A. Baltag
N. Gierasimczuk
S. Smets

Truth-Tracking by Belief Revision

Abstract. We study the learning power of iterated belief-revision methods. Successful learning is understood as convergence to correct, i.e., true, beliefs. We focus on the issue of universality: whether or not a particular belief-revision method is able to learn everything that in principle is learnable. We provide a general framework for interpreting belief-revision policies as learning methods. We focus on three popular cases: conditioning, lexicographic revision, and minimal revision. Our main result is that conditioning and lexicographic revision can drive a universal learning mechanism, provided that the observations include only and all true data, and provided that a non-standard, i.e., non-well-founded prior plausibility relation is allowed. We show that a standard, i.e., well-founded belief-revision setting is in general too narrow to guarantee universality of any learning method based on belief revision. We also show that minimal revision is not universal. Finally, we consider situations in which observational errors (false observations) may occur. Given a fairness condition, which says that only finitely many errors occur, and that every error is eventually corrected, we show that lexicographic revision is still universal in this setting, while the other two methods are not.

Keywords: Belief Revision, Dynamic Epistemic Logic, Formal Learning Theory, Truth-tracking

Introduction

At the basis of the modeling of intelligent behavior lies the idea that agents integrate new information into their prior beliefs and knowledge. Intelligent agents are assumed to be endowed with some learning methods, which allow them to change their beliefs on the basis of assessing new information. But how effective is an agent's learning method in eventually finding the truth? To make this question precise and to answer it, we borrow concepts from formal learning theory and adapt them to the commonly used model of beliefs, knowledge, and belief change, namely that of possible worlds.

A set S of possible worlds, let us call it a *state space*, together with a family $\mathcal{O}$ of observable properties, represents the agent's epistemic space, her knowledge. Note that the sets S and $\mathcal{O}$ do not have to be finite, in fact throughout the paper we will assume that both S and $\mathcal{O}$ are at most countable. Intuitively, the epistemic space represents the uncertainty range

Presented by **Name of Editor**; *Received* December 1, 2005

of the agent. She can consider some possible worlds to be more plausible than others. This is captured by a total preorder on possible worlds, called a *plausibility preorder*. It captures the agent's assessments concerning which of any two worlds s, s' is more plausible to be the actual one. Such an assessment can obviously be based on many different factors, in particular on the assessed level of simplicity, or on consistency with previous observations. An epistemic space, together with a plausibility preorder is called a *plausibility space*.

The above paragraph describes a static epistemic (plausibility) space. To represent the dynamic aspects of knowledge and belief we will define methods that, triggered by an incoming information, change the epistemic (plausibility) space. The change can occur through, e.g., removal of the states incompatible with the new information, or through a revision of the plausibility relation. Many belief-revision policies proposed in the literature are formulated, or can be reconstructed, within this setting. In this paper we investigate three basic policies: conditioning, minimal revision [14], and lexicographic revision [35, 34]. The goal is to see how they can be viewed as learning methods, and to investigate their learning power, i.e., the ability to identify the real world on the basis of the incoming information.

We obtain our results by defining learning methods which are based on belief-revision policies. We show that learning from positive data via conditioning and lexicographic revision is universal, i.e., those learning methods can uniformly learn the real world, when starting in any epistemic space in which the real world is learnable (via any learning method). However, this happens only if the agent's prior plans/dispositions for belief revision are suitably chosen; and not every such prior is suitable. Furthermore, the most conservative belief-revision method, minimal revision, is not universal.

Our approach brings together methods of formal learning theory [FLT, see, e.g., 33] and Dynamic Epistemic Logic [DEL, see 6, 5, 18, 12]. The interest in bringing together learning theory and belief-revision theory has appeared before within at least two lines of research. Firstly, in [30, 26, 27, 28, 29] some classical belief-revision policies were treated as learning strategies. Secondly, in [31, 32] the connection has been rooted in the classical AGM framework [1]. Finally, in [19, 20, 21, 23, 22] the set learning paradigm (also called *language learning*) has been connected with epistemic and doxastic logics of belief revision [3, 11, 17, 8, 7, 9]. The present paper is a continuation of the latter line of research, and is in fact a thorough presentation of results announced in a previously published extended abstract [4].

We are chiefly concerned with the possible-world based counterpart of one of the central notions in formal learning theory, namely *identifiability in* *the limit* [24]. Hence, we focus on *stabilizing* to a correct *belief.*[1] We hence investigate the *reliability* of mind-changing strategies, i.e., the possibility of converging to an accurate hypothesis after a finite number of mind-changes.

1. Notation and basic definitions

Let S be a possibly infinite set of possible worlds and let $\mathcal{O} \subseteq \mathcal{P}(S)$ be a possibly infinite (but at most countable) set of observable properties. An observable property is henceforth identified with the set of those possible worlds which make the property true. These properties can be observed by an agent and hence can be viewed as *data* or *evidence* for learning: an agent can learn whether or not they hold. This does not mean that they are necessarily all observable at the same time. Indeed, it is natural to assume that only finitely many of them could be observed at a given moment. To simplify we here assume that at each step of the learning process only one observable property is accessed by the learner.

The agent is represented by her epistemic space, i.e., a range of possible worlds that satisfy relevant observable properties.

Definition 1. *Let S be a set of possible worlds and $\mathcal{O} \subseteq \mathcal{P}(S)$. The pair $\mathbb{S} = (S, \mathcal{O})$ is then called an* epistemic space*.*

The epistemic space represents an agent who does not favor any possibility over others. We extend epistemic spaces to capture such a case by introducing a total preorder called a *plausibility relation*.

Definition 2. *Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space, and $\preceq \subseteq S \times S$ be a total preorder.*[2] *The structure $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ is called a* plausibility space*.*

Since we allow for the epistemic space to be infinite, the question of well-foundedness of the plausibility space becomes very relevant. We do not restrict our considerations to well-founded spaces, but we will take them into account as a special class of plausibility spaces. Because of their popularity in the literature we call them *standard plausibility spaces.*[3]

Definition 3. *A* standard plausibility space *$\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ is one whose plausibility relation $\preceq$ is well-founded (i.e., there is no infinite descending chain $s_0 \succ s_1 \succ \ldots \succ s_n \succ \ldots$, where $\prec$ is the strict plausibility relation, given by: $s \preceq t$ and $t \not\preceq s$.*

[1] The emergence of the stronger epistemic state of irrevocable knowledge can be linked to a more restrictive kind of identifiability, finite identifiability [see 15, 16, 23].

[2] In other words, the binary relation $\preceq$ is total, reflexive, and transitive in S.

[3] Note that we interpret $s \prec t$ as 's is *more* plausible than t'.

1.1. Knowledge and belief in epistemic spaces

Let us briefly discuss how this setting relates to that of epistemic and doxastic logic. The logical interpretation depends on the notion of observability that we are willing to employ. In general, the set of observable properties can be any set (of sets of possible worlds) closed under certain operations, e.g., under negation (if 'negative data' are observed), or under finite intersection (if 'conjunctions' are observed). Under the usual possible-world interpretation, $\mathcal{O}$ can be viewed as the set of (atomic) propositions, or, if the stress is put on the closure under certain operations (e.g., negation or conjunction), as a set encoding the valuation for a relevant logical language.

For the sake of shaping the right intuitions about the interpretation of logical notions of knowledge and belief in epistemic and plausibility spaces, we define the language of doxastic epistemic logic in the single agent case. Our purpose here is to clarify the (standard) meaning of the operators of knowledge and belief that we use.

Definition 4 (Syntax). *The language of doxastic logic for observable properties is given by the following syntax:*

$$\varphi := p \mid \neg\varphi \mid \varphi \Rightarrow \varphi \mid K\varphi \mid B\varphi,$$

for any $p \in \mathcal{O}$.

Definition 5 (Semantics). *Let us take a plausibility space* $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, $s \in S$, *and* $p, q \in \mathcal{O}$.

$$\begin{array}{lll}
\mathbb{B}_{\mathbb{S}}, s \models p & \text{iff} & s \in p \\
\mathbb{B}_{\mathbb{S}}, s \models \neg p & \text{iff} & \textit{it is not the case that } s \in p \\
\mathbb{B}_{\mathbb{S}}, s \models p \Rightarrow q & \text{iff} & \textit{if } s \in p, \textit{ then } s \in q \\
\mathbb{B}_{\mathbb{S}}, s \models Kp & \text{iff} & \textit{for all } t \in S, \textit{we have } t \in p \\
\mathbb{B}_{\mathbb{S}}, s \models B\varphi & \text{iff} & \exists w \preceq s\ \forall u \preceq w\ \mathbb{B}_{\mathbb{S}}, u \models p
\end{array}$$

In other words, we say that the agent knows φ iff φ is true in all possible worlds of the epistemic space $\mathbb{S}$. Since knowledge in our single-agent models is a global modality, the truth of knowledge of a formula in a particular world is equivalent to the validity of the formula in the whole epistemic space. We then write $\mathbb{B}_{\mathbb{S}} \models K\varphi$. The same holds for the belief operator since the $\preceq$ relation is total; in the single agent case the valuation of a belief formula cannot vary from world to world, we hence write $\mathbb{B}_{\mathbb{S}} \models B\varphi$.

As we mentioned before, most of the epistemic doxastic logic and belief-revision literature deals with standard, i.e., well-founded plausibility structures. The well-foundedness assumption gives at least two advantages.

