Part II

Learning and Definability

Chapter 4

Learning and Belief Revision

Learning can be described as a process of acquiring new information. This acquisition can take forms as different as things are that can be learned. We say: 'He learned that she cheated on him' or 'She learned about his disease', but also: 'She learned a language' or 'He finally learned how to behave'. The first two sentences are about a change in informational state induced by accepting a fact, getting to know something. The latter two are different, they describe a situation in which an inductive acquisition process came to a successful end.

The first kind of learning—getting to know about facts—is formalized and analyzed in the domain of belief revision and the diverse frameworks of epistemic and doxastic logics. The main aim here is to formalize the elementary dynamics of knowledge and epistemic attitudes towards incoming information.

The second kind—learning as a process—is studied within the framework of formal learning theory. In this framework a general concept (language, grammar, theory) gets to be identified by an agent on the basis of some elementary data (sentences, results of experiments) over a long period of time. The learning agent is allowed to change his mind on the way, and the process is successful if it results in convergence to an appropriate hypothesis. In a sense this kind of learning is built on top of the first kind, it consists of an iteration of simple getting-to-know events.

In this chapter we propose a way to use the framework of learning theory to evaluate belief-revision policies. Our interest is shared by at least two existing lines of research. Kelly, Schulte, & Hendricks (1995) and Kelly (1998a,b, 2004, 2008) focus on bringing together some classical belief-revision policies (among others those proposed by Boutilier, 1996; Darwiche & Pearl, 1997; Grove, 1988; Spohn, 1988) with the framework of function learning (see Chapter 2, Section 2.1.3, and for more details Blum & Blum, 1975). In this attempt the possible concepts to be learned or discovered are the possible sequential histories. The problem of prediction which seems to be at the heart of this approach is obviously useful for modeling certain kind of scientific inquiry. However in general, changes of epistemic states do not have to happen according to some prescribed sequence.

They are often governed by sequences of facts that are closed under permutation with respect to their informational content.

Martin & Osherson (1997, 1998) have also worked on establishing the connection between learning theory and belief revision. Their attempt has its roots in the classical AGM framework (Alchourrón et al., 1985) and treats belief revision as a two step process: the shrinking of the current belief state to accommodate the new information (belief contraction) and the incorporation of the data (see Levi, 1980). In this approach, the dominant features of modeling inductive learning as iterated belief revision are that the belief state is treated syntactically, as a set of sentences of a given language, and is assumed to be a full-blown theory (closed under the operation of consequence), incoming data get a fully trusted welcome, and last but not least, the agent does not explicitly consider other, perhaps counter-factual, possibilities.

Following Gierasimczuk (2009a,c) and Dégremont & Gierasimczuk (2009) we advance a different line of research. On the inductive inference side, we are interested in the paradigm of language learning which is more general than the aforementioned function learning approach. We assume that the data are observed in a random manner, so that in general predicting the future *sequence* is not feasible, or even relevant. As possible concepts that are inferred we take sets of atomic propositions. Therefore, receiving new data corresponds to getting to know about facts. On the side of belief revision we follow the lines of dynamic epistemic logic (see Van Benthem, 2007). Hence, we interpret current beliefs of the agent (hypothesis) as the content of those possible worlds that he considers most plausible. The revision does not only result in the change of the current hypothesis, but can also induce modification of the agent's plausibility order.

We are mainly concerned with *identifiability in the limit* (Gold, 1967). In the first part we restrict ourselves to learning from sound and complete streams of positive data. We show that learning methods based on belief revision via conditioning (update) and lexicographic revision are universal, i.e., provided certain prior conditions, those methods are as powerful as identification in the limit. Those prior conditions, the agent's prior dispositions for belief revision, play a crucial role here. We show that in some cases, these priors cannot be modeled using standard belief-revision models (as based on well-founded preorders), but only using generalized models (as simple preorders). Furthermore, we draw conclusions about the existence of tension between conservatism and learning power by showing that the very popular, most 'conservative' belief-revision methods, like Boutilier's minimal revision, fail to be universal. In the second part we turn to the case of learning from both positive and negative data. Here, along with information about facts the agent receives negative data about things that do not hold of the actual world. We again assume these streams to be truthful and we draw conclusions about iterated belief revision governed by such streams. This enriched framework allows us to consider the occurrence of erroneous information. Provided that errors occur finitely often and are always eventually corrected we show that the lexicographic revision method is still reliable, but more conservative methods fail. Before we get to the formal content of this chapter, let us first give two additional philosophical motivations for our work.

Evaluation of Belief-Revision Policies The traditional approach to the problem of belief revision (Alchourrón et al., 1985) can be summed up in the following way. The belief is taken to be a set of sentences, often assumed to be closed under some operation of consequence. Then, confronted with an incoming sentence the belief set has to undergo some transformation. If the sentence is consistent with the belief set, it is simply set-theoretically added, and the set is extended to include all consequences. If the sentence contradicts information contained in the set, the latter has to be modified by first removing inconsistency, and only then performing the addition. In general the belief-revision procedure takes the following form:

$$\langle \text{belief set, proposition} \rangle \rightarrow \text{revised belief set;}$$

in other words:

$$\langle S, \alpha \rangle \rightarrow S * \alpha.$$

Intuitively, belief revision, in order to be rational, has to conform to certain general rules. An intuitive set of rules of this kind, or axiomatization, if one prefers, of belief revision has been proposed by Alchourrón et al. (1985). Investigations into the properties of this type of revision led to the following difficulty: often there is more than one way to make a set consistent with some initially inconsistent input. For instance, if the belief set $S = \{\varphi, \varphi \rightarrow \psi\}$, and the incoming information is $\neg\psi$ then dropping either of the sentences in S would make the set S consistent with $\neg\psi$. How do we decide which one should be chosen? This problem indicates the need for some preference order that underlies beliefs and governs the order of potential elimination (Alchourrón et al., 1985). The postulate of ordering the beliefs according to their entrenchment has essentially enriched the framework. The results indicating the necessity of orders led to involving them explicitly as parts of belief states. A system that accounts for the ordering has been provided by Grove (1988), who represented AGM postulates in terms of systems of spheres. This modeling is a direct predecessor of the modal logic based approach to belief, as it presupposes a total well-founded preorder on the initial uncertainty range. The framework then conformed to a new scheme:

$$\langle (\text{belief set}, \leq),\ \text{proposition} \rangle \rightarrow (\text{revised belief set, revised} \leq);$$

in other words:

$$\langle (S, \leq), \alpha \rangle \rightarrow (S * \alpha, \leq^{\alpha}).$$

This approach led to many new questions, among others: the origin and justification of $\leq$; possible ways of transforming $\leq$ and the ways in which they are (or should be) chosen.

Although this approach has been shown to be quite powerful, it is also very controversial in the fact that it is very syntactic. Dependence on the specific language and closing off under a consequence relation present problems when comparing it to linguistic and cognitive reality. In the meantime an alternative, more semantic approach to belief change has been developed, in which modal logic turns out to be very useful (see Stalnaker, 2009). Despite those developments the old uneasiness remains. How are we supposed to judge and choose between different belief-revision policies? Perhaps, by introducing another level in which some new preference relation will order different policies and this preference ordering itself should become a matter of taste or character? Rott (2008) expects that this question will eventually lead to some sort of circularity in the domain of belief-revision theory. It can be argued that this problem could ultimately be solved only by psychological research in human cognitive tendencies. He poses the following two options as answers to this difficulty:

> One option is to insist on divide-and-conquer strategy: Researchers in belief revision should put their efforts into finding out which methods are best in which contexts. [...] [A] second option. This option assumes that there is some level in the belief state hierarchy up to which constraints of rationality reside, but above which, at all higher levels, we are just describing, in an idealized way, various ways that people happen to be. [...] In this perspective, there are no objective standards for rationality beyond the first level as characterized by AGM.

In the present chapter of this book we will challenge this position by showing that applying certain types of rules in certain contexts can be analyzed in terms of whether they can be relied upon in the 'quest for the truth' (the analysis of inductive inference in terms of reliability has been for the first time provided by Kelly, 1996). We will analyze certain belief-revision policies in terms of their dependability and show differences in their learning power. In our framework we can naturally treat the procedural aspect of iterated belief revision, address some intermediate stages of such iterations and relate them to the ultimate success of a belief-revision policy. Hence, belief-revision methods can get evaluated on the basis of their learning power. Finally, it can be argued that the above-mentioned 'different ways that people happen to be' can be traced as evolutionary equilibria that correspond to effectiveness and reliability of methods.

Safety and Stability of Beliefs The classical definition characterizes knowledge as true justified belief (see, e.g., Chisholm, 1982). In a modern setting this has been formalized as the state of certainty, because a decent justification of a theory should eliminate all other possibilities. Such definition is difficult to

accept from a philosophical standpoint, and many arguments against it can be (and have been) formulated (see, e.g., Gettier, 1963). One of them is that in fact knowledge is a dynamic phenomenon and it rarely occurs in the form of irrevocable states of certainty. Alternatives oscillate around the concept of knowledge as *safe belief*. The strength of safety is in the guarantee it gives: the safe belief is not endangered by the occurrence of true data. If we restrict our considerations to truthful information, or at least assume that mistakes happen rarely, safety can be reformulated in terms of stability. In other words, knowledge emerges when stability is reached. The need for such a notion appeared in many different frameworks: from reaching an agreement in a conversational situation (see, e.g., Lehrer, 1965, 1990) to the considerations in the domain of philosophy of science (see, e.g., Hendricks, 2001).

In this work we account for and characterize the emergence of both: the restrictive kind of knowledge (certainty) and stable belief. We explicitly formulate the conditions under which certain belief states give raise to the emergence of such epistemic and doxastic states.

Finally, inspired by Nozick (1981), Rott (2004) puts forward that perhaps:

> [...] knowledge [should] be made of still sterner stuff—stuff that also survives (a modest amount of) misinformation.

In the following sections we will show that under the requirement of convergence to stable belief some policies are still reliable if a finite number of errors occur and they are all corrected later in the process.

