \not\models \perp$. This cannot be expressed in Horn logic, though, because it is not possible to express the negation of the Horn sentence α (see also Section 3). Below we consider only *e-contraction*. Similar results have been obtained for *i-contraction* as well.

Delgrande's second contribution was to show that *e-contraction* for Horn theories should not satisfy the controversial **Recovery** postulate. As an example of the failure of **Recovery** for *e-contraction*, take $H = Cn_H(\{p \rightarrow r\})$ and let $\alpha = p \wedge q \rightarrow r$. Then any reasonable version of *e-contraction* will yield $H -_e \alpha = Cn_H(\emptyset)$. So $Cn_H(H -_e \alpha \cup \{\alpha\}) = Cn_H(\{p \wedge q \rightarrow r\})$ and therefore $H \not\subseteq Cn_H(H \cup \{\alpha\})$.

Delgrande's third contribution was to base the construction of Horn contraction operators on partial meet contraction. The definitions of remainder sets, selection functions, and partial meet contraction, as well as maxi-choice and full meet contraction all carry over directly to *e-contraction* and we will not repeat them here. We refer to these as *e-remainder sets* (denoted by $H \perp_e \alpha$), *e-selection functions*, partial meet *e-contraction*, maxi-choice *e-contraction* and full meet *e-contraction* respectively. As in the full propositional case, it is easy to verify that all *e-remainder sets* are also Horn theories, and that all partial meet *e-contractions* (and therefore the maxi-choice *e-contractions*, as well as full meet *e-contraction*) produce Horn theories.

In two subsequent papers, Booth et al. [15, 16] extended Delgrande's work in a number of interesting ways. They show that while Delgrande's partial meet constructions are all appropriate choices for *e-contraction* in Horn logic, they do not constitute *all* the appropriate forms of *e-contraction*. For example, let $H = Cn_H(\{p \rightarrow q, q \rightarrow r\})$. It can be verified that, for the

e -contraction of $p \rightarrow r$, maxichoice yields either $H_{mc}^1 = Cn_H(\{p \rightarrow q\})$ or $H_{mc}^2 = Cn_H(\{q \rightarrow r, p \wedge r \rightarrow q\})$, that full meet yields $H_{fm} = Cn_H(\{p \wedge r \rightarrow q\})$, and that these are the only three partial meet e -contractions. Now consider the Horn theory $H' = Cn_H(\{p \wedge q \rightarrow r, p \wedge r \rightarrow q\})$. It is clear that $H_{fm} \subseteq H' \subseteq H_{mc}^2$. But observe that H' is not a partial meet e -contraction. Booth et al. argue that H' ought to be regarded as an appropriate candidate for e -contraction and, more generally, that *every* Horn theory between full meet and some maxichoice e -contraction ought to be seen as an appropriate candidate for e -contraction.

Definition 27. For Horn theories H and H' , $H' \in H \downarrow_e \alpha$ iff there is some $H'' \in H \perp_e \alpha$ s.t. $(\bigcap H \perp_e \alpha) \subseteq H' \subseteq H''$. We refer to the elements of $H \downarrow_e \alpha$ as the infra e -remainder sets of H wrt α .

Definition 28. Let H be a Horn theory. An infra e -selection function is a function τ such that for every $\alpha \in L_H$, $\tau(H \downarrow_e \alpha) = H$ whenever $H \downarrow_e \alpha = \emptyset$, and $\tau(H \downarrow_e \alpha) \in H \downarrow_e \alpha$ otherwise. We use an infra e -selection function τ to define an infra e -contraction as $H -_\tau \alpha = \tau(H \downarrow_e \alpha)$.

Booth et al. show that infra e -contraction is captured precisely by the AGM postulates for theory contraction, except that **Recovery** is replaced by the **Core retainment** postulate we encountered earlier in the context of defining kernel contraction in Section 3.4.

Theorem 29 ([16]). Every infra e -contraction satisfies **Closure, Inclusion, Success, Extensionality, and Core retainment**. Conversely, every e -contraction which satisfies **Closure, Inclusion, Success, Extensionality, and Core retainment** is an infra e -contraction.