Firstly, it allows to canonically assign ordinal numbers to states [so-called *Spohn ordinals* or 'degrees of implausibility', see 35]. Secondly, it leads to a simplified definition of belief, which can then be understood as 'truth in all the most plausible worlds'. Indeed, in any plausibility model based on a standard plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ and any $s \in S$, we have:

$$\mathbb{B}_{\mathbb{S}}, s \models Bp \quad \text{iff} \quad \min_{\preceq} S \subseteq p,$$

where for any set $X \subseteq S$, $\min_{\preceq} X$ is the set of all its most plausible worlds[4], defined as $\{t \in X \mid t \preceq s \text{ for all } s \in X\}$.

We do not assume well-foundedness in this paper, simply because we cannot afford it. As we show, the class of standard plausibility structures is too narrow for obtaining universal learning via belief revision.

1.2. Observable properties

The agent is now identified with her plausibility space, and as such it is the static component of our modeling of the dynamic situation of learning and belief revision. She is allowed to observe properties and on the basis of those observations revise her beliefs. Before we get to the revision process itself, we will first devote some space to the nature of observations. We will also introduce some useful notation.

First, let us note that we are interested in unbounded sequences of events. Therefore, we have to consider infinite streams of information consisting of observable properties.

Definition 6. *Let $\mathcal{S} = (S, \mathcal{O})$ be an epistemic space and let $\mathbb{N} = \mathbb{N}^+ \cup \{0\}$ denote the set of all natural numbers.*
A data stream *is an infinite sequence $\vec{O} = (O_0, O_1, \ldots)$ of data $O_i \in \mathcal{O}, i \in \mathbb{N}$.*
A data sequence *is a finite sequence $\sigma = (O_0, \ldots, O_n)$. For any stage $n \in \mathbb{N}$ of a data stream $\vec{O} = (O_0, O_1, \ldots)$, the initial segment $(O_0, \ldots, O_{n-1})$ forms a data sequence.*

The intuition behind the streams of data is that at stage i, the agent observes the information in O_i. A data stream captures a possible future history of observations in its entirety, while a data sequence captures only a finite part of such a history.

Definition 7. *Let $\vec{O} = (O_0, O_1, \ldots)$ be any data stream, and let $\sigma = (\sigma_0, \ldots, \sigma_n)$ be any data sequence.*

[4] It is easy to see that, if $\preceq$ is well-founded, then $\min_{\preceq} X \neq \emptyset$ whenever $X \neq \emptyset$.

$\vec{O}_n$ *stands for the* n*-th observation in* $\vec{O}$.

$\vec{O}[n]$ *stands for the the initial segment of* $\vec{O}$ *of length* n, $(O_0, \ldots, O_{n-1})$.

$\text{set}(\vec{O}) := \{O \mid O$ *is an element of* $\vec{O}\}$ *stands for the set of all data in* $\vec{O}$*; we similarly define set*(σ)*, where* σ *is a finite data sequence.*

$\sigma * \vec{O} := (\sigma_0, \ldots, \sigma_n, O_0, O_2, \ldots)$ *is the concatenation of the finite sequence* σ *with the infinite stream* $\vec{O}$.

We pose restrictions on the observations. We require that only the observable properties that are true can be observed by the agent. Such data streams, those that include only data that is true in the actual world, are be called *sound data streams*.

Definition 8. *A data stream* $\vec{O}$ *is* sound with respect to state s *iff every element listed in* $\vec{O}$ *is true in* s*, i.e.,* $s \in \vec{O}_n$ *for all* $n \in \mathbb{N}$.

To create an ideal environment for learning we also assume that data streams are *complete*, i.e., they enumerate all observable properties that are true in the actual world.

Definition 9. *A data stream* $\vec{O}$ *is* complete with respect to state s *iff every observable true in* s *is listed in* $\vec{O}$*, i.e., for any* $O \in \mathcal{O}$*, if* $s \in O$ *then* $O = \vec{O}_n$ *for some* $n \in \mathbb{N}$.

Let us put $\mathcal{O}_s$ to stand for the set of all observable properties that are true in s, i.e., $\mathcal{O}_s = \{O \in \mathcal{O} \mid s \in O\}$. A data stream $\vec{O}$ is then *sound and complete with respect to state* s if and only if $\mathcal{O}_s = \text{set}(\vec{O})$.

In most of this paper we assume the data streams to be sound and complete with respect to the actual world, i.e., all observed data is true, and all true data will eventually be observed.[5] This assumption obviously means that, in the limit, the agent gets the most favorable conditions for learning the whole truth about the identity of the actual world.

2. Learning and belief-revision methods

We represent the learner as a function that, given the initial set of possibilities S and given any sequence of observed data, produces a *conjecture*: some subset of S (to which the actual world is conjectured to belong). In a natural way, we interpret the conjecture doxastically, as the agent's *current belief* about the real world (after observing the given sequence of data).

[5] In classical computational learning theory such sound and complete data streams are called 'texts' [25] or 'environments' [32].

Definition 10. *Let* $\mathbb{S} = (S, \mathcal{O})$ *be an epistemic space and let* $\sigma_0, \ldots, \sigma_n \in \mathcal{O}$*. A* learning method *is a function* L *that on the input of* $\mathbb{S}$ *and data sequence* $(\sigma_0, \ldots, \sigma_n)$ *outputs some set of worlds* $L(\mathbb{S}, (\sigma_0, \ldots, \sigma_n)) \subseteq S$*, called a* conjecture*.*

Such learning can have various properties, for instance the learner can be forgetful, or conservative in revising her conjectures. Below we list several properties of this type.

Definition 11. *Let* $\mathbb{S} = (S, \mathcal{O})$ *be an epistemic space, and let* $\sigma = (\sigma_0, \ldots, \sigma_n)$ *be any non-empty data sequence, for any* $n \in \mathbb{N}$*. A learning method* L *is called:*

weakly data-retentive *iff* $L(\mathbb{S}, \sigma) \neq \emptyset$ *implies* $L(\mathbb{S}, \sigma) \subseteq \sigma_n$*;*

strongly data-retentive *iff* $L(\mathbb{S}, \sigma) \neq \emptyset$ *implies* $L(\mathbb{S}, \sigma) \subseteq \bigcap_{i \in \{0,\ldots,n\}} \sigma_i$*;*

conservative *iff for every* $p \in \mathcal{O}$ *such that* $\emptyset \neq L(\mathbb{S}, \sigma) \subseteq p$*, we have* $L(\mathbb{S}, \sigma) = L(\mathbb{S}, \sigma * p)$*;*

data-driven *if it is both conservative and weakly data-retentive;*

memory-free *if for every two data sequences* σ, σ'*, and every* $p \in \mathcal{O}$*,* $L(\mathbb{S}, \sigma) = L(\mathbb{S}, \sigma')$ *implies* $L(\mathbb{S}, \sigma * p) = L(\mathbb{S}, \sigma' * p)$*.*

Let us now briefly compare the above properties with the AGM postulates [1]. Weak data retention means that the current conjecture always fits the most recently observed data. If we interpret conjectures as beliefs, this intuitively corresponds to the AGM *Success Postulate.* We are also able to formulate a stronger condition of a similar type: strong data retention says that the current conjecture always accounts for all data that have been encountered till now. The two limit cases, i.e., accounting for the last datum and accounting for all data received so far, open a way to analyzing a whole spectrum of retention levels. Conservativity requires that the agent keeps the same beliefs whenever the new piece of data is already entailed by her old beliefs. A learning method is memory-free if, at each stage, the new belief set depends only on the previous belief set and the new piece of data. The latter property was the original intention behind the AGM notation $T * \varphi$, for revision of a theory T with a new piece of data φ [see 1]; but, in fact, this assumption poses severe problems for *iterated* belief revision. Indeed, as we will see later, most standard belief-revision methods implementing the AGM postulates are not memory-free: the new belief depends in addition on some hidden parameter, namely the old plausibility relation.

We now turn to defining belief-revision methods. They are transformations of plausibility spaces triggered by the incoming data sequences.

Definition 12. *A* belief-revision method *is a function R that, for any plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ and any data sequence $\sigma = (\sigma_0, \ldots, \sigma_n)$, for any $n \in \mathbb{N}$, outputs a new plausibility space $R(\mathbb{B}_{\mathbb{S}}, \sigma) := (S^\sigma, \mathcal{O}, \preceq^\sigma)$.*[6]

To be more constructive about the belief-revision treatment of plausibility spaces, we define a special class of belief-revision methods, the *iterated* belief-revision methods.

Definition 13. *A* one-step revision method *is a function R_1 that, for any plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ and any $p \in \mathcal{O}$ outputs a new plausibility space $R_1((S, \mathcal{O}, \preceq), p) := (S^p, \mathcal{O}, \preceq^p)$.*

An iterated belief-revision method *is a belief-revision method R_1^∞ obtained by iterating a one-step revision method R_1, i.e., by recursively defining for any data sequence σ:*

$$R_1^\infty((S, \mathcal{O}, \preceq), \lambda) = (S, \mathcal{O}, \preceq), \text{ for the empty data sequence } \lambda,$$

$$R_1^\infty((S, \mathcal{O}, \preceq), \sigma * p) = R_1(R_1^\infty((S, \mathcal{O}, \preceq), \sigma), p).$$

Having defined the learning and the belief-revision methods, we are now ready to put the two things together and define learning based on belief-revision methods. It is enough for a learning method to be given some *prior* plausibility order on the initial epistemic space $\mathbb{S}$ (in this way obtaining a plausibility space $\mathbb{B}_{\mathbb{S}}$), and then to simulate how a chosen belief-revision method chooses a plausibility space $R(\mathbb{B}_{\mathbb{S}}, \sigma)$ for any data sequence σ. Hence, we can use this plausibility space to define a belief set (and hence a 'conjecture'). Obviously, we can only do that if there exist most plausible states with respect to the obtained plausibility preorder $\preceq^\sigma$.