4.1 Iterated Belief Revision

In our analysis of single agent information-update and belief revision we will redefine the framework of dynamic epistemic logic in order to simplify things. As we are here concerned with the single-agent case and moreover, we take the incoming information to be propositional, we will focus on the notion of *epistemic state*, i.e., *a set of possible worlds.*

Definition 4.1.1. *A* possible world *is a valuation over* PROP, *and it can be identified with a set* $s \subseteq \text{PROP}$. *We say that* p *is true in* s *(write* $s \models p$*) if and only if* $p \in s$.

The uncertainty range of an agent is represented as a set of worlds that the agent considers possible.

Definition 4.1.2. *An* epistemic state *is a set* $S \subseteq \mathcal{P}(\text{PROP})$ *of possible worlds. A pair* (S, s)*, where* S *is an epistemic state and* $s \in S$ *is called a* pointed epistemic state.

With respect to the setting defined in Chapter 2, epistemic states of the agent i associated to the epistemic model $\mathcal{M} = (W, (\sim_i)_{i \in \mathcal{A}}, V)$ are given by equivalence classes in $W/\sim$, in other words, a pointed epistemic state of an agent i is a pair $(\mathcal{K}_i[w], w)$.

In accordance with the semantics of basic epistemic logic, we will interpret the knowledge operator in the usual way.

Definition 4.1.3 (Semantics of $\mathcal{L}_{\mathrm{EL}}$ in epistemic states). *We interpret the single-agent $\mathcal{L}_{\mathrm{EL}}$ in the epistemic states in the following way.*

$$\begin{array}{lll} S, s \models p & \text{iff} & p \in s \\ S, s \models \neg\varphi & \text{iff} & \text{it is not the case that } s \models \varphi \\ S, s \models \varphi \vee \psi & \text{iff} & s \models \varphi \text{ or } s \models \psi \\ S, s \models K\varphi & \textit{iff} & S \subseteq \|\varphi\| \end{array}$$

Accordingly, our simplified approach will be extended to the doxastic framework. By enriching the epistemic state with a plausibility relation we consider epistemic plausibility states. To model beliefs, we need to specify some subset $S_0 \subseteq S$ of the epistemic state, consisting of the possible worlds that are consistent with the agent's beliefs. The intuition here is that although the agent considers all worlds in his epistemic state possible, some of them are seen as more 'desirable', those will be given as the minimal ones according to the plausibility order.

Definition 4.1.4. *A prior plausibility assignment $S \mapsto \leq_S$ assigns to any epistemic state some plausibility order based on the original epistemic state.*

Definition 4.1.5. *A* plausibility state *is a pair $(S, \leq)$ of an epistemic state S and a total preorder $\leq$ on S, called a plausibility relation.*

An epistemic state together with some prior plausibility assignment constitute a plausibility state. Here as in the case of plausibility models we will assume the plausibility relations to be arbitrary total preorders. We will sometimes essentially require their non-well-foundedness.

The language $\mathcal{L}_{\mathrm{DOX}}$ is interpreted on plausibility states in the same way as $\mathcal{L}_{\mathrm{EL}}$. The missing clause of belief is given in the following way:

$$S, \leq, s \models B\varphi \text{ iff } \exists w \leq s\ \forall u \leq w\ u \models \varphi.$$

In the case when $\leq$ is well-founded, the usual definition of 'belief as truth in all the most plausible worlds' holds, i.e., if $(S, \leq)$ is a plausibility state, then for all $s \in S$:

$$S, \leq, s \models B\varphi \text{ iff } \min_{\leq} S \subseteq \|\varphi\|.$$

Our aim now is to reconstruct iterated belief revision in a strict correspondence with identifiability in the limit. We will analyze the epistemic and doxastic properties of limiting learning and the influence of various epistemic attitudes on

the process of convergence. Later we will compare the power of learning in the limit with the capabilities of various belief-revision policies. Before we get to them, we need to set some basic notions of incoming data. As mentioned before, we are interested in (possibly indefinite) iterations—to get the kind of full generality we need to consider infinite streams of information.

Our streams of data consist of chunks of information—every such chunk is a finite set of atomic propositions.

Definition 4.1.6. *A positive non-deterministic data stream is an infinite sequence $\varepsilon = (\varepsilon_1, \varepsilon_2, \ldots)$ of finite sets ε_i of propositions from* Prop.

The intuition is that, at stage i, the agent observes the data in ε_i. The data set $\emptyset$ corresponds to making no observation. For clarity, we will call finite parts of such data streams *data sequences.*

Definition 4.1.7. *A data sequence is a finite sequence $\sigma = (\sigma_1, \ldots, \sigma_n)$, where for every $0 < i \leq n$, σ_i is a finite subset of* Prop.

Besides the usual notation given for texts in Definition 2.1.2, we will also use the concatenation of data sequences.

Definition 4.1.8. *Let σ and π be data sequences. We write $\sigma * \pi$ to denote the concatenation of the two strings, i.e., if $\sigma = (\sigma_1, \ldots, \sigma_n)$ and $\pi = (\pi_1, \ldots \pi_k)$, then $\sigma * \pi = (\sigma_1, \ldots, \sigma_n, \pi_1, \ldots \pi_k)$. For simplicity, if ρ is a finite set of propositions, then $\sigma * \rho = (\sigma_1, \ldots, \sigma_n, \rho)$.*

As explained in Section 2.1, data streams are not entirely arbitrary, they should reflect reality, be consistent with the actual world. The analogy with scientific inquiry can be used here: one can base theories on the results of experiments if the results are assumed to be consistent with reality. This property of data streams will be called 'soundness'.

Definition 4.1.9. *A positive data stream ε is* sound *with respect to world s iff all data in ε are true in s, i.e.,* $\mathrm{set}(\varepsilon) \subseteq s$.

Another restriction on data streams is that, since they are infinite, they should enumerate all elements that are true in the actual world. In other words, if we wait long enough we will see it all. This property of data streams will be called 'completeness'.

Definition 4.1.10. *A positive data stream ε is* complete *with respect to world s iff all the true atomic propositions in s are in ε, i.e., if $s \subseteq \mathrm{set}(\varepsilon)$.*

Throughout the most of this chapter we will assume the data streams for some world s to be sound and complete with respect to s, i.e., we will assume that $s = \mathrm{set}(\varepsilon)$.

In standard learning theory such positive, sound and complete data streams are called 'texts' (see Chapter 2). They are restricted only to the streams in which all the observed data ε_i are either singletons $\{p\}$ (consisting of a positive atom $p \in$ PROP) or $\emptyset$ ('no observation'). The above definitions allow observing more than one atomic fact at a time. In our formalism each piece of information ranges over finite ε_i, and therefore the classical learning theory setting is equivalent to ours.

Until now learning methods have been described generally as ways of converting epistemic states into belief sets in a way dependent on the incoming information. In order to approach the subject of learning as an iterated belief-revision process, we will now turn to the more constructive part of our paradigm—the belief-revision methods themselves. The long-term aim that we have in mind is to define and investigate learning methods that are governed by belief-revision policies.

We define a belief-revision method as a function that, given some data sequence, transforms plausibility states.

Definition 4.1.11. *A* belief-revision method *is a function R that given any plausibility state $(S, \leq)$ and a data sequence $\sigma = (\sigma_1, \ldots, \sigma_n)$ (of any finite length n), outputs a new plausibility state*

$$R((S, \leq), \sigma) := (S^\sigma, \leq^\sigma).$$

Our notion of belief-revision method is more general than the one of classical belief-revision policies. The latter are memory-free, can account by default only for one step of revision. Hence, each time they take only one piece of incoming information. The above definition makes our methods dependent on a finite history of events, but obviously it accounts for the classical policies as a special case.

As in the case of learning methods there are some basic requirements that belief-revision methods might be expected to fulfill. This time most of the properties will be defined in terms of belief operator B, as given in Section 4.1. First we give two versions of data-retention, the property that states that beliefs are expected to reflect the incoming information.

Definition 4.1.12. *A belief-revision method is* weakly data-retentive *if after the revision the most recent piece of data is believed, i.e., for $\sigma = (\sigma_1, \ldots, \sigma_n)$, we have*

$$\text{if } p \in \sigma_n \text{ then } (S^\sigma, \leq^\sigma) \models Bp.$$

Definition 4.1.13. *A belief-revision method is* strongly data-retentive *if all the observed data are believed, i.e., if $\sigma = (\sigma_1, \ldots, \sigma_n)$ then for every $1 \leq i \leq n$:*

$$\text{if } p \in \sigma_i \text{ then } (S^\sigma, \leq^\sigma) \models Bp.$$

In the case of belief-revision methods we can define two types of conservatism. Unlike general learning methods, belief-revision methods output the whole revised

plausibility state. So, conservatism can take a weak form in which the belief itself does not change if the new piece of data has already been believed, or a strong form in which the whole plausibility state does not change under new information, that has been already believed.

Definition 4.1.14. *A belief-revision method is* weakly conservative *if it keeps the same belief when it is confirmed by the new information, i.e., for every finite $\rho \subseteq$ PROP such that $(S^\sigma, \leq^\sigma) \models B(\bigwedge \rho)$ and for every formula θ, we have that:*

$$(S^\sigma, \leq^\sigma) \models B\theta \text{ iff } (S^{\sigma * \rho}, \leq^{\sigma * \rho}) \models B\theta.$$

Definition 4.1.15. *A belief-revision method is* strongly conservative *if it does not change the plausibility state when the new data has already been believed, i.e., for every finite $\rho \subseteq$ PROP s.t. $(S^\sigma, \leq^\sigma) \models B(\bigwedge \rho)$, we have*

$$(S^\sigma, \leq^\sigma) = (S^{\sigma * \rho}, \leq^{\sigma * \rho}).$$

We define the notion of data-drivenness as in the case of learning methods:

Definition 4.1.16. *A belief-revision method is data-driven if it is both weakly data-retentive and weakly conservative.*

As mentioned before, belief-revision methods work on whole plausibility states. This allows a refined notion of keeping track of past events. History-independent belief-revision methods do not distinguish between the same plausibility states that have different pasts.