It is possible to define a version of kernel contraction for Horn logic, simply by closing under Horn consequence the results obtained from kernel contraction for bases.

Definition 30. Given a Horn theory H and an incision function σ for H , the kernel e -contraction for H is defined as $H \approx_\sigma^e \alpha = Cn_H(H -_\sigma \alpha)$, where $-_\sigma$ is the base kernel contraction for H obtained from σ .

Booth et al. prove that kernel e -contraction corresponds *exactly* to infra e -contraction. From these results it seems that the contraction of Horn theories exhibits a kind of “hybrid” behaviour, somewhere between classical base contraction and classical theory contraction. As evidence for this, recall firstly that in the classical case, partial meet contraction and kernel contraction coincide for theories, but that kernel contraction is more general than partial meet contraction when dealing with the contraction of bases.

Furthermore, Horn e -contraction for theories does not satisfy the **Recovery** postulate, unlike classical contraction for theories, but similar to classical base contraction. And finally, the set of postulates provided by Booth et al. to characterise infra e -contraction (and kernel e -contraction) bears a close resemblance to the postulates for characterising Horn contraction for bases in the classical case.

To summarise, these recent investigations into Horn contraction have highlighted the fact that a move away from propositional logic as the underlying logic for belief change can yield interesting and unexpected results. Interestingly enough, although the motivation for initiating research on Horn contraction was partially motivated by an interest in Horn logic in its own right, another reason for doing so is that propositional Horn logic forms the backbone of a group of *description logics*, the class of logics to which we turn to next.

9.2 Belief change for description logics

Description Logics (or DLs for short) are a well-known family of logics used for knowledge representation [6]. They have become the formalism of choice for representing formal ontologies [46]. DLs are decidable fragments of first-order logic, mainly characterised by constructors that allow complex concepts (unary predicates) and roles (binary predicates) to be built from atomic ones. We provide a brief description of two well-known DLs referred to as \mathcal{ALC} and \mathcal{EL} , and show how they relate to belief change.

In the description logic \mathcal{ALC} [69], concept descriptions are built from concept names using the constructors disjunction ($C \sqcup D$), conjunction ($C \sqcap D$), negation ($\neg C$), existential restriction ($\exists R.C$) and value restriction ($\forall R.C$), where C, D stand for concepts and R for a role name. To define the semantics of concept descriptions, concepts are interpreted as subsets of a domain of interest, and roles as binary relations over this domain. An interpretation I consists of a non-empty set Δ^I (the domain of I) and a function \cdot^I (the *interpretation function* of I) which maps every concept name A to a subset A^I of Δ^I , and every role name R to a subset R^I of $\Delta^I \times \Delta^I$. The interpretation function is extended to arbitrary concept descriptions as follows. Let C, D be concept descriptions and R a role name, and assume that C^I and D^I are already defined. Then $(\neg C)^I = \Delta^I \setminus C^I$, $(C \sqcup D)^I = C^I \cup D^I$, $(C \sqcap D)^I = C^I \cap D^I$, $(\exists R.C)^I = \{x \mid \exists y \text{ s.t. } (x, y) \in R^I \text{ and } y \in C^I\}$, and $(\forall R.C)^I = \{x \mid \forall y, (x, y) \in R^I \text{ implies } y \in C^I\}$. The distinguished concept name \top is always interpreted as $\top^I = \Delta^I$. Similarly, the distinguished concept name \perp is always interpreted as $\perp^I = \emptyset$. A

DL *Tbox* contains statements of the form $C \sqsubseteq D$ (*inclusions*) where C and D are (possibly complex) concept descriptions. Tboxes are used to represent the terminology part of ontologies in different application areas. The semantics of Tbox statements is as follows: an interpretation I *satisfies* $C \sqsubseteq D$ iff $C^I \subseteq D^I$. I is a *model* of a Tbox iff it satisfies every statement in it. A Tbox statement φ is a *logical consequence* of a Tbox T , written as $T \vDash \varphi$, iff every model of T is a model of φ .