Definition 14. *Let $\mathbb{S} = (S, \mathcal{O})$ be any epistemic space. A* prior plausibility assignment *$f_\preceq$ is a map $\mathbb{S} \mapsto \preceq_{\mathbb{S}}$ that assigns to $\mathbb{S}$ some plausibility relation $\preceq_{\mathbb{S}}$ on S (i.e., a total pre-order on S), thus converting it into a plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq_{\mathbb{S}})$.*

Definition 15. *Every belief-revision method R, together with a prior plausibility assignment $f_\preceq$, generates in a canonical way a learning method $L_R^\preceq$, called a* belief-revision-based learning method*, and given by:*

$$L_R^\preceq(\mathbb{S}, \sigma) := \min_{\preceq_{\mathbb{S}}} R((S, \mathcal{O}, \preceq_{\mathbb{S}}), \sigma).$$

In the particular case of iterated belief-revision methods R_1^∞, for simplicity we denote by $L_{R_1}^\preceq := L_{R_1^\infty}^\preceq$ the learning method generated by R_1^∞.

[6] As we will see below, the exact structure of S^σ and $\preceq^\sigma$ will vary depending on the chosen belief revision method.

The previously defined properties of learning functions (Definition 11) can be naturally applied to belief-revision methods.

Definition 16. *A belief-revision method R is called* weakly data-retentive (strongly data-retentive, conservative*, or* data-driven) *iff for any prior plausibility assignment $f_{\preceq}$, the induced learning method $L_R^{\preceq}$ is weakly data-retentive (strongly data-retentive, conservative, or data-driven).*

The properties that the belief-revision methods inherit from their corresponding learning methods can be characterized in doxastic logic (the latter was discussed in Section 1.1).

Proposition 1. *Let $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ be a plausibility space, R be a belief-revision method, $\sigma = (\sigma_0, \ldots, \sigma_n)$ be any non-empty data sequence of any length $n \in \mathbb{N}$, for $0 \leq i \leq n$, $\sigma_i \in \mathcal{O}$.*

(1) R is weakly data-retentive iff $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models B\sigma_n$;

(2) R is strongly data-retentive iff $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models B\sigma_i$, for any $i \in \{0, \ldots, n\}$;

(3) if R is conservative then, for every $p, q \in \mathcal{O}$ such that $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bp$ we have $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bq$ iff $R(\mathbb{B}_{\mathbb{S}}, \sigma * p) \models Bq$.

Proof. The left-to-right implications of (1) and (2) are trivial, given the semantics of the belief operator. For the right-to-left implication of (1), let us take a belief-revision method R and some epistemic space together with a prior plausibility assignment $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq_{\mathbb{S}})$. Assume that R is such that, for every data sequence $\sigma = (\sigma_0, \ldots, \sigma_n)$, we have $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models B\sigma_n$. To prove that R is weakly data-retentive, we need to show that if $L_R(\mathbb{S}, \sigma) \neq \emptyset$ then $L_R(\mathbb{S}, \sigma) \subseteq \sigma_n$. Let us take $R(\mathbb{B}_{\mathbb{S}}, \sigma) = (S^{\sigma}, \mathcal{O}, \preceq_{\mathbb{S}}^{\sigma})$ then assume that $L_R(S, \sigma) \neq \emptyset$, i.e., there are $\preceq_{\mathbb{S}}^{\sigma}$-minimal elements in S^{σ}. Then in every world minimal with respect to $\preceq_{\mathbb{S}}$ σ_n holds: $\min_{\preceq_{\mathbb{S}}^{\sigma}} R(\mathbb{B}_{\mathbb{S}}, \sigma) \subseteq \sigma_n$. Since $\min_{\preceq_{\mathbb{S}}^{\sigma}} R(\mathbb{B}_{\mathbb{S}}, \sigma) = L_R(\mathbb{S}, \sigma)$, we have that $L_R^{\preceq_{\mathbb{S}}}(\mathbb{S}, \sigma) \subseteq \sigma_n$. The proof of the right-to-left implication in the second assertion is analogous.

For (3), let us take a belief-revision method R and assume that it is conservative, i.e., that for any $\mathbb{S} = (S, \mathcal{O})$ and any total preorder $\preceq \subseteq S \times S$ the canonical learning method $L_R^{\preceq}$ is conservative. This means that for any $p \in \mathcal{O}$ such that $L_R^{\preceq}(\mathbb{S}, \sigma) \subseteq p$, we have $L_R^{\preceq}(\mathbb{S}, \sigma) = L_R^{\preceq}(\mathbb{S}, \sigma * p)$. We need to show that then for every $p, q \in \mathcal{O}$ such that $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bp$ we have $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bq$ iff $R(\mathbb{B}_{\mathbb{S}}, \sigma * p) \models Bq$.

Let us then take $p, q \in \mathcal{O}$, such that $\emptyset \neq L_R^{\preceq}(\mathbb{S}, \sigma) \subseteq p$, $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bp$, and assume that $L_R^{\preceq}(\mathbb{S}, \sigma) = L_R^{\preceq}(\mathbb{S}, \sigma * p)$. We need to show that

$R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bq$ iff $R(\mathbb{B}_{\mathbb{S}}, \sigma * p) \models Bq$. For the left to right direction assume that $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bq$. Since $L_R^{\preceq}(\mathbb{S}, \sigma) \neq \emptyset$, by the definition of B, we also know that $L_R^{\preceq}(\mathbb{S}, \sigma) \subseteq q$. Then, since $L_R^{\preceq}(\mathbb{S}, \sigma) = L_R^{\preceq}(\mathbb{S}, \sigma * p)$ and by the definition of belief-revision based learning method, we have that $L_R^{\preceq}(\mathbb{S}, \sigma * p) = \min_{\preceq} R(\mathbb{S}, \sigma * p) \subseteq q$, which by the semantics of B means that $R(\mathbb{B}_{\mathbb{S}}, \sigma * p) \models Bq$. For the reverse the argument is analogous. ■

We also define additional, specific to belief-revision properties: strong conservativity and history independence.

DEFINITION 17. *Let us take $\mathbb{B}_{\mathbb{S}} = (S, \preceq, \mathcal{O})$. A belief-revision method R is:*

a) strongly conservative *iff for every $p \in \mathcal{O}$ such that $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bp$, we have $R(\mathbb{B}_{\mathbb{S}}, \sigma) = R(\mathbb{B}_{\mathbb{S}}, \sigma * p)$, i.e., R does not change the plausibility space at all, if the new piece of data has already been believed.*

b) history-independent *iff for every $p \in \mathcal{O}$ and all data sequences σ, π we have that if $R(\mathbb{B}_{\mathbb{S}}, \sigma) = R(\mathbb{B}_{\mathbb{S}}, \pi)$ then $R(\mathbb{B}_{\mathbb{S}}, \sigma * p) = R(\mathbb{B}_{\mathbb{S}}, \pi * p)$, i.e., R's output at any stage depends only on the previous output and the most recently observed data.*

In the light of consistent data, strongly conservative belief-revision methods not only keep the old conjecture the same, but, when receiving truthful information, they do not change anything within the plausibility space. As we will see below, the classical belief-revision methods are not necessarily strongly conservative. However, every iterated belief-revision method R_1^{∞} must be history-independent. History-independent methods do not require the agent to keep in memory all the past events: only the last plausibility space and the new piece of data are enough to determine the next plausibility space. However, as we will show in the next section, the corresponding learning is not necessarily memory-free.

3. Some iterated belief-revision methods

Below we consider three basic iterated belief-revision methods that received considerable attention within the belief-revision and the DEL literature. All three satisfy all the AGM postulates.[7] We only need to define the one-step revision method that canonically generates the iterated versions of each of them.

[7] Note that in general, when updating with epistemic information, the AGM success postulate is not always satisfied. However, in our setting this problem is avoided because all the new incoming data for an agent consists only of atomic sentences representing basic ontic facts about the world.

Conditioning First we focus on the revision by *conditioning* [35, 34], also called *update* in DEL [11, 9]. This method operates by deleting those worlds that do not satisfy the newly observed data.

DEFINITION 18. Conditioning *is a one-step belief-revision method* `Cond` *that takes a plausibility space* $\mathbb{B}_{\mathbb{S}}$ *and a proposition* $p \in \mathcal{O}$*, and outputs a new plausibility space in the following way:* $\texttt{Cond}(\mathbb{B}_{\mathbb{S}}, p) = \mathbb{B}_{\mathbb{S}}{}' = (S', \mathcal{O}, \preceq')$*, where* $S' = S \cap p$*, and* $\preceq' = \preceq \cap (S' \times S')$*.*

It is easy to see that conditioning is weakly data-retentive. Moreover, one might say that this method takes the incoming information 'very seriously', it deletes all worlds inconsistent with it. The deletion cannot be reversed, hence, in a way, conditioning is the ultimate belief-revision method to memorize past observations.