Definition 4.1.17. *A belief-revision method is* history-independent *if its output at any stage depends only on the previous output and the most recently observed data, i.e., for every finite $\rho \subseteq$ PROP and all finite data sequences σ, π, we have*

$$\text{if } (S^\sigma, \leq^\sigma) = (S^\pi, \leq^\pi) \text{ then } (S^{\sigma * \rho}, \leq^{\sigma * \rho}) = (S^{\pi * \rho}, \leq^{\pi * \rho}).$$

4.2 Iterated DEL-AGM Belief Revision

All revision methods satisfying the AGM postulates are data-driven. This follows from the second AGM postulate: If S is a belief state and $S * \varphi$ represents the set of beliefs resulting from revising S with new belief φ, then φ belongs to $S * \varphi$ (Alchourrón et al., 1985). However, as we will see below, AGM methods are not necessarily strongly data-retentive, nor strongly conservative. Below we will consider three basic qualitative belief-revision methods that met considerable attention within dynamic epistemic logic research: conditioning (update), lexicographic revision (radical upgrade) and minimal revision (conservative upgrade) (for details see Chapter 2). We will investigate the properties of homogeneous iterated revision, i.e., sequences of revisions governed by one particular belief-revision policy.[1]

4.2.1 Conditioning

We want to focus now on the conditioning revision method, which corresponds to *update* in dynamic epistemic logic (see Van Benthem, 2007, and Chapter 2). To briefly recall the notion, update operates by deleting those worlds that do not satisfy all the new data. The minimal requirement for rational application of update is that the incoming information is truthful. We redefine the notion of update for our epistemic states in the following way.

Definition 4.2.1. Conditioning *is a belief-revision method* Cond *that takes an epistemic state S together with a prior plausibility assignment $\leq_S$, i.e., a plausibility state, and a finite set of propositions ρ and outputs a new plausibility state in the following way:*

$$\mathtt{Cond}((S, \leq_S), \rho) = (S^\rho, \leq_S^\rho),$$

where $S^\rho = \{s \in S | s \models \bigwedge \rho\}$, and $\leq_S^\rho = \leq_S \upharpoonright S^\rho$.

The conditioning revision method is obviously weakly data-retentive. Moreover, one can say that it treats the incoming information very seriously—it deletes all worlds inconsistent with it. The deletion cannot be reversed—in this sense conditioning is the ultimate way to memorize things. Below we prove that conditioning is strongly data-retentive.

Proposition 4.2.2. *Conditioning revision method on $(S, \leq_S)$ is strongly data-retentive.*

Proof. Let us take $\sigma = (\sigma_1, \ldots, \sigma_n)$ and assume that $\mathtt{Cond}((S, \leq), \sigma) = (S^\sigma, \leq^\sigma)$. We have to show that the conditioning revision method Cond is strongly data-retentive, i.e., if , for every $1 \leq i \leq n$:

$$\text{if } p \in \sigma_i \text{ then } (S^\sigma, \leq^\sigma) \models Bp.$$

Each time the new information σ_i comes in all worlds that do not satisfy it are eliminated, therefore $S^\sigma = \| \bigwedge \bigcup \sigma \|$. Hence for every world $s \in S^\sigma$, we have that $s \models \bigwedge \bigcup \sigma$. So in the resulting model every proposition that ever occurred in σ is believed. □

Conditioning, being an AGM revision method, is weakly conservative. We will show that it is not strongly conservative.

[1]An alternative, complementary view is to alternate belief-revision policies depending on the status of the incoming information. In such a case the level of doubt (or conservatism) with which one can accept the incoming data depends on the level of the reliability of the incoming information (see Baltag & Smets, 2008a; Van Benthem, 2007). We will not be concerned with such heterogeneous policies, but we view them as an interesting topic for future work.

Proposition 4.2.3. *Conditioning is not strongly conservative.*

Proof. Let us take a sequence of data σ and assume that $\mathtt{Cond}((S, \leq), \sigma) = (S^\sigma, \leq^\sigma)$. We have to show that the conditioning revision method Cond is not strongly conservative, i.e., it is not necessarily the case that it keeps the same plausibility state when the new data is already believed. In other words for every finite $\rho \subseteq \text{PROP}$ such that $(S^\sigma, \leq^\sigma) \models B(\bigwedge \rho)$, we have

$$(S^\sigma, \leq^\sigma) = (S^{\sigma * \rho}, \leq^{\sigma * \rho}).$$

Let us consider the following example. Assume that $S^\sigma = \{\{p, q\}, \{p\}\}$, $\rho = \{q\}$, and the plausibility gives the following order: $\{p, q\} \leq^\sigma \{p\}$. Then clearly

$$(S^\sigma, \leq^\sigma) \models B(\bigwedge \rho).$$

However, after receiving ρ, the revision method Cond will eliminate world $\{p\}$ and therefore:

$$(S^\sigma, \leq^\sigma) \neq (S^{\sigma * \rho}, \leq^{\sigma * \rho}).$$

□

We have shown that apart from being data-driven (weakly data-retentive and weakly conservative) conditioning is strongly data-retentive but not strongly conservative.

Conditioning on Epistemic States Conditioning can change the underlying plausibility order only by deletion of possible worlds. If update is performed on epistemic states that lack plausibility structure, in some cases, while the range of uncertainty of the agent shrinks upon new data, the emergence of full certainty can occur. Conditioning can be considered successful if the actual guess is *finitely identified* (see Chapter 5 and Dégremont & Gierasimczuk, 2009). In this case the iteration on any data stream consistent with any world s allows eliminating uncertainty in a finite number of steps.

4.2.2 Lexicographic Revision

Lexicographic revision corresponds to radical upgrade in dynamic epistemic logic. When facing new information, it does not delete states, it just makes all the worlds satisfying the new piece of data more plausible than all the worlds that do not satisfy it and within the two parts, the old order is kept.

Definition 4.2.4. Lexicographic revision *is a belief-revision method* Lex *that takes an epistemic state S together with a prior plausibility assignment $\leq_S$, i.e., a plausibility state, and a finite set of propositions ρ and outputs a new plausibility state in the following way:*

$$\mathtt{Lex}((S, \leq_S), \rho) = (S, \leq_S^\rho),$$

where for all $t, w \in S$:

$$t \leq_S^\rho w \text{ iff } (t \leq_S^\rho w \text{ or } t \leq_S^{\bar{\rho}} w \text{ or } (t \in \|\bigwedge\rho\| \wedge w \in \|\neg\bigwedge\rho\|)),$$

where: $\leq_S^\rho = \leq_S\upharpoonright\|\bigwedge\rho\|$, *and* $\leq_S^{\bar{\rho}} = \leq_S\upharpoonright\|\neg\bigwedge\rho\|$.

Lexicographic revision is not strongly data-retentive on arbitrary sequences of data. However, if the data sequence is sound with respect to a world in the epistemic state, strong data retention holds. Moreover, this type of revision is not strongly conservative. Let us go through the arguments for each case.

Proposition 4.2.5. *Lexicographic revision is not strongly data-retentive on arbitrary data streams.*

Proof. Let us take a finite sequence of data $\sigma = (\sigma_1, \ldots, \sigma_n)$ and assume that $\texttt{Lex}((S, \leq), \sigma) = (S, \leq^\sigma)$. We have to show that the lexicographic revision method is not strongly data-retentive, i.e. it is not the case that for every $1 \leq i \leq n$:

$$\text{if } p \in \sigma_i \text{ then } (S, \leq^\sigma) \models Bp.$$

Let us take $S = \{\{p\}, \{q\}\}$, $\sigma = (\{p\}, \{q\})$, and assume any initial ordering on S, e.g., $\{p\} \leq \{q\}$. First $\sigma_1 = \{p\}$ comes in, and p starts to be believed. After receiving $\sigma_2 = \{q\}$ the most plausible state becomes $\{q\}$, so p is no longer believed, i.e., $\neg B \bigwedge \sigma_1$. □

Observe that σ in the above proof is not sound and complete with respect to any possible world in S. Therefore, in the learning-theoretic setting that we described in Chapter 2, σ cannot possibly appear in the first place. In this sense the lexicographic revision is especially well suited to learning and scientific inquiry—it's behaviour improves on data streams that are assumed to be consistent with reality. Let us see that it is so.

Proposition 4.2.6. *Lexicographic revision method on $(S, \leq_S)$ is strongly data-retentive on data sequences that are sound with respect to some $s \in S$.*

Proof. We have to show that the lexicographic revision method $\texttt{Lex}$ is strongly data-retentive on sound data sequences. Let us take a plausibility state $(S, \leq_S)$, $s \in S$ and σ—a data sequence that is sound with respect to s, i.e., $\text{set}(\sigma) \subseteq s$. After reading σ, for all the worlds t that are most plausible with respect to $\leq_S$ in S it is the case that $\|\bigwedge\bigcup\sigma\| \subseteq t$, $t \models B\bigwedge\bigcup\sigma$. It is so because by assumption there is at least one such world, s. □

While adhering to the desired form of strong data-retention, lexicographic revision is not strongly conservative.

Proposition 4.2.7. *The lexicographic revision is not strongly conservative.*

Proof. Let us take a sequence of data σ and assume that $\texttt{Lex}((S, \leq), \sigma) = (S, \leq^\sigma)$. We have to show that the lexicographic revision method $\texttt{Lex}$ is not strongly conservative, i.e., it is not necessarily the case that it keeps the same plausibility state when the new data is already believed. Formally, for every finite $\rho \subseteq \text{PROP}$ such that $(S^\sigma, \leq^\sigma) \models B(\bigwedge\rho)$, we have

$$(S, \leq^\sigma) = (S, \leq^{\sigma*\rho}).$$

Let us consider the following example. Assume that $S = \{\{p, q\}, \{p\}, \{q\}\}$, $\rho = \{q\}$, and the plausibility gives the following order: $\{p, q\} \leq^\sigma \{p\} \leq^\sigma \{q\}$. Then clearly $(S, \leq^\sigma) \models B(\bigwedge\rho)$. However, after getting ρ, the revision method will put world $\{q\}$ to be more plausible than $\{p\}$, and therefore

$$(S, \leq^\sigma) \neq (S, \leq^{\sigma*\rho}).$$

□

Summary Lexicographic revision is strongly data-retentive on streams that are sound with respect to some possible world. On the other hand, it is not strongly conservative.

4.2.3 Minimal Revision

The minimal revision method corresponds to conservative upgrade in dynamic epistemic logic (see Van Benthem, 2007). The most plausible worlds satisfying the new data become the most plausible overall. In the remaining part the old order is kept.