A concept name A is *concept-satisfiable* wrt to a Tbox T iff there is a model, say I , of T in which $A^I \neq \emptyset$. This turns out to be an important property for ontology construction—if some concept names are *concept-unsatisfiable* wrt a Tbox T it is usually an indication of modelling errors made during the construction of T . For example, Schlobach et al. [68] show the following part of a Tbox for the DICE medical terminology:

```
brain  $\sqsubseteq$  CentralNervousSystem
brain  $\sqsubseteq$  BodyPart
CentralNervousSystem  $\sqsubseteq$  NervousSystem
NervousSystem  $\sqsubseteq$   $\neg$ BodyPart
```

According to this, a brain is a body part as well as a central nervous system, while the latter is a type of nervous system, which, in turn, is not a body part. Formally, the concept *brain* is concept-unsatisfiable wrt the Tbox. Checking for concept-satisfiability is closely related to checking for logical consequence. Indeed, for many DLs, including \mathcal{ALC} , checking for concept-satisfiability can be reduced to checking for logical consequence. DL reasoners such as RACER [34] and FaCT++ [72] are able to detect concept-unsatisfiability quite efficiently.

The link with belief change comes in with attempts to deal with concept-unsatisfiability in appropriate ways. *Ontology debugging* [48, 68] is concerned with determining the cause of concept-unsatisfiability in a Tbox T , while *ontology repair* [67, 54] aims to modify T in such a way that all concept names become concept-satisfiable. It turns out that the techniques used for ontology debugging are closely related to the special case of kernel contraction for belief bases known as safe contraction, which was mentioned in Section 3. Recall that the α -kernels of a base B are the minimal subsets of B implying α . Similarly, techniques for ontology debugging identify the minimal subsets of a Tbox T with respect to which at least one concept name is concept-unsatisfiable.

In ontology debugging the Tbox T isn't modified automatically. Instead, the ontology engineer, when presented with the “kernels” of Tbox statements, is expected to use this information to modify T manually in order

achieve concept-satisfiability. In contrast, the aim of ontology repair is to modify T automatically to ensure concept-satisfiability. This is achieved by removing exactly one element from each of the “kernels” of Tbox statements, an approach that can be seen as safe contraction applied to concept-satisfiability. Ontology repair, in this sense, has more in common with belief *base* contraction than with *theory* contraction, since it is the Tbox statements occurring explicitly in the Tbox that are used to obtain the TBox “kernels”, and not statements in the theory obtained from the the Tbox.

A different application of belief contraction, this time one that is more closely related to *theory* contraction, occurs in ontologies represented in one of the \mathcal{EL} family of DLs [5]. In \mathcal{EL} itself, the basic member of this DL family, concept descriptions are built up from concept names using just conjunction ($C \sqcap D$) and existential restriction ($\exists R.C$). As in \mathcal{ALC} , Tbox statements have the form $C \sqsubseteq D$, where C and D are (possibly) complex concepts. The lack of expressivity in \mathcal{EL} is made up for by the efficiency of reasoning algorithms for it. In particular, the task of *computing the subsumption hierarchy* for an \mathcal{EL} Tbox T (determining whether $T \models A \sqsubseteq B$ for all concept names A and B) has polynomial complexity (in the size of the Tbox). Moreover, it turns out that a member of the \mathcal{EL} family is sufficiently expressive to represent a number of biomedical ontologies, including the widely used medical ontology SNOMED [70].

As with \mathcal{ALC} , the application of belief change to \mathcal{EL} is also related to the construction of ontologies. In this case, however, it does not address concept-unsatisfiability. Indeed, since \mathcal{EL} does not have negation, concept-unsatisfiability can only occur if the bottom concept \perp is used explicitly. Instead, it relates to a different method for testing the quality of a constructed ontology: asking a domain expert to inspect and verify the computed subsumption hierarchy. Correcting such errors involves the expert pointing out that certain subsumptions are missing from the subsumption hierarchy, while others currently occurring in the subsumption hierarchy ought not to be there. A concrete example of this involves the medical ontology SNOMED [70] which erroneously classified the concept `Amputation-of-Finger` as being subsumed by the concept `Amputation-of-Arm`. Finding a solution to problems such as these can be seen as an instance of *theory contraction*, in this case by the statement `Amputation-of-Finger` \sqsubseteq `Amputation-of-Arm`. The scenario also illustrates why we are concerned with contraction of theories and not bases. In general, ontologies are not constructed by writing down DL axioms, but rather using ontology editing tools such as SWOOP¹² or Protégé¹³, from which the

¹²<http://code.google.com/p/swoop>

¹³<http://protege.stanford.edu>

axioms are generated automatically. Because of this, it is the theory obtained from a Tbox that is important, not the axioms from which the theory is generated.