PROPOSITION 2. *Conditioning is strongly data-retentive.*

PROOF. Let us take $\sigma = (\sigma_0, \ldots, \sigma_n)$ and $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$. By Proposition 1, it is enough to show that for every $0 \preceq i \preceq n$, $\texttt{Cond}(\mathbb{B}_{\mathbb{S}}, \sigma) \models B\sigma_i$. Let us refer to the new plausibility space $\texttt{Cond}(\mathbb{B}_{\mathbb{S}}, \sigma)$ as $\mathbb{B}_{\mathbb{S}}{}^{\sigma} = (S^{\sigma}, \mathcal{O}, \preceq^{\sigma})$. Each time the new information σ_i comes in, all worlds that do not satisfy it are eliminated, therefore $S^{\sigma} = \bigcap set(\sigma)$. We get that $\texttt{Cond}(\mathbb{B}_{\mathbb{S}}, \sigma) \models B(\bigwedge set(\sigma))$, i.e., in the resulting belief space every proposition that ever occurred in σ is believed: $\texttt{Cond}(\mathbb{B}_{\mathbb{S}}, \sigma) \models B\sigma_i$, for $i \in \{0, \ldots, n\}$. ■

While conditioning is obviously conservative, it does not satisfy the strong version of this condition.

PROPOSITION 3. *Conditioning is not strongly conservative.*

PROOF. We need to demonstrate that for a plausibility space $\mathbb{B}_{\mathbb{S}}$ and an observable $p \in \mathcal{O}$ we have that $R(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bp$ but $R(\mathbb{B}_{\mathbb{S}}, \sigma) \neq R(\mathbb{B}_{\mathbb{S}}, \sigma * p)$. Let us consider $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, with $\mathcal{O} = \{p, q\}$, and $S = \{s, t\}$ such that $p = \{s, t\}$, $q = \{s\}$, and the plausibility $\preceq$ gives $s \prec t$ (see Figure 1). Then $\mathbb{B}_{\mathbb{S}} \models Bq$. However, after receiving q, the revision method `Cond` will eliminate world t and therefore $R(\mathbb{B}_{\mathbb{S}}, \sigma) \neq R(\mathbb{B}_{\mathbb{S}}, \sigma * q)$. ■

Lexicographic Revision Lexicographic revision [34, 35], also known as *radical upgrade* in Dynamic Epistemic Logic [11, 9] does not delete any worlds. Instead, it 'promotes' all the worlds satisfying the new piece of data, making them more plausible than all the worlds that do not satisfy it; while within the two zones, the old order is kept the same.

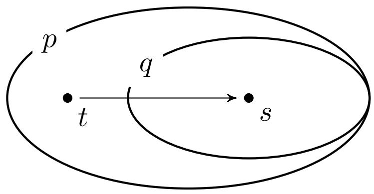

Figure 1. Plausibility space from the proof of Proposition 3. The arrow points to the more plausible world, in this case $s \prec t$.

DEFINITION 19. Lexicographic revision *is a one-step belief-revision method* `Lex` *that takes a plausibility space* $\mathbb{B}_{\mathbb{S}}$ *and a proposition* $p \in \mathcal{O}$*, and outputs a new plausibility space in the following way:* $\texttt{Lex}(\mathbb{B}_{\mathbb{S}}, p) = \mathbb{B}_{\mathbb{S}}{}' = (S, \mathcal{O}, \preceq')$ *where for all* $t, w \in S$*,* $t \preceq' w$ *iff* $(t \preceq_p w$ *or* $t \preceq_{\bar{p}} w$ *or* $(t \in p \wedge w \notin p))$*, where:* $\preceq_p = \preceq \cap (p \times p)$*,* $\preceq_{\bar{p}} = \preceq \cap (\bar{p} \times \bar{p})$*, and* $\bar{p}$ *stands for the complement of* p *in* S*.*

It is very easy to see that lexicographic revision is weakly data-retentive and conservative. However, this revision method does not satisfy the strong versions of these properties:

PROPOSITION 4. *Lexicographic revision is not strongly data-retentive on arbitrary data streams.*

PROOF. By Proposition 1, it is enough to show that there is a $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, $\sigma = (\sigma_0, \ldots, \sigma_n)$, and $i \in \{0, \ldots, n\}$ such that $\texttt{Lex}(\mathbb{B}_{\mathbb{S}}, \sigma) \not\models B\sigma_i$. Let us take $S = \{s, t\}$, $\mathcal{O} = \{p, q\}$ such that $p = \{s\}$ and $q = \{t\}$. Let us also assume any initial plausibility ordering on S, e.g., $s \preceq t$, and take $\sigma = (p, q)$. First, $\sigma_0 = p$ comes in, and now p is believed. After receiving $\sigma_1 = q$ the most plausible state becomes t, so p is no longer believed, i.e., $\texttt{Lex}(\mathbb{B}_{\mathbb{S}}, \sigma) \not\models B\sigma_0$. ■

PROPOSITION 5. *Lexicographic revision is strongly data-retentive on sound data streams.*

PROOF. Let us take $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$. Let us also fix $s \in S$ and assume that $\sigma = (\sigma_0, \ldots, \sigma_n)$, is sound with respect to s, i.e., $\mathbb{B}_{\mathbb{S}}, s \models \bigwedge \text{set}(\sigma)$. Let us take $\texttt{Lex}(\mathbb{B}_{\mathbb{S}}, \sigma) = (S^{\sigma}, \mathcal{O}, \preceq^{\sigma})$. After reading σ, for all the worlds t that are most plausible with respect to $\preceq^{\sigma}$ it is the case that $\mathbb{B}_{\mathbb{S}}, t \models \bigwedge \text{set}(\sigma)$, and hence that $\mathbb{B}_{\mathbb{S}} \models B(\bigwedge \text{set}(\sigma))$. It is so because by assumption there is at least one such world, s. ■

Proposition 6. *Lexicographic revision is not strongly conservative.*

Proof. Let us consider the following example of a plausibility space (see Figure 2). Assume that $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, where $S = \{s, t, u\}$, $\mathcal{O} = \{p, q\}$, and $p = \{s, u\}$, $q = \{s, t\}$. Let us also assume that the plausibility $\preceq$ gives the following order: $s \prec t \prec u$, and that $\sigma = (p)$. Clearly, $\mathbb{B}_{\mathbb{S}} \models Bp$. However, after receiving $\sigma_0 = p$, the revision method will still put world u to be more plausible than t, and therefore $\mathbb{B}_{\mathbb{S}} \neq \texttt{Lex}(\mathbb{B}_{\mathbb{S}}, p)$. ■

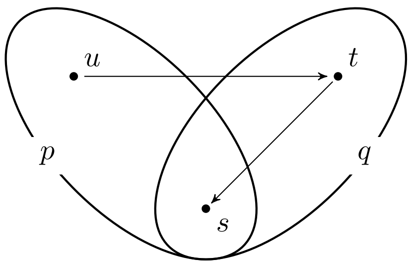

Figure 2. Plausibility space from the proof of Proposition 6

Proposition 7. *A learning method $L_R^{\preceq}$ generated from a history-independent belief-revision method R does not have to be memory-free.*

Proof. We prove this proposition by showing an example, a belief-revision method R that is history-independent but the learning method that it induces is not memory-free (see Figure 3). Let R be Lex. Lex is clearly history-independent, because it is an iterated one-step revision method. To see that $L_{\texttt{Lex}}$ is not memory-free consider the following two plausibility spaces $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ and $\mathbb{B}_{\mathbb{S}}' = (S, \mathcal{O}, \preceq')$ with $\mathcal{O} = \{p, q\}$ and $S = \{s, t, u\}$, such that $p = \{s, u\}$, $q = \{t, u\}$. Assume that for some σ and σ':

(1) $\texttt{Lex}(\mathbb{B}_{\mathbb{S}}, \sigma)$ gives the plausibility order: $s \prec_\sigma u \prec_\sigma t$;

(2) $\texttt{Lex}(\mathbb{B}_{\mathbb{S}}', \sigma')$ gives the plausibility order: $s \prec_\sigma' t \prec_\sigma' u$.

We have that $L^{\preceq}_{\texttt{Lex}}(\mathbb{S}, \sigma) = L^{\preceq'}_{\texttt{Lex}}(\mathbb{S}, \sigma') = \{s\}$. Assume now that the next observation is q. Then clearly $L^{\preceq}_{\texttt{Lex}}(\mathbb{S}, \sigma * q) = \{u\}$, while $L^{\preceq'}_{\texttt{Lex}}(\mathbb{S}, \sigma' * q) = \{t\}$. Therefore, for the belief-revision method Lex there is a $p \in \mathcal{O}$ such that: $L^{\preceq}_{\texttt{Lex}}(\mathbb{S}, \sigma) = L^{\preceq'}_{\texttt{Lex}}(\mathbb{S}, \sigma')$ but $L^{\preceq}_{\texttt{Lex}}(\mathbb{S}, \sigma * p) \neq L^{\preceq'}_{\texttt{Lex}}(\mathbb{S}, \sigma' * p)$. ■

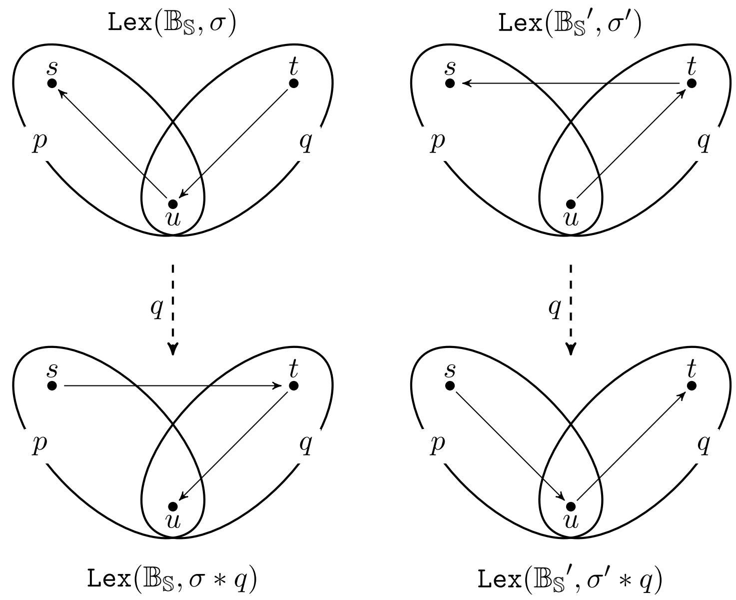

Figure 3. The transformations of the plausibility space from the proof of Proposition 7

Minimal Revision The minimal revision method [14, 34], known as *conservative upgrade* in DEL [11, 9], is 'conservative' in the sense that it keeps as much as possible of the old structure. More precisely, the most plausible states satisfying the new piece of data become the most plausible overall; while in the rest of the space, the old order is kept the same.