Definition 4.2.8. *Minimal revision is a belief-revision method* $\texttt{Mini}$ *that takes an epistemic state S together with a prior plausibility assignment $\leq_S$ (i.e., a plausibility state) and a finite set of propositions ρ and outputs a new plausibility state in the following way:*

$$\texttt{Mini}((S, \leq_S), \rho) = (S, \leq_S^\rho),$$

where for all $t, w \in S$:

$$t \leq_S^\rho w \text{ iff } t \leq_S^{\text{rest}\rho} w \text{ or } t \in \min_{\leq_S}(S, \leq_S)$$

where: $\leq_S^{\text{rest}\rho} = \leq_S\upharpoonright\{t \in S \mid t \notin \min_{\leq_S}\|\bigwedge\rho\|\}$.

The minimal revision method is not strongly data-retentive (not even on sound data streams). But, on the other hand, it is strongly conservative.

Proposition 4.2.9. *Minimal revision on $(S, \leq_S)$ is not strongly data-retentive on all data sequences that are sound with respect to some $s \in S$.*

Proof. Let us take a sequence of data $\sigma = (\sigma_1, \ldots, \sigma_n)$ and assume that $\texttt{Mini}((S, \leq), \sigma) = (S, \leq^\sigma)$. We have to show that Mini is not strongly data-retentive, i.e., is not the case that for every $1 \leq i \leq n$:

$$\text{if } p \in \sigma_i \text{ then } (S, \leq^\sigma) \models Bp.$$

Let us take $S = \{\{p\}, \{q\}, \{p, q\}\}$, $\sigma = (\{p\}, \{q\})$ a data sequence consistent with world $\{p, q\}$, and assume that the initial ordering on S is $\{q\} \leq \{p\} \leq \{p, q\}$. After receiving $\sigma_1 = \{p\}$ the plausibility ordering becomes $\{p\} \leq^{\sigma_1} \{q\} \leq^{\sigma_1} \{p, q\}$. Then $\sigma_2 = \{q\}$ comes in—now our method gives the ordering $\{q\} \leq^{(\sigma_1, \sigma_2)} \{p\} \leq^{(\sigma_1, \sigma_2)} \{p, q\}$. So p is no longer believed although it was included in σ_1, i.e., after the second piece of data $\neg B(\bigwedge \sigma_1)$. □

Proposition 4.2.10. *Minimal revision is strongly conservative.*

Proof. Let us take a sequence of data σ and assume that $\texttt{Mini}((S, \leq), \sigma) = (S, \leq^\sigma)$. We have to show that the minimal revision method is strongly conservative, i.e., it keeps the same plausibility state when the new data is already believed. Formally, for every finite $\rho \subseteq$ Prop such that $(S, \leq^\sigma) \models B(\bigwedge \rho)$, we have

$$(S, \leq^\sigma) = (S, \leq^{\sigma * \rho}).$$

Let us take $\rho \subseteq$ Prop such that $(S, \leq^\sigma) \models B(\bigwedge \rho)$, we have to show that

$$(S, \leq^\sigma) = (S, \leq^{\sigma * \rho}).$$

Let us assume that $(S, \leq^\sigma) \neq (S, \leq^{\sigma * \rho})$. This means that after receiving ρ the plausibility order has been rearranged. By the definition of Mini, this could happen only in the case when among the most plausible in $(S, \leq^\sigma)$ there was no world t such that $t \in \|\rho\|$. But then also $(S, \leq^\sigma) \not\models B(\bigwedge \rho)$. Contradiction. □

The precise relation between the minimal revision method and the notion of conservatism is an interesting subject of further investigation. Our definition of strong conservatism indicates that minimal revision is the only strongly conservative belief-revision method. Hence, the concepts of minimal revision and strongly conservative revision are equivalent.

4.3 Learning Methods

In learning theory the learner is taken to be a function that on each finite sequence of data outputs a conjecture—a hypothesis from the initially given set of possibilities. We will follow this intuition here—our learning is performed by a

function that, given some initial uncertainty range and a data sequence outputs a belief set. The latter contains worlds that are in some way consistent with the received sequence. In other words, a learning method is a way of converting any epistemic state, S, into a *belief set* $S_0 \subseteq S$ on the basis of any given data sequence.[2]

Definition 4.3.1. *A learning method is a function L that assigns some belief set $L(S, (\sigma_1, \ldots, \sigma_n)) \subseteq S$ to an epistemic state S (of any epistemic model) and a data sequence $\sigma = (\sigma_1, \ldots, \sigma_n)$ (of any finite length n).*

In contrast to our setting, in learning theory the learning functions are assumed to be deterministic, i.e., $L(S, \sigma_1, \ldots, \sigma_n) \subseteq S$ is either a singleton $\{h\}$ (for some $h \in S$, called the current hypothesis) or $\emptyset$ (when no conjecture is made).

In principle learning methods can be completely arbitrary. However, (to keep learners reasonable) we might well want them to satisfy certain requirements. Let us now define and briefly describe such basic properties of learning methods.

The first feature of learning methods is their data-retention. It formalizes the intuition that new beliefs should be consistent with the information received. Minimal requirement of this kind is that the belief set accounts for the most recent piece of data.

Definition 4.3.2. *A learning method is* weakly data-retentive *if for all data sequences $\sigma = (\sigma_1, \ldots, \sigma_n)$, we have that if $L(S, \sigma) \neq \emptyset$ then:*

$$\sigma_n \subseteq \bigcap L(S, \sigma).$$

The maximal requirement of data-retention is that the current conjecture always accounts for all data that have been encountered.

Definition 4.3.3. *A learning method is* strongly data-retentive *if for all data sequences $\sigma = (\sigma_1, \ldots, \sigma_n)$ and for every $1 \leq i \leq n$, we have that if $L(S, \sigma) \neq \emptyset$, then:*

$$\sigma_i \subseteq \bigcap L(S, \sigma).$$

In formal learning theory the corresponding property is known under the name of consistent learning.

Another quite intuitive assumption is that agents change their beliefs only when the observed information clearly contradicts their current beliefs. In other words, unless the agent is forced to, he does not change his mind.

[2]Learning functions output 'absolute beliefs'. In this respect they seem to be closer to the AGM-style belief revision than to DEL operations which account for 'conditional beliefs'. In fact, this analogy is not full, because learning functions are allowed to base their conjectures on whole sequences of events.

Definition 4.3.4. *A learning method is* weakly conservative *if for all data sequences $\sigma = (\sigma_1, \ldots, \sigma_n)$ and a finite $\rho \subseteq$ PROP, we have:*

$$\textit{if } \rho \subseteq \bigcap L(S, \sigma) \textit{ then } L(S, \sigma) = L(S, \sigma * \rho).$$

The analogous concepts: conservatism, in learning theory has been shown to restrict the class of languages identifiable in the limit, as has been consistency (see, e.g., Jain et al., 1999). We will not go into details of these arguments here. Let us just mention that our concept of learning method is different from that of the learning function in formal learning theory. These are assumed to output integers that are indices of sets in the initial class. Our learning methods are working directly on sets and output the entire set corresponding to current beliefs. In this respect our approach is more 'semantic' and accordingly learning methods may turn out to be more powerful.

For brevity's sake we will sometimes combine the two weak conditions of conservatism and data-retention together under the name of data-drivenness.

Definition 4.3.5. *A learning method is* data-driven *if it is both weakly data-retentive and weakly conservative.*

Last but not least, an important aspect of learning methods is their memory concerning past conjectures. Below we define the limit case of a memory-free learning method.

Definition 4.3.6. *A learning method is* memory-free *if, at each stage, the new belief set depends only on the previous belief set and the new data, i.e., for any finite $\rho \subseteq$ PROP:*

$$\text{if } L(S, \sigma) = L(S', \sigma') \text{ then } L(S, \sigma * \rho) = L(S', \sigma' * \rho).$$

A condition analogous to this in Definition 4.3.6 has been considered in formal learning theory and is known under the name of memory limitations (see, e.g., Jain et al., 1999).

4.4 Belief-Revision-Based Learning Methods

Finally, we are ready to put all the pieces together and describe learning that is based on belief-revision methods. We build a learning method from a belief-revision strategy in the following way. We take an epistemic state, put some plausibility order on it, and simulate a certain belief-revision method while receiving new information. The answer of the learning method each time consists of the most plausible worlds in the plausibility state. Such a learning method still outputs just the belief states but it bases its conjectures on the constructive work executed in the background by the belief-revision method.

Definition 4.4.1. *A belief-revision method R, together with a prior plausibility assignment $S \mapsto \leq_S$, generates a learning method L_R, called a* belief-revision-based learning method, *and given by:*

$$L_R(S, \sigma) := \min_{\leq_S} R(S, \leq_S, \sigma),$$

where $\min_{\leq'}(S', \leq')$ is the set of all the least elements of S' with respect to $\leq'$ (if such least elements exist) or $\emptyset$ (otherwise).

Now, an interesting set of questions arises. Is it the case that data-retention and conservatism of belief-revision method is inherited by the corresponding belief-revision-based learning methods? Do history-independent belief-revision methods generate memory-free learning methods? Below we list and discuss several dependencies between belief-revision methods and learning methods generated from them.

Proposition 4.4.2. *If a belief-revision method R is weakly data-retentive then the generated learning method L_R is weakly data-retentive.*

Proof. Let us take a belief-revision method R and some epistemic state together with a prior plausibility assignment $(S, \leq_S)$. Assume that R is weakly data-retentive, i.e., if $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a data sequence then:

$$\forall p \in \sigma_n \ (S^\sigma, \leq_S^\sigma) \models Bp.$$

We need to show that if $L_R(S, \sigma) \neq \emptyset$, then $\sigma_n \subseteq \bigcap L_R(S, \sigma)$. Let us then assume that $L_R(S, \sigma) \neq \emptyset$, i.e., there is a $\leq_S^\sigma$ minimal element in S^σ. Then in every world minimal with respect to $\leq_S$ every p from σ_n holds:

$$\forall p \in \sigma_n \ \min_{\leq_S^\sigma}(S^\sigma, \leq_S^\sigma) \subseteq \|p\|,$$

where $\|p\|$ stands for the set of possible worlds that include p. Therefore, in every minimal world the conjunction of the σ_n holds:

$$\min_{\leq_S^\sigma}(S^\sigma, \leq_S^\sigma) \subseteq \| \bigwedge \sigma_n \|,$$

or equivalently:

$$\sigma_n \subseteq \bigcap \min_{\leq_S^\sigma}(S^\sigma, \leq_S^\sigma).$$

Since $(S^\sigma, \leq_S^\sigma) = R((S, \leq_S), \sigma) = L_R(S, \sigma)$, we have that

$$\sigma_n \subseteq \bigcap L_R(S, \sigma).$$

□

The next two propositions are proved in a similar way.