It is only recently that researchers have started to pay attention to theory contraction for \mathcal{EL} [14]. Indeed, much of the work relevant to this topic does not address the \mathcal{EL} family of DLs directly at all. In particular, the work on propositional Horn contraction is of importance in this context. Horn clauses correspond closely to subsumption statements in DLs, since both Horn logic and the \mathcal{EL} family lack full negation and disjunction. In this respect, there is still much work to be done before a claim can be made that belief contraction for \mathcal{EL} has been addressed properly.

Finally, in this section we have focused on recent work related to belief *contraction* for descriptions logics, but it must be pointed out that there has also been some recent work on belief *revision* and related questions [53, 61, 62, 76].

10 Conclusion

In conclusion, we hope that this brief overview of belief change has convinced the reader that research in this area has come a long way over the past 30 years, with the fundamentals of the topic now firmly in place. The main challenge ahead is to build on the established fundamentals and extend the work that has been done to new application areas. As we have seen in Section 9, this is already taking place. And although much remains to be done in this regard with, for example, different underlying logics raising interesting and unexpected questions, it seems clear that the existing body of work provides an appropriate springboard for finding solutions to those new issues that are cropping up.

Acknowledgements

Thanks are due to Guillaume Aucher and an anonymous reviewer for some helpful comments.

References

- [1] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

- [2] C. Alchourrón and D. Makinson. Hierarchies of regulations and their logic. In Hilpinen, editor, *New Studies in Deontic Logic*, pages 125–148. D. Reidel Publishing Company, 1981.
- [3] C. Alchourrón and D. Makinson. On the logic of theory change: Safe contraction. *Studia logica*, 44(4):406–422, 1985.
- [4] G. Aucher. Private announcement and belief expansion: an internal perspective. *Journal of Logic and Computation*, to appear.
- [5] F. Baader, S. Brandt, and C. Lutz. Pushing the \mathcal{EL} envelope. In *Proceedings of IJCAI*, pages 364–369, 2005.
- [6] F. Baader and W. Nutt. Basic description logics. In Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider, editors, *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
- [7] A. Baltag and S. Smets. The logic of conditional doxastic actions: A theory of dynamic multi-agent belief revision. In *Proceedings of the Workshop on Rationality and Knowledge, ESSLLI'06*, pages 13–30, 2006.
- [8] A. Baltag, H. van Ditmarsch, and L. Moss. Epistemic logic and information update. In *Handbook of Philosophy of Information*, volume 8 of *Handbook of Philosophy of Science*, pages 361–455. 2008.
- [9] O. Board. Dynamic interactive epistemology. *Games and Economic Behavior*, 49(1):49–80, 2004.
- [10] R. Booth, S. Chopra, A. Ghose, and T. Meyer. Belief liberation (and retraction). *Studia Logica*, 79(1):47–72, 2005.
- [11] R. Booth, S. Chopra, T. Meyer, and A. Ghose. Double preference relations for generalised belief change. *Artificial Intelligence*, 174(16-17):1339–1368, 2010.
- [12] R. Booth and T. Meyer. Admissible and restrained revision. *Journal of Artificial Intelligence Research (JAIR)*, 26:127–151, 2006.
- [13] R. Booth and T. Meyer. How to revise a total preorder. *Journal of Philosophical Logic*, to appear, 2011.