DEFINITION 20. Minimal revision *is a one-step belief-revision method* `Mini` *that takes a plausibility space* $\mathbb{B}_{\mathbb{S}}$ *and a proposition* $p \in \mathcal{O}$*, and outputs a new plausibility space in the following way:* $\texttt{Mini}(\mathbb{B}_{\mathbb{S}}, p) = \mathbb{B}_{\mathbb{S}}{}' = (S, \mathcal{O}, \preceq')$ *where for all* $t, w \in S$*, if* $t \in min_{\preceq}p$ *and* $w \notin min_{\preceq}p$*, then* $t \preceq' w$*, otherwise* $t \preceq' w$ *iff* $t \preceq w$*.*

Minimal revision is obviously weakly data-retentive—it leads to a belief that accounts for the last datum. However, it does not retain more than that.

PROPOSITION 8. *Minimal revision is not strongly data-retentive.*

PROOF. Let consider the plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, such that $S = \{s, t, u\}$, $\mathcal{O} = \{p, q\}$, $p = \{s, u\}$, $q = \{t, u\}$, the sequence $\sigma = (p, q)$ (which is sound with respect to world u), and assume that the initial ordering on S is $t \prec s \prec u$. After receiving $\sigma_0 = p$ the plausibility ordering $\prec^{\sigma_0}$ becomes $s \prec^{\sigma_0} t \prec^{\sigma_0} u$. Then $\sigma_1 = q$ comes in and now our method gives the ordering $t \prec^{(\sigma_0,\sigma_1)} s \prec^{(\sigma_0,\sigma_1)} u$. So p is no longer believed after the second piece of data was given, hence $\texttt{Mini}(\mathbb{B}_{\mathbb{S}}, \sigma) \not\models B\sigma_0$. ■

Moreover, the minimal revision is conservative—as long as the incoming information is consistent with the currently held belief, the belief does not change, since the minimal worlds do not change. In this case we can say even more: nothing about the plausibility space changes.

Proposition 9. *Minimal revision is strongly conservative.*

Proof. Let us take $\sigma = (\sigma_0, \ldots, \sigma_n)$, $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, and let $\texttt{Mini}(\mathbb{B}_{\mathbb{S}}, \sigma) = (S, \mathcal{O}, \preceq^{\sigma})$. Let us assume, towards contradiction, that there is a $p \in \mathcal{O}$ such that $\texttt{Mini}(\mathbb{B}_{\mathbb{S}}, \sigma) \models Bp$ but $\texttt{Mini}(\mathbb{B}_{\mathbb{S}}, \sigma) \neq \texttt{Mini}(\mathbb{B}_{\mathbb{S}}, \sigma * p)$. This means that after receiving p the plausibility order of $\mathbb{B}_{\mathbb{S}}$ has been rearranged. By the definition of Mini, this could happen only in the case when among the most plausible worlds in $\texttt{Mini}(\mathbb{B}_{\mathbb{S}}, \sigma)$ there was a world t such that $t \notin p$. But then also $\texttt{Mini}(\mathbb{B}_{\mathbb{S}}, \sigma) \not\models Bp$. Contradiction. ■

The properties that we introduced in this section capture some interesting differences between belief-revision methods. While conditioning and lexicographic revision are quite similar, differing only with respect to their strong retention capacity, minimal revision is different in two respects. It is not strongly data retentive, even on sound data streams. As the only one it is also strongly conservative, necessarily preserving the old plausibility spaces upon receiving information that is in the first place already believed. In the next section we will see that this combination of properties negatively affects learning.

4. Convergence to truth

Formal learning theory is concerned with *reliable* learning methods, i.e., those that can be relied upon (when observing a sound and complete data stream) to find in finite time the real world, no matter what the real world is, as long as it is among the possibilities allowed by the initial epistemic space $\mathbb{S}$.[8] By 'reliability' we mean the learner's ability to converge to the right hypothesis, i.e., the requirement that at a finite stage the answers of the learning

[8]For a discussion of reliability of belief-revision methods see [30].

method stabilize on the correct conjecture. Following the learning-theoretic terminology, we say in this case that the real world has been *identified in the limit*.

DEFINITION 21. *Given an epistemic space* $\mathbb{S} = (S, \mathcal{O})$*, a world* $s \in S$ *is* identified in the limit by a learning method L *if, for every sound and complete data stream for* s*, there exists a finite stage after which* L *outputs the singleton* $\{s\}$ *from then on.*

We say that the epistemic space $\mathbb{S}$ *is* identified in the limit by L *iff all its worlds are identified in the limit by* L*.*

An epistemic space $\mathbb{S}$ *is* identifiable in the limit (learnable) *if there exists a learning method* L *that can identify it in the limit.*

Learning methods differ in their learning power. We are interested in the *most powerful* among them, those that are *universal*—they can learn any epistemic space that is learnable.

DEFINITION 22. *A learning method* L *is* universal on a class $\mathcal{C}$ of epistemic spaces *if it can identify in the limit every epistemic space in* $\mathcal{C}$ *that is identifiable in the limit. A* universal learning method *is one that is universal on the class of all epistemic spaces.*

In the remainder of this paper we focus on learning methods that are generated by iterated belief-revision methods. For brevity we will attribute the ability of identification in the limit also to belief-revision policies.

DEFINITION 23. *An epistemic space* $\mathbb{S}$ *is* identified in the limit by a belief-revision method R *if there exists a prior plausibility assignment* $f_{\preceq}$ *such that the generated learning method* $L_R^{\preceq}$ *identifies* S *in the limit.*

The epistemic space $\mathbb{S}$ *is* standardly identified in the limit *by* R *if there exists a well-founded prior plausibility assignment* $f_{\preceq}$ *(thus inducing a standard plausibility space on* $\mathbb{S}$*) such that* $L_R^{\preceq}$ *identifies* S *in the limit.*

DEFINITION 24. *A revision method* R *is* universal on a class $\mathcal{C}$ of epistemic spaces *if it can identify in the limit every epistemic space in* $\mathcal{C}$ *that is identifiable in the limit.*

R *is* standardly universal on a class $\mathcal{C}$ *if it can standardly identify in the limit every epistemic space in* $\mathcal{C}$ *that is identifiable in the limit.*

Our main result is the existence of AGM-like universal belief-revision methods. The main technical difficulty of this part is the construction of the appropriate prior plausibility order. To define it we use the concept of locking sequences introduced in [13] and that of finite tell-tale sets proposed in [2].

We adjust 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, generates universal learning methods.

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 25. *Let an epistemic space* $\mathbb{S} = (S, \mathcal{O})$*, 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* $s \in \bigcap \mathrm{set}(\sigma)$ *and for each data sequence* τ *with* $s \in \bigcap \mathrm{set}(\tau)$*,* $L(\sigma * \tau) = L(\sigma)$*.*

LEMMA 1. *If a learning method* L *identifies possible world* s *in the limit then there exists a locking sequence for* s *and* L*.*

PROOF. Assume L identifies s without there being a locking sequence for L and s on which L gives $\{s\}$.

First let us consider the case in which a locking sequence σ does exist, but that $L(\sigma)$ is not $\{s\}$. This is clearly absurd, since the L would keep giving an answer $\{t\}$ for $s \neq t$ on any data stream sound and complete wrt s that starts with σ, and hence L would not identify s.

Now it is sufficient to show that a contradiction follows from the assumption that there exists no locking sequence at all. We construct in stages a data stream $\vec{O}$ for s on which L does not converge. Let $x_1, x_2, x_3, \ldots$ enumerate propositions true in s (recall that $\mathcal{O}$ is at most countable).

> *Stage* 1. The string (x_1) is not a locking sequence, so for some τ, sound data sequence for s, $L((x_1) * \tau) \neq L((x_1))$. Take $(x_1) * \tau$ as the initial segment σ_1 of $\vec{O}$.
>
> *Stage* $n+1$. Assume the initial segment σ_n of $\vec{O}$ has been constructed in stage n. By assumption, the sequence $\sigma_n * (x_{n+1})$ is not a locking sequence, so there is a sequence τ sound for s such that $L(\sigma_n * (x_{n+1}) * \tau) \neq L(\sigma_n * (x_{n+1}))$. Take $\sigma_{n+1} = \sigma_n * (x_{n+1}) * \tau$.

Because each x_i occurs in $\vec{O}$, $\vec{O}$ is a sound and complete data stream for s. But learner L keeps changing value on $\vec{O}$, it does not converge. ∎

The characterization of identifiability in the limit requires the existence of finite sets that allow drawing a conclusion without the risk of overgeneralization. The characterization theorem is adapted from [2] and [25].