Proposition 4.4.3. *If R is data-retentive then the induced learning method L_R is data-retentive.*

Proposition 4.4.4. *If a belief-revision method R is weakly conservative then the induced learning method L_R is weakly conservative.*

It remains to show how belief-revision and learning methods relate to each other with respect to their memory limitations.

Proposition 4.4.5. *A learning method generated from a history-independent belief-revision method does not have to be memory-free.*

Proof. We prove this proposition by showing an example—a belief-revision method that is history-independent but the learning method that it induces is not memory-free. Let R be the lexicographic revision method (that corresponds to *lexicographic upgrade* in DEL, see Chapter 2), all the worlds satisfying the new data become more plausible than all the worlds not satisfying them; and within the two zones, the old order is kept. R is clearly history-independent. Each time the revision takes into account only the last output in the form of an epistemic plausibility state and the new incoming information. To see that L_R is not memory-free consider the following two plausibility orders on $S = S' = \{\{p\},\{q\},\{p,q\}\}$. Assume that for some σ and σ':

1. $R((S,\leq_S),\sigma)$ gives the plausibility order: $\{p\} <_S \{p,q\} <_S \{q\}$;
2. $R((S',\leq'_S),\sigma')$ gives the plausibility order: $\{p\} <_{S'} \{q\} <_{S'} \{p,q\}$.

It is easy to observe that $L_R(S,\sigma) = L_R(S',\sigma')$. Assume now that the next observation $\rho = \{q\}$. Then clearly $L_R(S,\sigma * \rho) = \{p,q\}$, while $L_R(S',\sigma' * \rho) = \{q\}$. Therefore, for the belief-revision method R there is a data sequence ρ such that:

$$L_R(S,\sigma) = L_R(S',\sigma'), \text{ but } L_R(S,\sigma * \rho) \neq L_R(S',\sigma' * \rho).$$

□

Summary Let us briefly summarize the results we have obtained so far. Data-retention and weak conservatism are preserved when a learning method is generated from a belief-revision method. However history-free belief-revision methods are still able to remember more than just the last conjecture of the generated learning method. This is so, because they 'keep' the whole plausibility order for 'further use'.

4.5 Convergence

What does it mean for a learning method to be *reliable* with respect to the initial epistemic state S? It means that it is possible to rely upon it to find the real world in finite time, no matter what the real world is, as long as it belongs to the given initial epistemic state S and as long as the data stream is sound and complete (for a discussion of reliability in belief-revision see Kelly et al., 1995). In this section we investigate reliability with respect to convergence to the correct belief. The expected result is not knowledge understood as full certainty, but rather a kind of belief that is guaranteed to persist under true information. In this setting, an agent can be right in believing something but he might not know it.

Identification in the limit guarantees the *convergence* to the right hypothesis, i.e., at a finite stage the answers of the learning method *stabilize* on the correct conjecture.[3]

Definition 4.5.1. *An epistemic state S is* identified in the limit *on positive data by learning method L if and only if for every world $s \in S$ and every sound and complete positive data stream for s, there exists a finite stage after which L outputs the singleton $\{s\}$ from then on.*[4]

In general we can attribute identifiability to the epistemic states by requiring that there is a learning method that identifies the state.

Definition 4.5.2. *An epistemic state is* identifiable in the limit (resp. finitely identifiable) *on positive data if there exists a learning method that can identify it in the limit (resp. finitely identify it) on positive data.*

Particular learning methods differ in their power. The most powerful among them are those that are universal, i.e., they can identify in the limit every class identifiable in the limit.

Definition 4.5.3. *A learning method L is* universal *on positive data if and only if it can identify in the limit on positive data every epistemic state that is identifiable in the limit.*

We are especially interested in learning methods that are generated from belief-revision policies. For brevity's sake we will use the notion of identification in the limit while talking about belief-revision policies. By a belief-revision method identifying S in the limit, we mean that the belief-revision method together with some prior plausibility assignment generates a learning method that identifies S in the limit (as given in Definition 4.5.1).

[3] In this chapter we will focus on identification in the limit. Finite identification is investigated in the context of epistemic logic in Chapter 5.

[4] In terms of belief, it means that the agent's conjectures stabilize to the complete true belief about the actual world.

Definition 4.5.4. *An epistemic state S is identified in the limit on positive data by a belief-revision method R if there exists a prior plausibility assignment $S \mapsto \leq_S$ such that the generated belief-revision-based learning method L_R identifies S in the limit on positive data.*

The above definition requires the existence of an appropriate initial plausibility assignment. In principle it can be a completely arbitrary preorder. However, we might want this prior plausibility assignment to satisfy certain assumptions of cognitive realism or rationality. The properties that are often required of such priors are *well-foundedness* and *totality*. Well-foundedness assures that a minimal state is always exists and it is possible to point to it as to the current belief. Totality guarantees that whenever two possibilities are considered, they are comparable with respect to the plausibility assignment. With respect to identifiability in the limit one can accordingly demand a prior plausibility assignment to satisfy those as standard assumptions of preference relations in doxastic epistemic logic (see, e.g., Dégremont, 2010).

Definition 4.5.5. *An epistemic state S is* standardly identified in the limit *on positive data by a belief-revision method R if there exists a (total) well-founded prior plausibility assignment $S \mapsto \leq_S$ such that the induced belief-revision-based learning method L identifies S in the limit on positive data.*

We define the analogous notion of universality for standard identifiability.

Definition 4.5.6. *A revision method is* standardly universal *on positive data if it can standardly identify in the limit on positive data every epistemic state that is identifiable.*

Our aim now is to show that some of the DEL-AGM revision methods generate a universal learning methods. The main technical difficulty of this part is the construction of the appropriate prior plausibility order. To define it we will use the concepts of locking sequences introduced by Blum & Blum (1975) and finite tell-tale sets proposed by Angluin (1980). For the latter we will use the simple non-computable version. We will refine the classical notion of finite tell-tales and use it in the construction of the suitable prior plausibility assignment that, together with conditioning and lexicographic revision, will generate universal learning method.

The first observation is that if convergence occurs, then there is a finite sequence of data that 'locks' the corresponding sequence of conjectures on a correct answer. This finite sequence is called a 'locking sequence'.

Definition 4.5.7 (Blum & Blum 1975). *Let an epistemic state S, a possible world $s \in S$, a learning method L and a finite data sequence of propositions, σ, be given. The sequence σ is called a* locking sequence *for s and L if:*

1. $\text{set}(\sigma) \subseteq s$;
2. $L(S, \sigma) = \{s\}$;
3. *for any data sequence α, if* $\text{set}(\alpha) \subseteq s$, *then* $L(S,\sigma) = L(S, \sigma * \alpha)$.

Lemma 4.5.8 (Blum & Blum 1975). *If a learning method L identifies possible world s in the limit then there exists a locking sequence for L on s.*

The characterization of identifiability in the limit (see Theorem 2.1.14) can be generalized to account for arbitrary classes, by dropping the assumption of computability. It requires the existence of finite sets that allow drawing a conclusion without the risk of overgeneralization.

Lemma 4.5.9 (Angluin 1980). *Let S be an epistemic state over a set* PROP *of atomic sentences, such that* PROP *and S are at most countable. If S is identifiable in the limit on positive data, then there exists a total map $D : S \to \mathcal{P}^{<\omega}(\text{PROP})$, given by $s \mapsto D_s$, such that D_s is a finite tell-tale for s, i.e.,*

1. *D_s is finite,*
2. *$D_s \subseteq s$,*
3. *if $D_s \subseteq t \subseteq s$ then $t = s$.*

Proof. Let S be an epistemic state over a set PROP of atomic sentences, such that PROP and S are at most countable. Let us also assume that S is identifiable in the limit on positive data by the learning method L, i.e., for every world $s \in S$ and every sound and complete positive data stream for s, there exists a finite stage after which L outputs the singleton $\{s\}$ from then on. By Lemma 4.5.8, for every $s \in S$ we can take a locking sequence σ_s for L on s. For any $s \in S$ we define $D_s := \text{set}(\sigma_s)$.

1. D_s is finite because locking sequences are finite.
2. $D_s \subseteq s$, because $\text{set}(\sigma_s) \subseteq S$.
3. if $D_s \subseteq t \subseteq s$ then $t = s$. Assume that there are $s, t \in S$, such that $s \neq t$ and $D_s \subseteq t \subsetneq s$. Let us take a positive sound and complete data stream ε for t, such that for some $n \in \mathbb{N}$, $\varepsilon \restriction n = \sigma_s$. Because σ_s is a locking sequence for L on s and $\text{set}(\varepsilon) = t \subseteq s$, L converges to s on ε. Therefore, L does not identify t, a state from S. Contradiction.

This concludes the proof. □

We will use the notion of finite 'tell-tale' to construct an ordering of S. The aim is to find a way of assigning the prior plausibility order that allows reliable belief revision. We will base the construction on finite tell-tales, but we will introduce one additional condition (see point 4 of Definition 4.5.10, below).

Definition 4.5.10. *Let S be a countable epistemic state with an injective map $i : S \to \mathbb{N}$, and D' be a total map such that $D' : S \to \mathcal{P}^{<\omega}(\text{PROP})$, given by $s \mapsto D'_s$ having the following properties:*

1. *D'_s is finite,*
2. *$D'_s \subseteq s$,*
3. *if $D'_s \subseteq t \subseteq s$ then $t = s$,*
4. *if $D'_s \subseteq t$ but $s \not\subseteq t$ then $i(s) < i(t)$.*

We call D' an ordering tell-tale map, *and D'_s an* ordering tell-tale set *of s.*

Definition 4.5.11. *For $s, t \in S$, we put*

$$s \preceq_{D'} t \text{ if and only if } D'_s \subseteq t.$$

We take $\leq_{D'}$ to be the transitive closure of the relation $\preceq_{D'}$.