- [14] R. Booth, T. Meyer, and I.J. Varzinczak. First steps in \mathcal{EL} contraction. In *Proceedings of the IJCAI 2009 Workshop on Automated Reasoning about Context and Ontology Evolution (ARCOE 2009)*, 2009.
- [15] R. Booth, T. Meyer, and I.J. Varzinczak. Next steps in propositional horn contraction. In *Proceedings of IJCAI*, pages 702–707, 2009.
- [16] R. Booth, T. Meyer, I.J. Varzinczak, and R. Wassermann. A contraction core for horn belief change: Preliminary report. In *Proceedings of the 12th International Workshop on Nonmonotonic Reasoning (NMR 2010)*, 2010.
- [17] C. Boutilier. Revision sequences and nested conditionals. In *Proceedings of IJCAI*, pages 519–519, 1993.
- [18] C. Boutilier. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25(3):263–305, 1996.
- [19] J. Cantwell. Some logics of iterated belief change. *Studia Logica*, 63(1):49–84, 1999.
- [20] S. Chopra, A. Ghose, and T. Meyer. Non-prioritized ranked belief change. *Journal of Philosophical Logic*, 32(4):417–443, 2003.
- [21] S. Chopra, A. Ghose, T. Meyer, and K.S. Wong. Iterated belief revision and the recovery axiom. *Journal of Philosophical Logic*, 37(5), 2008.
- [22] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
- [23] J.P. Delgrande. Horn clause belief change: Contraction functions. In *Proceedings of KR*, pages 411–421, 2008.
- [24] M.A. Falappa, E.L. Fermé, and G. Kern-Isberner. On the logic of theory change: Relations between incision and selection functions. In *Proceedings of ECAI*, pages 402–406, 2006.
- [25] E. Fermé and S. O. Hansson. Shielded contraction. In *Frontiers in Belief Revision*, pages 85–107. Kluwer, 2001.
- [26] G. Flouris, D. Plexousakis, and G. Antoniou. On Generalizing the AGM postulates. In *STAIRS 2006: Proceedings of the third starting AI researchers' symposium*, pages 132–143, 2006.

- [27] A. Fuhrmann and S.O. Hansson. A survey of multiple contractions. *Journal of Logic, Language and Information*, 3(1):39–75, 1994.
- [28] P. Gärdenfors. *Knowledge in Flux : Modeling the Dynamics of Epistemic States*. The MIT Press, Cambridge, Massachusetts, 1988.
- [29] P. Gärdenfors and D.. Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of TARK*, pages 83–95. 1988.
- [30] P. Gärdenfors and D.. Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65:197–245, 1994.
- [31] J. Gerbrandy and W. Groeneveld. Reasoning about information change. *Journal of Logic, Language and Information*, 6(2):147–169, 1997.
- [32] P. Gochet and P. Gribomont. Epistemic logic. In *Handbook of the History of Logic*, volume 7, pages 99–195. 2006.
- [33] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [34] V. Haarslev and R. Möller. Racer system description. In *Proceedings of IJCAR 2001*, volume LNAI 2100, 2001.
- [35] S. O. Hansson. Changes on disjunctively closed bases. *Journal of Logic, Language and Information*, 2:255–284, 1993.
- [36] S. O. Hansson. *A Textbook of Belief Dynamics*. Kluwer Academic Publishers, 1999.
- [37] S. O. Hansson, E. Fermé, J. Cantwell, and M. Falappa. Credibility-limited revision. *Journal of Symbolic Logic*, 66(4):1581–1596, 2001.
- [38] S.O. Hansson. A dyadic representation of belief. In *Belief Revision*, pages 89–121. Cambridge Univ Press, 1992.
- [39] S.O. Hansson. Reversing the levi identity. *Journal of Philosophical Logic*, 22(6):637–669, 1993.
- [40] S.O. Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59(3):845–859, 1994.
- [41] S.O. Hansson. A survey of non-prioritized belief revision. *Erkenntnis*, 50(2):413–427, 1999.