LEMMA 2. *Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space. $\mathbb{S}$ is identifiable in the limit iff there exists a total map $D : S \to \mathcal{P}(\mathcal{O})$, 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) $s \in \bigcap D_s$,

(3) for any $t \in S$, if $t \in \bigcap D_s$ and $\mathcal{O}_t \subseteq \mathcal{O}_s$, then $t = s$.

PROOF. [⇒] Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space. Recall that S and $\mathcal{O}$ are at most countable. Let us also assume that $\mathbb{S}$ is identifiable in the limit 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 1, 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) $s \in \bigcap D_s$, because $s \in \bigcap \text{set}(\sigma_s)$.

(3) for any $t \in S$, if $t \in \bigcap D_s$ and there is no $p \in \mathcal{O}$ such that $t \in p$ and $s \notin p$, then $t = s$. Assume that there are $s, t \in S$, such that $s \neq t$ and for all $p \in \mathcal{O}$ such that $t \in p$ we have $s \in p$. Let us take a positive sound and complete data stream $\vec{O}$ for t, such that for some $n \in \mathbb{N}$, $\vec{O}[n] = \sigma_s$. Because σ_s is a locking sequence for L on s and $t \in \bigcap \text{set}(\vec{O})$, L converges to s on $\vec{O}$. Therefore, L does not identify t, a space from S. Contradiction.

[⇐] Assume $\mathbb{S} = (S, \mathcal{O})$ has tell-tales. Let us enumerate $S = \{s_0, s_1, \ldots\}$ and define L in the following way:

$$L(\mathbb{S}, \sigma) = \{s_i\} \text{ where } i \text{ is minimal such that } D_{s_i} \subseteq set(\sigma)$$
$$\text{if such } i \text{ exists, otherwise } L(\mathbb{S}, \sigma) = \emptyset.$$

Assume $\vec{O}$ is a sound and complete data stream for s, and that i is the least number for s in the enumeration of S. It is sufficient to show that, for k large enough, $L(\mathbb{S}, \vec{O}[k]) = \{s_i\}$. We can fix n large enough so that $D_{s_i} \subseteq set(\vec{O}[n])$. Nevertheless, we cannot conclude that $L(\mathbb{S}, \vec{O}[n]) = \{s_i\}$, because there may be (finitely many) other worlds $s_1, \ldots, s_m$, with $m < i$ and therefore different from s, that satisfy the same condition. Take any such s_j, $j \in \{0, \ldots, m\}$. We now also have $D_{s_j} \subseteq set(\vec{O}[n]) \subseteq \mathcal{O}_s$, but then, by the properties of tale-tales, there will be a $p \in \mathcal{O}$ such that $s \in p$ and $s_j \notin p$. As $\vec{O}$ is sound and complete wrt s there will be a k_j such that $p \in set(\vec{O}[k_j])$.

We now take k to be the maximum of all the k_j and j, and then only s will satisfy $D_s \subseteq set(\vec{O}[k]) \subseteq \mathcal{O}_s$, and this will remain so for all $n > k$.

This concludes the proof. ∎

We use the notion of finite tell-tales to construct an ordering on S. This way we are able to talk about learnability in the context of plausibility spaces. The aim is to find a way of assigning the prior plausibility order that allows reliable belief revision. We base the construction on finite tell-tales, but we introduce one additional condition (see (2) in Definition 26, below).

DEFINITION 26. *Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space (with countable S and $\mathcal{O}$) with an injective map $i : S \to \mathbb{N}$, and D' be a total map*[9] *such that $D' : S \to \mathcal{P}(\mathcal{O})$, given by $s \mapsto D'_s$ having the following properties:*

(1) D'_s is finite tell-tale for s;

(2) if $t \in \bigcap D'_s$, but $\mathcal{O}_s \nsubseteq \mathcal{O}_t$, then $i(s) < i(t)$.

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

LEMMA 3. *Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space (with countable S and $\mathcal{O}$). If $\mathbb{S}$ is identifiable in the limit, then $\mathbb{S}$ has an ordering tell-tale map.*

PROOF. First let us assume that $\mathbb{S} = (S, \mathcal{O})$, an epistemic space, is identifiable in the limit. Let us then take any injective map $i : S \to \mathbb{N}$ and $j : \mathcal{O} \to \mathbb{N}$ (recall that S and $\mathcal{O}$ are countable). By Lemma 2 we can assume the map D that gives tell-tales for any $s \in S$. On the basis of D we construct a new map $D' : S \to \mathcal{P}(\mathcal{O})$. We proceed step by step according to the enumeration of S given by i and the enumeration of $\mathcal{O}$ given by j (when $i(s) = n$ we will simply write s_n, similarly for j and $p \in \mathcal{O}$).

(1) For s_1 we set $D'_{s_1} := D_{s_1}$.

(2) For s_n we proceed in the following way. First, for every $k < n$ we set P^n_k:

$$P^n_k = \begin{cases} \{p_\ell \mid \ell \text{ is smallest s.t. } s_n \in p_\ell \text{ and } s_k \notin p_\ell\} & \text{if } D_{s_n} \subseteq \mathcal{O}_{s_k}, \\ \emptyset & \text{otherwise.} \end{cases}$$

Finally, we set $D'_{s_n} = D_{s_n} \cup \ (P^n_1 \cup \ldots \cup P^n_{n-1})$.

We now check if D' satisfies the conditions of Definition 26.

[9] We use D', to distinguish from the original tell-tale function D.

(1) For any $s \in S$, D'_s is finite, because D_s is finite, $i(s) = n$ for some $n \in \mathbb{N}$, and there are only finitely many P^n_k such that $k < n$, each of them being either a singleton or an empty set;

(2) moreover, $s_n \in \bigcap D'_{s_n}$, because $s_n \in \bigcap D_{s_n}$ and $s_n \in \bigcap(P^n_1 \cup \ldots \cup P^n_{n-1})$.

(3) For any $s \in S$, if $D'_s \subseteq \mathcal{O}_t \subseteq \mathcal{O}_s$, then $D_s \subseteq D'_s \subseteq \mathcal{O}_t \subseteq \mathcal{O}_s$, and hence, by the definition the finite tell-tale set $t = s$.

It remains to check the condition 4: If $t \in \bigcap D'_s$ and $\mathcal{O}_s \nsubseteq \mathcal{O}_t$ then $i(s){<}i(t)$. Towards contradiction assume that $t \in \bigcap D'_s$, $\mathcal{O}_s \nsubseteq \mathcal{O}_t$, and $i(s){\geq}i(t)$. If $i(s){=}i(t)$, then $\mathcal{O}_s = \mathcal{O}_t$, contradiction. If $i(s){>}i(t)$ then by the construction of D'_s we know that there is a $p \in D'_s$ such that $s \in p$ and $t \notin p$ (if $D_s \subseteq \mathcal{O}_t$ we added such p in the process of obtaining D'_s, otherwise it had been already there to start with). But then $t \notin \bigcap D'_s$. Contradiction. ∎

The next step is to use the ordering tell-tales to define a preorder on an epistemic space.

Definition 27. *For $s, t \in S$, we put*

$$s \preceq^1_{D'} t \text{ iff } t \in \bigcap D'_s.$$

We take $\preceq_{D'}$ to be the transitive closure of the relation $\preceq^1_{D'}$.

We want to show that indeed the above construction generates an order, i.e., that $\preceq_{D'}$ is reflexive, transitive, and antisymmetric. The latter will require proving that $\preceq_{D'}$ includes no proper cycles (see Figure 4).

Definition 28. *A* proper cycle *in $\preceq_{D'}$ is a sequence of distinct worlds $s_1, \ldots, s_n$, with $n \geq 2$, and such that:*

(1) for all $i = 1, \ldots, n-1$, we have $s_i \in \bigcap D'_{s_{i+1}}$,

(2) $s_n \in \bigcap D'_{s_1}$.

Lemma 4. *For any identifiable epistemic space $\mathbb{S}$ and any ordering tell-tale map D', the relation $\preceq_{D'}$ is an order, i.e., $\preceq_{D'}$ is reflexive, transitive, and antisymmetric.*

Proof. The fact that $\preceq_{D'}$ is a preorder is trivial: reflexivity follows from the fact that s is always in $\bigcap D'_s$, and transitivity is imposed by construction (by taking the transitive closure).

We need to prove that $\preceq_{D'}$ is antisymmetric. In order to do that we will show (by induction on n) that $\preceq_{D'}$ does not contain proper cycles of any length $n \geq 2$.

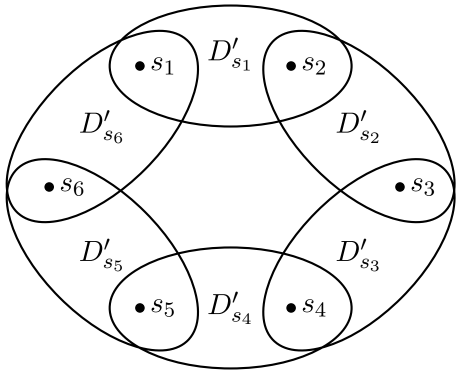

Figure 4. A visualization of a proper cycle of length 6, here we write D'_{s_i} instead of $\bigcap D'_{s_i}$

(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 $s_1 \neq s_2$, $s_2 \in \bigcap D'_{s_1}$, and $s_1 \in \bigcap D'_{s_2}$.