Lemma 4.5.12. *For any identifiable epistemic state S and any ordering tell-tale map D', the relation $\leq_{D'}$ is an order, i.e., $\leq_{D'}$ is reflexive, transitive and antisymmetric.*[5]

Before we give the proof let us introduce the notion of a proper cycle in $\leq_{D'}$.

Definition 4.5.13. *A* proper cycle *in $\leq_{D'}$ is a sequence of worlds $s_1, \ldots, s_n$, with $n \geq 2$, and such that:*

1. *D'_{s_i} is included in s_{i+1} (for all $i = 1, \ldots, n-1$).*
2. *$s_1 = s_n$, but*
3. *$s_1 \neq s_2$.*

Proof. The fact that $\leq_{D'}$ is a preorder is trivial: reflexivity follows from the fact that D'_s is always included in s, and transitivity is imposed by construction (by taking the transitive closure). We need to prove that $\leq_{D'}$ is antisymmetric. In order to do that we will show (by induction on n) that $\leq_{D'}$ does not contain proper cycles of any length $n \geq 2$.

1. For the initial step ($n = 2$): Suppose we have a proper cycle of length 2. As we saw, this means that there exist two states s_1, s_2 such that $s_1 \neq s_2$, D'_{s_1} is included in s_2 and D'_{s_2} is included in s_1. There are three cases:

[5] We use D', to distinguish from the original tell-tale function D (satisfying only the conditions of Angluin's Theorem).

Case 1: s_1 is included in s_2. In this case, D'_{s_2} is included in s_1, and s_1 is included in s_2, so (by Condition 3 of Definition 4.5.10), we have that $s_1 = s_2$. Contradiction.

Case 2: s_2 is included in s_1. This case is similar: D'_{s_1} is included in s_2 and s_2 is included in s_1, so (by Condition 3 of Definition 4.5.10), we have that $s_2 = s_1$. Contradiction.

Case 3: s_1 is not included in s_2, and s_2 is not included in s_1. In this case, from the assumption that D'_{s_1} is included in s_2, and that s_1 is not included in s_2, we can infer (by Condition 4 of Definition 4.5.10), that $i(s_1) < i(s_2)$. But, in a completely similar manner (from D'_{s_2} included in s_1, and s_2 not included in s_1), we can also infer that $i(s_2) < i(s_1)$. Putting these together, we get $i(s_1) < i(s_2) < i(s_1)$. Contradiction.

2. For the inductive step ($n + 1$): Suppose $s_1 . s_2 . \ldots . s_{n+1}$ is a proper cycle of length $n + 1$. We consider two cases:

Case 1: There exists k with $1 \leq k < n$ such that s_k is included in s_{k+1}. If $1 < k$, then it is easy to see that the sequence $s_1, \ldots, s_{k-1}, s_{k+1}, \ldots$ (obtained by deleting s_k from the above proper cycle of length $n + 1$) is also a proper cycle, but of smaller length (n). Contradiction. Similarly, if $k = 1$, it is easy to see that the sequence $s_1, s_3, \ldots, s_{n+1}$ (obtained by deleting s_2) is a proper cycle of smaller length (n). Contradiction.

Case 2: s_k is not included in s_{k+1} for any $1 \leq k < n$. In this case, we have that for all $1 \leq k < n$, D'_{s_k} is included in s_{k+1} but s_k is not included in s_{k+1}. By Condition 4 of Definition 4.5.10, it follows that we have $i_{s_k} < i_{s_{k+1}}$, for all $k = 1, \ldots, n$. By the transitivity of $\leq_{D'}$, it follows that $i_{s_1} < i_{s_{n+1}}$. But by Condition 2 of Definition 4.5.10, $s_1 = s_{n+1}$, hence $i_{s_1} > i_{s_{n+1}}$. Contradiction.

□

We will now show that $\leq_{D'}$, used by the conditioning revision method, guarantees convergence to the right belief whenever the underlaying epistemic state is identifiable in the limit.

Theorem 4.5.14. *The conditioning-based learning method is universal on positive data.*

Proof. We have to show that an epistemic model S is identifiable in the limit iff S is identifiable in the limit by the conditioning-based learning method. Obviously, if S is identifiable in the limit by the conditioning-based learning method, then S is identifiable in the limit. We will therefore focus on the other direction, i.e., we will show that if S is identifiable in the limit by any learning method, then it is identifiable in the limit by the conditioning-based learning method.

First let us assume that S, an epistemic state, is identifiable in the limit and hence it is at most countable. Let us then take an injective map $i : S \to \mathbb{N}$. By Lemma 4.5.9 we can assume the map D that gives tell-tales for any $s \in S$. On the basis of D we will now construct a new map $D' : S \to \mathcal{P}^{\leq\omega}(\text{PROP})$. We will proceed step by step according to the enumeration of S given by i.

1. For s_1 we set $D'(s_1) := D(s_1)$.
2. For s_n: For every $k < n$ such that $D_{s_n} \subseteq s_k$ and $s_n \not\subseteq s_k$, we choose an atomic proposition p_k such that $p_k \in s_n$ and $p_k \notin s_k$. We define *Rest* in the following way.

$$Rest := \{p_k \mid k < n \;\&\; p_k \in s_n \;\&\; p_k \notin s_k \;\&\; D_{s_n} \subseteq s_k \;\&\; s_n \not\subseteq s_k\}.$$

Then, we set $D'_{s_n} = D_{s_n} \cup Rest$.

We have to check if D' satisfies conditions of Definition 4.5.10.

1. D'_s is finite, because D_s and *Rest* are both finite.
2. $D'_s \subseteq s$, because D_s and *Rest* are subsets of s.
3. If $D'_s \subseteq t \subseteq s$ then $t = s$, because then $D_s \subseteq D'_s \subseteq t \subseteq s$, and hence, by the definition the finite tell-tale set $t = s$.

What remains is to check the condition 4: If $D'(s) \subseteq t$ and $s \not\subseteq t$ then $i(s) < i(t)$. Let us assume the contrary: $D'(s) \subseteq t$ and $s \not\subseteq t$ and $i(t) \leq i(s)$. There are two possibilities:

1. $i(t) = i(s)$, but then $s = t$ and hence $s \subseteq t$. Contradiction.
2. $i(t) < i(s)$. Then, there is a proposition $p \in D'(s)$ such that $p \in s$ and $p \notin t$. Therefore, by the inductive step of the construction of D', $D'(s) \not\subseteq t$. Contradiction.

We now have that D' satisfies all conditions of Definition 4.5.10, and therefore it is an ordering tell-tale map. Hence, by Lemma 4.5.12, the corresponding $\leq_{D'}$ is an order.

It remains to show that S is identifiable in the limit by the learning method generated from the conditioning belief-revision method and the prior plausibility assignment $\leq_{D'}$. Let us then take any $s \in S$ and the corresponding $D'(s)$. Since $D'(s) \subseteq s$ for every ε—a sound and complete positive data stream for s, there is $n \in \mathbb{N}$ for which $D'(s) \subseteq \text{set}(\varepsilon{\restriction}n)$. Our aim is now to demonstrate that after receiving the elements of $\varepsilon{\restriction}n$, s is the minimal element in $S^{\varepsilon\restriction n}$ with respect to $\leq_{D'}$. Let us assume for contradiction that it is not, i.e., there is $t \in S^{\varepsilon\restriction n}$ such that $t \neq s$ and $t \leq_{D'} s$. Since $t \in S^{\varepsilon\restriction n}$, we get that $D'(s) \subseteq t$, but then, by Definition

4.5.10, $s \leq_{D'} t$. And because by Lemma 4.5.12, $\leq_{D'}$ is antisymmetric we get that $s = t$. Contradiction.

To see that the conditioning process stabilizes on $\{s\}$, it is enough to observe that ε is sound with respect to s, and therefore no further information from ε can eliminate s (because conditioning is weakly conservative). So for any $k > n$, $\min_{\leq_{D'}} \text{Cond}((S, \leq'_D), \varepsilon{\restriction}k) = \{s\}$. □

Theorem 4.5.15. *The lexicographic belief-revision method is universal on positive data streams.*

The proof is analogous to the proof of Theorem 4.5.14. As far as simple beliefs are concerned, radical upgrades with true information do exactly what updates do. The only difference is that the rest of the doxastic structure might not stabilize, but only the minimal elements stabilize (on worlds indistinguishable from the real one).

The preorder defined in the proof of Theorem 4.5.14 is not necessarily well-founded. It is impossible to improve on this without losing the universality property. This is why as background setting we need generalized plausibility models, in which the plausibility is a preorder, without assuming well-foundedness. Belief can still be defined as 'truth in all the states that are plausible enough' (this requires three quantifiers: For every state s there exists some state $t \leq s$ such that φ is true in all states $w \leq t$). This is equivalent to the standard definition in the case that there exist minimal states (i.e., states $\leq$ than all others).

Let us now turn to the negative result concerning the minimal revision method.