- [42] S.O. Hansson and R. Wassermann. Local change. *Studia Logica*, 70(1):49–76, 2002.
- [43] W.L. Harper. Rational conceptual change. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 1976, pages 462–494. Philosophy of Science Association, 1976.
- [44] M. Hild and W. Spohn. The measurement of ranks and the laws of iterated contraction. *Artificial Intelligence*, 172(10):1195–1218, 2008.
- [45] J. Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- [46] I. Horrocks, P. Patel-Schneider, and F. van Harmelen. From *SHIQ* and RDF to OWL: The Making of a Web Ontology Language. *Journal of Web Semantics*, 1(1):7–26, 2003.
- [47] Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.
- [48] A. Kalyanpur, B. Parsia, E. Sirin, and J. Hendler. Debugging unsatisfiable classes in OWL ontologies. *Journal of Web Semantics - Special Issue of the Semantic Web Track of WWW2005*, 3(4):268–293, 2005. (To Appear).
- [49] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
- [50] I. Levi. Subjunctives, dispositions and chances. *Synthese*, 34(4):423–455, 1977.
- [51] I. Levi. *Mild contraction: Evaluating loss of information due to loss of belief*. Oxford University Press, USA, 2004.
- [52] T. Meyer, J. Heidema, W. Labuschagne, and L. Leenen. Systematic withdrawal. *Journal of Philosophical Logic*, 31(5):415–443, 2002.
- [53] T. Meyer, K Lee, and R. Booth. Knowledge integration for description logics. In *Proceedings of AAAI05*, pages 645–650, 2005.
- [54] T. Meyer, K. Lee, R. Booth, and J. Pan. Finding maximally satisfiable terminologies for the description logic *alc*. In *Proceedings of AAAI*, pages 269–274. AAAI Press, 2006.

- [55] R. Moore. Possible-world semantics for autoepistemic logic. In *Proceedings of the AAAI Workshop on Non-Monotonic Reasoning*, New Paltz, NY, pages 344–354, 1984.
- [56] A. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41:353–390, 1994.
- [57] A. Nayak, P. Nelson, and H. Polansky. Belief change as change in epistemic entrenchment. *Synthese*, 109:143–174, 1996.
- [58] A. Nayak, M. Pagnucco, and P. Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146:193–228, 2003.
- [59] A.C. Nayak, R. Goebel, and M.A. Orgun. Iterated belief contraction from first principles. In *Proceedings of IJCAI*, pages 2568–2573. Morgan Kaufmann Publishers Inc., 2007.
- [60] J. Plaza. Logics of public communications. In *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.
- [61] G. Qi, W. Liu, and D. Bell. Knowledge base revision in description logics. In *Proceedings of JELIA*, pages 386–398, 2006.
- [62] G. Qi, W. Liu, and D. Bell. A revision-based approach for handling inconsistency in description logics. *Artificial Intelligence Review*, 2006.
- [63] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [64] H. Rott. Two dogmas of belief revision. *The Journal of Philosophy*, 97(9):503–522, 2000.
- [65] H. Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. *Towards Mathematical Philosophy*, pages 269–296, 2006.
- [66] H. Rott and M. Pagnucco. Severe withdrawal (and recovery). *Journal of Philosophical Logic*, 28:501–547, 1999.
- [67] S. Schlobach. Diagnosing terminologies. In *Proceedings of AAAI05*, pages 670–675, 2005.

- [68] S. Schlobach and R. Cornet. Non-standard reasoning services for the debugging of description logic terminologies. In *Proceedings of IJCAI*, pages 355–360. Morgan Kaufmann, 2003.
- [69] M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48:1–26, 1991.
- [70] S. Schulz, B. Suntisrivaraporn, F. Baader, and M. Boeker. SNOMED reaching its adolescence: Ontologists’ and logicians’ health check. *International Journal of Medical Informatics*, 78(1):S86–S94, 2009.
- [71] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. *Causation in decision, belief change, and statistics*, 2:105–134, 1988.
- [72] D. Tsarkov and I. Horrocks. Description logic reasoner: System description. In *Proceedings of IJCAR*, pages 292–297, 2006.
- [73] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.
- [74] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Public announcements and belief expansion. In *Proceedings of AiML-2004 (Advances in Modal Logic)*, pages 62–73, 2004.
- [75] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007.
- [76] Z. Wang, K. Wang, and R. Topor. Revision of DL-Lite knowledge bases. In *Proceedings of the Description Logics Workshop*, 2009.