   (a) If $\mathcal{O}_{s_1} \subseteq \mathcal{O}_{s_2}$, then $s_1 \in \bigcap D'_{s_2}$, so (by Condition 1 of Definition 26, i.e., the fact that D'_{s_2} is a tell-tale for s_2), we have that $s_1 = s_2$. Similarly, if $\mathcal{O}_{s_2} \subseteq \mathcal{O}_{s_1}$, then $s_2 \in \bigcap D'_{s_1}$, so again we have that $s_2 = s_1$. Contradiction.

   (b) $\mathcal{O}_{s_1} \nsubseteq \mathcal{O}_{s_2}$ and $\mathcal{O}_{s_2} \nsubseteq \mathcal{O}_{s_1}$. From the assumption that $s_2 \in \bigcap D'_{s_1}$, and that $\mathcal{O}_{s_1} \nsubseteq \mathcal{O}_{s_2}$, we can infer (by Condition 2 of Definition 26), that $i(s_1) < i(s_2)$. But, in the same way (from $s_1 \in \bigcap D'_{s_2}$, and $\mathcal{O}_{s_2} \nsubseteq \mathcal{O}_{s_1}$), 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 that there is no proper cycle of length n, and, towards contradiction, that $s_1, s_2, ..., s_{n+1}$ is a proper cycle of length $n + 1$. We consider two cases:

Case 1: There exists k with $1 \leq k \leq n$ such that $\mathcal{O}_{s_k} \subseteq \mathcal{O}_{s_{k+1}}$. This intuitively means that whatever observable property is made true by s_k is also made true by s_{k+1}. Then, 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 (shorter) proper cycle (of length n). Contradiction.

Case 2: $\mathcal{O}(s_k) \nsubseteq \mathcal{O}(s_{k+1})$, for all $1 \leq k \leq n$. We have that for all $1 \leq k \leq n$, $s_k \in \bigcap D'_{s_{k+1}}$. By Condition 2 of Definition 26, it follows that we have $i(s_k) < i(s_{k+1})$, for all $k = 1, \ldots, n$, and hence $i(s_1) < i(s_{n+1})$. But $s_n \in \bigcap D'_{s_1}$ and $\mathcal{O}(s_{n+1}) \nsubseteq \mathcal{O}(s_1)$ (since otherwise s_{n+1} could be eliminated and $s_1, \ldots, s_n$ would give a proper cycle of length n), hence $i_{s_1} > i_{s_{n+1}}$. Contradiction.

■

We now show that $\preceq_{D'}$, when used by the conditioning revision method, guarantees convergence to the right belief whenever the underlying epistemic space is identifiable in the limit.

Theorem 1. *The conditioning belief-revision method* (`Cond`) *is universal.*

Proof. We have to show that an epistemic space $\mathbb{S}$ is identifiable in the limit iff $\mathbb{S}$ is identifiable in the limit by conditioning. Obviously, if $\mathbb{S}$ is identifiable in the limit by conditioning, then $\mathbb{S}$ is identifiable in the limit. We therefore focus on the other direction, i.e., we show that if $\mathbb{S}$ is identifiable in the limit by any learning method, then it is identifiable in the limit by conditioning.

By Lemma 3 we know there exists an ordering tell-tale map for $\mathbb{S}$ and by Lemma 4, the corresponding $\preceq_{D'}$ is a (partial) order on S. By the Order-Extension Principle, every partial order $\preceq_{D'}$ on a set S can be extended to a total order on the same set, i.e., there exists a total order $\preceq$ on S such that, for all $s, t \in S$, we have that $s \preceq_{D'} t$ implies $s \preceq t$.[10]

It remains to show that $\mathbb{S}$ is identifiable in the limit by the learning method generated from the conditioning belief-revision method and the prior plausibility assignment $\preceq$. Let $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ and let us take any $s \in S$ and the corresponding D'_s. Since $s \in \bigcap D'_s$, it follows that for every sound and complete positive data stream $\vec{O}$ for s, there exists $n \in \mathbb{N}$ such that $D'_s \subseteq \text{set}(\vec{O}[n])$. Let $\texttt{Cond}(\mathbb{B}_{\mathbb{S}}, \vec{O}[n]) = (S', \mathcal{O}, \preceq')$. Our aim is now to demonstrate that $\min_{\preceq'} S' = \{s\}$. By the antisymmetry of the order relation $\preceq$ and hence also of $\preceq'$, the minimal element of S' is unique, so it is sufficient to show that $s \in \min_{\preceq} S'$. For this, let $t \in S'$ be arbitrary. We need to show that $s \preceq t$. Since $t \in S'$, we get that $D'_s \subseteq \text{set}(\vec{O}[n]) \subseteq \mathcal{O}_t$, so, by Definition 26, we have $s \preceq_{D'} t$, and hence $s \preceq t$. It remains to show that `Cond` stabilizes on $\{s\}$. Observe that $\vec{O}$ is sound with respect to s, and therefore no further information from $\vec{O}$ can eliminate s (because conditioning is conservative), and hence for any future data the set of minimal elements will remain $\{s\}$. ■

Theorem 2. *The lexicographic belief-revision method* (`Lex`) *is universal.*

[10] In general, the proof of this principle uses the Axiom of Choice. But here we only need the special case in which the support set S is *countable*, and this special case is provable without the Axiom of Choice.

The proof is analogous to the proof of Theorem 1. Within our learning setting lexicographic revision with true information does exactly what conditioning does. The only difference is that the rest of the doxastic structure might not stabilize, but only the minimal elements stabilize.

THEOREM 3. *The minimal belief-revision method* (`Mini`) *is not universal.*

PROOF. Let us give a counter-example, an epistemic space that is identifiable in the limit, but is not identifiable by the minimal revision method (see Figure 5). Let $\mathbb{S} = (S, \mathcal{O})$, where $S = \{s_1, s_2, s_3\}$, $\mathcal{O} = \{p, q\}$, and $p = \{s_1, s_3\}$, $q = \{s_2, s_3\}$. The epistemic space $\mathbb{S}$ is identifiable in the limit by the conditioning revision method: just assume the ordering $s_1 \prec s_2 \prec s_3$. However, there is no ordering that allows identification in the limit of S by the minimal revision method. If s_3 occurs in the ordering before s_1 (or before s_2), then the minimal revision method fails to identify s_1 (s_2, respectively). If both s_1 and s_2 precede s_3 in the ordering then the minimal revision method fails to identify s_3 on any data stream consisting of singletons of propositions from s_3. On all such data streams for s_3 the minimal state will alternate between s_1 and s_2, or stabilize on one of them. The last case is that at least one of s_1 and s_2 is equi-plausibile to s_3. In such case s_3 is not identifiable because for any single proposition from s_3 there is more than one possible world consistent with it. ■

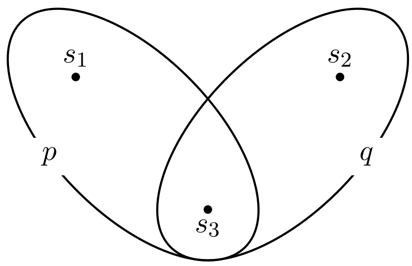

Figure 5. Epistemic space from the proof of Theorem 3. It is impossible to find a plausibility order that would allow learnability via minimal belief-revision method, `Mini`.

THEOREM 4. *No conservative belief-revision method is standardly universal.*

PROOF. There is an epistemic space $\mathbb{S}$ that is identifiable in the limit by a learning method, but is not standardly identified in the limit by any conservative belief-revision method. The following epistemic space constitutes such counter-example. Let $\mathbb{S} = (S, \mathcal{O})$ such that $S = \{s_n \mid n \in \mathbb{N}\}$, $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k = \{s_i \mid 0 \leq i \leq k\}$, see Figure 6.

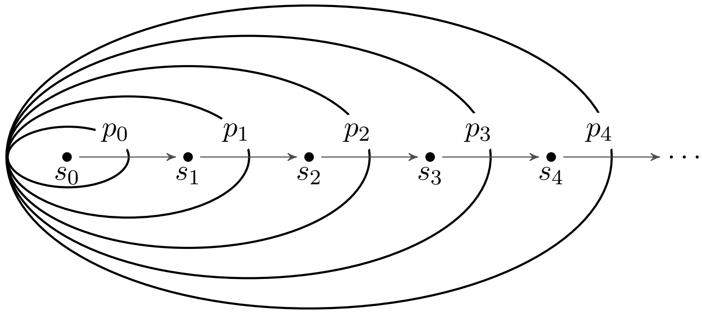

Figure 6. Epistemic space from the proof of Theorem 4. The grey arrows show the non-well-founded plausibility order appropriate for successful, non-standard learning.

$\mathbb{S}$ is identifiable in the limit[11] by the following learning method L:

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

Let us now assume (towards contradiction) that $\mathbb{S}$ is standardly identifiable in the limit by a conservative belief-revision method R, i.e., there exists a well-founded total preorder $\preceq$ on S, such that the learning method $L_R^{\preceq}$ generated from R and $\preceq$ identifies S in the limit and is conservative.

If $\preceq$ is well-founded we can choose some minimal $s_k \in \min_{\preceq} S$ and set $L_R^{\preceq}(S, \lambda) = \{s_k\}$, where λ is the empty data sequence. Take now some $m > k$, and notice that $\mathcal{O}_{s_m} \subset \mathcal{O}_{s_k}$ (by our construction of $\mathbb{S}$). Let $\vec{O}$ by a sound and complete data stream for s_m. By assumption, $L_R^{\preceq}$ identifies s_m in the limit, hence there must exists some k such that $L_R^{\preceq}(S, \vec{O}[k]) = \{s_m\}$. But since $\mathcal{O}_{s_m} \subseteq \mathcal{O}_{s_k}$, the stream $\vec{O}$ is sound for s_k as well: $\text{set}(\vec{O}[n]) \subseteq \mathcal{O}_{s_k}$, for all $n \in \mathbb{N}$.