Proposition 4.5.16. *Minimal revision is not universal.*

Proof. Let us give a counter-example, an epistemic state that is identifiable in the limit, but is not identifiable by the minimal revision method:

$$S = \{\{p\}, \{q\}, \{p, q\}\}.$$

The epistemic state S is identifiable in the limit by the conditioning revision method: just assume the ordering $\{p\} < \{q\} < \{p, q\}$. However, there is no ordering that will allow identification in the limit of S by the minimal revision method. If $\{p, q\}$ occurs in the ordering before $\{p\}$ (or before $\{q\}$), then the minimal revision method fails to identify $\{p\}$ ($\{q\}$, respectively). If both $\{p\}$ and $\{q\}$ precede $\{p, q\}$ in the ordering then the minimal revision method fails to identify $\{p, q\}$ on any data stream consisting of singletons of propositions from $\{p, q\}$. On all such data streams for $\{p, q\}$ the minimal state will alternate between $\{p\}$ and $\{q\}$, or stabilize on one of them. The last case is that at least one of $\{p\}$ and $\{q\}$ is equiplausibile to $\{p, q\}$. In such case $\{p, q\}$ is not identifiable because for any single proposition from $\{p, q\}$ there is more than one possible world consistent with it. □

Proposition 4.5.17. *There is no standardly universal belief-revision method.*

Proof. There is an epistemic state S that is identifiable in the limit by a learning method, but is not standardly identified in the limit by any belief-revision method, i.e., there is no belief-revision method that would, together with a well-founded order $\leq$ generate a learning method that identifies S in the limit. The following epistemic state constitutes such counter-example:

$$S = \{s_n = \{p_k \mid k \geq n\} \mid n \in \mathbb{N}\}.$$

S is identifiable in the limit by learning method L, that is defined in the following way:

$$L(S, \sigma) = s_n \text{ iff } n \text{ is the smallest such that } \mathrm{set}(\sigma) \subseteq s_n.$$

Moreover, S is identifiable in the limit by a revision-based learning method. We take the conditioning revision method and $\leq \subseteq S \times S$ defined in the following way: For any $s_n, s_m \in S$, $s_n \leq s_m$ iff $n \geq m$. It is easy to observe that $\leq$ is not well-founded.

Let us now assume that S is standardly identifiable in the limit, i.e., there is a belief-revision method R and a well-founded order $\leq$ on S, such that the learning method generated from R and $\leq$ identifies S in the limit. If $\leq$ is well-founded we can choose the $\leq$-minimal element. Let as assume that it is s_k for some $k \in \mathbb{N}$. Obviously, for all $n > k$, $s_n \subseteq s_k$, in fact there are infinitely many $n \in \mathbb{N}$ such that $s_k \leq s_n$ and $s_n \subseteq s_k$. Therefore, all positive, sound and complete data steams for such s_n are also sound with respect to s_k. If we accept the minimal assumption of data-drivenness of belief-revision methods, we can easily see that R will not change the $\leq$-minimal state from s_k to any of s_n, for any sound and complete data streams for s_n. Therefore R fails to identify in the limit s_n, for all $n > k$. □

Summary In this section we considered a notion of reliability of a belief-revision method. We used the concept of identifiability in the limit to define success of an iterated belief-revision process. We have shown that some belief-revision methods are universal, i.e., they identify in the limit all epistemic states that are identifiable by arbitrary learning methods. Such very powerful learning methods are generated from conditioning and lexicographic revision (update and conservative (radical) upgrade in dynamic epistemic logic). More conservative methods turn out not to be universal. This indicates the existence of a tension between learning power and conservatism. We can see that the weakness of the minimal revision method lies in ignoring information that is already believed. Universal belief-revision methods perform operations on plausibility states even if they do not influence the current beliefs immediately. These operations pay off as the process continues.

4.6 Learning from Positive and Negative Data

We will now extend our framework to account for revising with negation. Let us consider the stream ε that consists of both positive and negative data:[6]

$$\mathrm{set}(\varepsilon) \subseteq \text{PROP} \cup \{\overline{p} \mid p \in \text{PROP}\}.$$

All notions defined in Sections 4.1 and 4.5 (soundness and completeness of a stream, identifiability in the limit, universality, etc.) are analogous for this case.

Let us recall the definition of the epistemic state, together with the additional explanation how to interpret the negative information.

Definition 4.6.1. *Let* PROP *be the a (possibly infinite) set of atomic propositions.* A possible world *is a valuation over* PROP*, and it can be identified with a set* $s \subseteq$ PROP. *We say that* p *is true in* s *(write* $s \models p$*) iff* $p \in s$*, we say that* $\overline{p}$ *is true in* s *(write* $s \models \overline{p}$*) iff* $p \notin s$.

Proposition 4.6.2. *Conditioning and lexicographic revision generate standardly universal learning methods for positive and negative data.*

Proof. In fact, any ω-type order on S gives a suitable prior plausibility assignment. Let us take $s \in S$. Since $\leq$ is ω-type it is well-founded, there are only finitely many more plausible worlds. For each such world $t \in S$ we collect a $p_n \in t$ such that $p_n \in s \dot{-} t$ ($\dot{-}$ stands for the symmetric difference of two sets). Then we construct a finite data sequence σ enumerating the all the information obtained in this manner, including p_n if $p_n \in s$ or $\overline{p}_n$ if $p_n \notin s$. Obviously σ is an initial segment of some data stream ε for s, hence $\mathrm{set}(\sigma)$ is enumerated in finite time by every data stream ε for s. After $\mathrm{set}(\sigma)$ has been observed all worlds that are more plausible than s will be deleted (in the case of conditioning) or will become less plausible than s. Hence, conditioning and lexicographic revision generate universal learning methods. □

Proposition 4.6.3. *Minimal revision is not universal for positive and negative data.*

Proof. We will give a counterexample, an epistemic state that is identifiable in the limit on positive and negative data streams, but is not identifiable in the limit by the minimal revision method. Let us first introduce the sets crucial for constructing the counterexample. Let $S_{\mathbb{N}} := \{p_i \mid i \in \mathbb{N}\}$, $S_{pos} := \{S_i = \{p_0, \ldots, p_i\} \mid i \in \mathbb{N}\}$, $S_{neg} := \{T_i = S_{\mathbb{N}} - \{p_0 \ldots p_i\} \mid i \in \mathbb{N}\}$. Now we define our epistemic state in the following way:

$$S := \{S_{\mathbb{N}}, \emptyset\} \cup S_{pos} \cup S_{neg}.$$

First let us observe that S is countable, and hence it is identifiable in the limit from positive and negative data (from the proof of Proposition 4.6.2). We will now

[6] In learning theory such streams are called 'informants', see Jain et al., 1999.

show that for any total preorder $\leq$ on S there is a set in S that is not identifiable in the limit by the minimal revision method. We will consider three basic cases: $\emptyset < S_{\mathbb{N}}$, $S_{\mathbb{N}} < \emptyset$ and $S_{\mathbb{N}} \sim \emptyset$.

1. $\emptyset < S_{\mathbb{N}}$. Let $B \subset S$ be the set of all C such that $S_{\mathbb{N}} < C$. There are two cases:

   (a) $B \neq \emptyset$. Then there is a set C such that $\emptyset < S_{\mathbb{N}} < C$ and $C \in S_{pos} \cup S_{neg}$. Then C is not identifiable in the limit by the minimal revision method.

   (b) $B = \emptyset$. Then all sets from S_{pos} are at least as plausible as $S_{\mathbb{N}}$. Then $S_{\mathbb{N}}$ is not identifiable in the limit.

2. $S_{\mathbb{N}} < \emptyset$. Again, let $B \subset S$ be the set of all C such that $\emptyset < C$. Let us again consider two cases.

   (a) $B \neq \emptyset$. Then there is a set C such that $S_{\mathbb{N}} < \emptyset < C$ and $C \in S_{pos} \cup S_{neg}$. Then C is not identifiable in the limit by the minimal revision method.

   (b) $B = \emptyset$. Then all sets from S_{neg} are at least as plausible as $\emptyset$. Then $\emptyset$ is not identifiable in the limit.

3. $\emptyset \sim S_{\mathbb{N}}$. With this assumption the elements of $S_{pos} \cup S_{neg}$ can find themselves in one of the three parts of the preorder. We can have elements that are more plausible than $\emptyset$ (we will call the set of such sets B_1), equally plausible as $\emptyset$ (set of those will be called B_2) or less plausible than $\emptyset$ (B_3). Since our epistemic set is infinite, one of B_1, B_2 and B_3 has to be infinite. Let us again consider three cases:

   (a) B_1 is infinite. Then B_1 has to contain infinitely many sets from S_{pos}, in which case $S_{\mathbb{N}}$ is not identifiable, or infinitely many sets from S_{neg}, in which case $\emptyset$ is not identifiable.

   (b) B_2 is infinite. Then the argument from the above case holds, here for B_2.

   (c) B_3 is infinite. Then B_3 has to contain infinitely many sets from S_{pos}, in which case all sets from $S_{pos} \cap B_3$ are not identifiable, or infinitely many sets from S_{neg}, in which case all sets from $S_{neg} \cap B_3$ are not identifiable.

□

4.6.1 Erroneous Information

If data streams are known to be sound and complete with respect to the actual world, the most economical strategy is to shrink the uncertainty range by deleting those possibilities that contradict the data. This strategy is based on total trust of the information source. However, in belief revision errors might be encountered

(in the form of mistakes or lies). Eliminating worlds that contradict the incoming information is then risky and irrational. It is better to change beliefs via some upgrade method, that does not have any built-in mechanism of deletion. Let us compare the performance of upgrading strategies on erroneous data.

To consider errors, we will give up the soundness of data streams, i.e., we will allow data that are false in the real world. To still keep the identification of the real world possible, the data streams are required to be 'fair': there are only finitely many errors, and every error is eventually corrected.

Definition 4.6.4. *A stream ε of positive and negative data is* fair *with respect to the world s iff:*

- *ε is complete with respect to s,*
- *there is $n \in \mathbb{N}$ such that for all $k \geq n$, all the data in $s \models \bigwedge \varepsilon_k$, and*
- *for every $i \in \mathbb{N}$ and for every $\varphi \in \varepsilon_i$ such that $s \not\models \varphi$, there exists some $k \geq i$, such that $\overline{\varphi} \in \varepsilon_k$.*

Notions defined in Subsection 4.5 (identifiability in the limit, universality, etc.) are defined analogously for fair data streams.

We will now demonstrate that lexicographic revision deals with errors in a skillful manner. Before we get to that we will introduce and discuss the notion of *propositional upgrade* (which is a special case of generalized upgrade, see Baltag & Smets, 2009b). Such an upgrade is a transformation of an epistemic-plausibility state that can be given by any finite sequence of mutually disjoint propositional sentences $x_1, \ldots, x_n$. The corresponding propositional upgrade $(x_1, \ldots, x_n)$ acts on an epistemic-plausibility state $(S, \leq_S)$ by changing the preorder $\leq_S$ as follows: all worlds that satisfy x_1 become less plausible than all satisfying x_2, all the worlds satisfying x_2 become less plausible than all x_3 worlds, etc., up to the worlds which satisfy x_n. Moreover, for any k such that $1 \leq k \leq n$, among the worlds satisfying x_k the old order $\leq_S$ is kept. In particular, our lexicographic revision is a special case of such propositional upgrade, namely in these terms lexicographic revision with φ can be identified with the propositional upgrade $(\neg\varphi, \varphi)$.

Lemma 4.6.5. *The class of propositional upgrades is closed under sequential composition.*

Proof. We need to show that the sequential composition of any two propositional upgrades gives a propositional upgrade. Let us take $X := (x_1, \ldots, x_n)$ and $Y := (y_1, \ldots, y_m)$. The sequential composition XY is equivalent to the following propositional upgrade:

$$(x_1 \wedge y_1, \ldots, x_1 \wedge y_n, x_1 \wedge y_2, \ldots, x_n \wedge y_2, \ldots, x_1 \wedge y_m, \ldots, x_n \wedge y_m).$$

To show this let us take an arbitrary epistemic-plausibility state $(S, \leq_S)$ and apply upgrades X and Y successively. First, we apply to $(S, \leq_S)$ the upgrade X.

We obtain the new preorder $\leq_S^X$, in which all worlds satisfying x_1 are less plausible than all x_2-worlds, etc., and within each such partition the old order $\leq_S$ is kept. Now, to this new epistemic-plausibility state we apply the second upgrade, Y, obtaining the new preorder $\leq_S^{XY}$, in which all y_1-worlds are less plausible than all y_2-worlds, etc. However, since the upgrade Y has been applied to the preorder $\leq_S^X$ we also know that the new preorder $\leq_S^{XY}$ has the following property: for each j, such that $1 \leq j \leq m$, within the partition given by y_j, we have that all x_1-worlds are less plausible than all x_2-worlds, etc. At the same time in each j and k, such that $1 \leq j \leq m$ and $1 \leq k \leq n$, in the partition $(y_j \wedge x_k)$ the preorder $\leq_S$ is maintained.

Putting this together, we get that $\leq_S^{XY}$ has the following structure:

$$\|(x_1 \wedge y_1)\| \geq_S^{XY} \ldots \geq_S^{XY} \|(x_n \wedge y_1)\| \geq_S^{XY}$$
$$\|(x_1 \wedge y_2)\| \geq_S^{XY} \ldots \geq_S^{XY} \|(x_n \wedge y_2)\| \geq_S^{XY} \ldots \geq_S^{XY} \|(x_n \wedge y_m)\|.$$

Moreover, within each such partition, the old preorder $\leq_S$ is kept.

The final observation is that the above setting can be obtained directly by the propositional upgrade of the following form:

$$(x_1 \wedge y_1, \ldots, x_1 \wedge y_n, x_1 \wedge y_2, \ldots, x_n \wedge y_2, \ldots, x_1 \wedge y_m, \ldots, x_n \wedge y_m).$$

□

Now we are ready to show that lexicographic revision is well-behaved on fair streams.

Proposition 4.6.6. *Lexicographic revision generates a standardly universal belief-revision-based learning method for fair streams of positive and negative data.*

Proof. First let us recall that lexicographic revision, Lex, is standardly universal on positive and negative data. For the above conjecture it is left to be shown that it retains its power on fair streams. It is sufficient to show that lexicographic revision is 'error-correcting': the effect of revising with the stream $\varphi, \sigma, \overline{\varphi}$ is exactly the same as with the stream $\sigma, \overline{\varphi}$, where σ is a sequence of propositions. The proof uses the properties of sequential composition for propositional upgrade.

Let us assume that length$(\sigma) = n$. In terms of generalized upgrade we need to demonstrate that the sequential composition $(\neg\varphi, \varphi)(\neg\sigma_1, \sigma_1) \ldots (\neg\sigma_n, \sigma_n)(\varphi, \neg\varphi)$ is equivalent to $(\neg\sigma_1, \sigma_1) \ldots (\neg\sigma_n, \sigma_n)(\varphi, \neg\varphi)$.

From Lemma 4.6.5 we know that propositional upgrade is closed under sequential composition. Hence, in the equivalence to be shown, we can replace the composition $(\neg\sigma_1, \sigma_1) \ldots (\neg\sigma_n, \sigma_n)$ by only one generalized upgrade, which we will denote by $(x_1, \ldots, x_m)$. Now, we have to show that: $(\neg\varphi, \varphi)(x_1, \ldots, x_m)(\varphi, \neg\varphi)$ is equivalent to: $(x_1, \ldots, x_m)(\varphi, \neg\varphi)$.

By the proof of Lemma 4.6.5, the composition $(x_1, \ldots, x_n)(\varphi, \neg\varphi)$ has the following form:

$$(x_1 \wedge \varphi, \ldots, x_n \wedge \varphi, x_1 \wedge \neg\varphi, \ldots, x_n \wedge \neg\varphi).$$

Accordingly, the other upgrade, $(\neg\varphi, \varphi)(x_1, \ldots, x_n)(\varphi, \neg\varphi)$, has the following form:

$$(\neg\varphi \wedge x_1 \wedge \varphi, \varphi \wedge x_1 \wedge \varphi, \ldots, \neg\varphi \wedge x_n \wedge \varphi, \varphi \wedge x_n \wedge \varphi, \neg\varphi \wedge x_1 \wedge \neg\varphi, \varphi \wedge x_1 \wedge \neg\varphi, \ldots,$$
$$\neg\varphi \wedge x_n \wedge \neg\varphi, \varphi \wedge x_n \wedge \neg\varphi).$$

Let us observe that some of the terms in the above upgrade are inconsistent. We can eliminate them since they correspond to empty subsets of the epistemic-plausibility state. We obtain:

$$(x_1 \wedge \varphi, \ldots, x_n \wedge \varphi, x_1 \wedge \neg\varphi, \ldots. x_n \wedge \neg\varphi).$$

The observation that the two propositional upgrades turn out to be the same concludes the proof. □

Proposition 4.6.7. *Conditioning and minimal revision are not universal for fair streams.*

Proof. Conditioning does not tolerate errors at all. On any ε_i such that $\varepsilon_i \not\subseteq s$ conditioning will remove s and it does not provide a way to revive it. Minimal revision, as it has been shown, is not universal on regular positive and negative data streams, which are a special case of fair streams. □

Summary In this section we have shown how the framework of iterated belief revision can be enriched by the use of negative information. First, we investigated positive and negative information that is sound and complete. In this case, both conditioning and lexicographic revision are standardly universal, i.e., there are well-founded total orders that, together with either of the two mentioned belief-revision methods, generate universal learning methods. Minimal revision again turns out to be insufficient. Secondly, we define fair data streams that use both positive and negative information. Such fair streams contain a finite number of errors and every error is eventually corrected later in the stream. The conditioning revision method again proves to be universal on fair streams, because it overrides inconsistent information. Conditioning and minimal revision lack this error-correcting property.

4.7 Conclusions and Perspectives

We have considered iterated belief-revision policies of conditioning, lexicographic and minimal revision. We have identified certain features of those methods relevant in the context of iterated revision: data-retention, conservatism, and history-independence. We defined learning methods based on those revision policies and we have shown how the aforementioned properties influence the learning process. Throughout this chapter we have been mainly interested in convergence to the actual world on the basis of infinite data streams. In the setting of positive, sound, and complete data streams we have exhibited that conditioning and lexicographic revision generate universal learning methods. Minimal revision fails to be universal, and the crucial property that makes it weaker is its strong conservatism. Moreover, we have shown that the full power of learning cannot be achieved when the underlying prior plausibility assignment is assumed to be well-founded. In the case of positive and negative information, both conditioning and lexicographic revision are universal. Minimal revision again is not. Finally, in the setting of fair streams (containing a finite number of errors that all get corrected later in the stream) lexicographic revision again turns out to be universal. Both conditioning and minimal revision lack the 'error-correcting' property.

Future work consists in multi-level investigation of the relationship between learning theory, belief revision, and dynamic epistemic logic. There surely are many links still to be found, with interesting results for everyone involved. What seems to be especially interesting is the multi-agent extension of our results. In terms of the efficiency of convergence it would enrich the multi-agent approach to information flow, an interesting subject for epistemic and doxastic logic. The interactive aspect would probably be appreciated in formal learning theory, where the single-agent perspective is clearly dominating. Another way to extend the framework is to allow revision with more complex formulae. This would perhaps link to the AGM approach, and to the philosophical investigation into the process of scientific inquiry, where possible realities have a more 'theoretical' character.

Chapter 5
Epistemic Characterizations of Identifiability

In this chapter we will further investigate the connection between formal learning theory and modal-temporal logics of belief change. We will again focus on the language-learning paradigm, in which languages are treated as sets of positive integers.[1] In the previous chapter we focused on the semantic analysis of identifiability in the limit. Now, we will devote more attention to the syntactic counterparts of our logical approach to identifiability, focusing on both finite identifiability and identifiability in the limit. We will show how the previously chosen semantics can be reflected in an appropriate syntax for knowledge, belief, and their changes over time. The corresponding notions of learning theory and dynamic epistemic logic are given in Chapter 2.

Our approach to inductive learning in the context of dynamic epistemic and epistemic temporal logic is as follows. As in the previous chapter, we take the initial class of sets to be possible worlds in an epistemic model, which mirrors Learner's initial uncertainty over the range of sets. The incoming pieces of information are taken to be events that modify the initial model. We will show that iterated update on epistemic models based on finitely identifiable classes of sets is bound to lead to the emergence of irrevocable knowledge. In a similar way identifiability in the limit leads to the emergence of stable belief. Next, we observe that the structure resulting from updating the model with a sequence of events can be viewed as an epistemic temporal forest. We explicitly focus on protocols that are assigned to worlds in set-learning scenarios. We give a temporal characterization of forests that are generated from learning situations of finite identifiability and identifiability in the limit. We observe that a special case of this protocol-based setting, in which only one stream of events is allowed in each state, can be used to model the function-learning paradigm. We show that the simple setting of iterated epistemic update cannot account for all possible learning situations. In

[1] In this chapter we are concerned with logical characterizations of learning, hence we will often refer to *languages* of certain logics. To avoid confusion for the time being we will replace the name *language learning* with *set learning* (see Section 2.1).