We prove by induction that $s_k \in L_R^{\preceq}(S, \vec{O}[n])$ for all $n \in \mathbb{N}$. Note that this leads to contradiction, namely to $s_k \in L_R^{\preceq}(S, \vec{O}[n]) = \{s_m\}$, and hence to $s_k = s_m$, which contradicts our choice of $m > k$. The proof by induction goes as follows. The base case is already established, since $s_k \in L(\mathbb{S}, \lambda) = L_R^{\preceq}(\mathbb{S}, \vec{O}[0])$. For the inductive case, let us assume that $s_k \in L_R^{\preceq}(\mathbb{S}, \vec{O}[n])$ for

[11] And, as a consequence of our previous results, it is identifiable in the limit by conditioning. Indeed, it is enough take the prior plausibility given by: $s_n \preceq s_m$ iff $n \geq m$. But notice that $\preceq$ is not well-founded, so this is not a standard prior.

some n, then $\mathrm{set}(\vec{O}[n+1]) \subseteq \mathcal{O}_{s_k}$ and L is conservative to conclude that $s_k \in L_R^{\preceq}(\mathbb{S}, \vec{O}[n] * \vec{O}_{n+1}) = L_R^{\preceq}(\mathbb{S}, \vec{O}[n+1])$. ■

The above result concerns the type of preorders that facilitate identifiability in the limit. We show that for universality results our non-standard setting (involving non-well-founded plausibility orders) is essential: assuming that AGM-like belief revision must be conservative, no such method is universal with respect to well-founded plausibility spaces.

Corollary 1. *No AGM-like belief-revision method is standardly universal.*

5. Learning from positive and negative data

One may wonder what would happen if the revision process was governed not only by arbitrary sets of observable properties, but by observables which are closed under certain logical operations. One simple adjustment of the set $\mathcal{O}$ is to assume its closure on negation.[12]

First let us extend our framework to account for situations in which both positive and negative data can be observed.

Definition 29. *An epistemic space* $\mathbb{S} = (S, \mathcal{O})$ *is* negation-closed *if the set* $\mathcal{O}$ *of all data is negation-closed, i.e., if for every* $p \in \mathcal{O}$ *there exists some* $\overline{p} \in \mathcal{O}$ *such that* $p = S \setminus \overline{p}$ *(i.e., for every* $s \in S$*, we have* $s \in \overline{p}$ *iff* $s \notin p$*).*

Proposition 10. *Conditioning and lexicographic revision generate standardly universal learning methods on the class of negation-closed epistemic spaces.*

Proof. We prove that every negation-closed epistemic space $\mathbb{S}$ (with $\mathcal{O}$ and S countable) is identifiable in the limit by conditioning and by lexicographic revision.

Let us assume that $\mathbb{S}$ is countable and negation closed. In fact, any ω-type order $\preceq$ on S gives a suitable (well-founded) prior plausibility assignment. Let us take an $s \in S$. Since $\preceq$ is ω-type it is well-founded, so there are only finitely many worlds that are more plausible than s. For each such world $t \prec s$ we collect one $O^t \in \mathcal{O}$ such that $s \in O^t$ but $t \notin O^t$ or *vice versa*. Such an O^t must exist, since $s \neq t$ implies that either there exists some $O \in \mathcal{O}$ such that $s \in O$ and $t \notin O$ or there exists some $O' \in \mathcal{O}$ such that $t \in O'$ and $s \notin O'$. In the first case, we put $O^t = O$, while in the second case

[12] In a follow up work we will consider closure on finite intersections, which will allow us to view the learnability on epistemic spaces as topological properties.

we put $O^t = \overline{O}$.) Then the data set $\{O^t \mid t \prec s\}$ is finite. For every data stream $\vec{O}$ that is sound and complete with respect to s, there must exist a stage $n \in \mathbb{N}$ by which all data in $\{O^t \mid t \prec s\}$ have been observed. After this stage, all worlds that are more plausible than s will have been deleted (in the case of conditioning) or will have become less plausible than s (in the case of lexicographic revision), so from then on the (only) most plausible state is s. Hence conditioning and lexicographic revision identify any world $s \in S$ in the limit. ■

PROPOSITION 11. *Minimal revision is not universal on negation-closed spaces.*

PROOF. We show a negation-closed epistemic space that is identifiable in the limit, but is not identifiable in the limit by the minimal revision method. We take $\mathbb{S} = (S, \mathcal{O})$, where: $S = \{s_1, s_2, s_3, s_4\}$ and $\mathcal{O} = \{p, q, r\}$, with $p = \{s_2, s_4\}$, $q = \{s_3, s_4\}$, and $r = \{s_1, s_2, s_3, s_4\}$, see Figure 7.

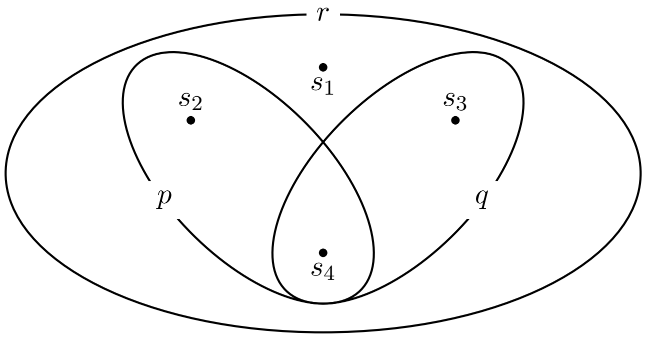

Figure 7. Epistemic space from the proof of Proposition 11

The epistemic space $\mathbb{S}$ is identifiable in the limit by `Cond`, just assume the plausibility order $s_1 \prec s_2 \prec s_3 \prec s_4$. However, there is no plausibility order that allows identification in the limit of $\mathbb{S}$ by the minimal revision method. Whichever ordering is assumed, the least plausible element will not be identifiable. It is so because each piece of data consistent with s is also consistent with one of the $\prec$-smaller sets. ■

6. Erroneous information

With the introduction of negative information, we can now allow for occasional observational *errors*, and for their corrections. To consider erroneous data we now give up the soundness of data streams, i.e., we allow that the learner can observe data that may be false in the real world. In order to still give the agent a chance to learn the real world, we need to impose some limitation on errors. We do this by requiring the data streams to be 'fair'.[13]

DEFINITION 30. *Let* $\mathbb{S} = (S, \mathcal{O})$ *be a negation-closed epistemic space. A stream* $\vec{O}$ *of data from* $\mathcal{O}$ *is* 'fair' *with respect to the world s if* $\vec{O}$ contains only finitely many errors and every such error is eventually corrected *in* $\vec{O}$*, in other words:*

- $\vec{O}$ *is complete with respect to* s,
- *there is* $n \in \mathbb{N}$ *such that for all* $k \geq n$, $s \in O_k$, *and*
- *for every* $i \in \mathbb{N}$ *such that* $s \notin O_i$, *for some* $k > i$ *we have* $O_k = \overline{O_i}$.

Unsurprisingly, conditioning (which assumes absolute veracity of the new observations) is no longer a good strategy. If erroneous observations are possible, then eliminating worlds that do not satisfy these observations is risky and irrational.

PROPOSITION 12. *Conditioning and minimal revision are not universal for fair streams.*

PROOF. Conditioning does not tolerate errors at all. On any $\vec{O}_i$ such that $s \notin \vec{O}_i$ conditioning will remove s and there is no way to revive it. Minimal revision, as it has been shown, is not universal on negation-closed epistemic spaces even with respect sound and complete data streams, which are a special case of fair streams. ∎

We will demonstrate that lexicographic revision deals with errors in a skillful manner. Before we get to that we introduce and discuss the notion of *propositional upgrade* [which is a special case of generalized upgrade, see 10]. Such an upgrade is a transformation of a plausibility space 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 a plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ by changing the preorder $\preceq$ 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 $\preceq$ is kept the same. In particular, our lexicographic revision is a special case of such propositional upgrade, $(\neg p, p)$.

[13]Notions defined in Section 4 (identifiability in the limit, universality, etc.) are similar for fair data streams.

LEMMA 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 $X * Y$ is equivalent to the following propositional upgrade:

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

To show this let us take an arbitrary plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ and apply upgrades X and Y successively. First, we apply to the upgrade X, and we obtain the new preorder $\preceq^X$, in which all worlds satisfying x_1 are less plausible than all x_2-worlds, etc., and within each such partition the old order $\preceq$ is kept the same. Now, to this new plausibility space we apply the second upgrade, Y, obtaining the new preorder $\preceq^{XY}$, 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 $\preceq^X$ we also know that the new preorder $\preceq^{XY}$ 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 $\preceq$ is maintained. Putting these together, we get that $\preceq^{XY}$ has the following structure:

$$(x_1 \cap y_1) \succeq^{XY} \ldots \succeq^{XY} (x_n \cap y_1) \succeq^{XY}$$
$$(x_1 \cap y_2) \succeq^{XY} \ldots \succeq^{XY} (x_n \cap y_2) \succeq^{XY} \ldots \succeq^{XY} \ldots \succeq^{XY} (x_n \cap y_m),$$

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

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

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

■

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

PROPOSITION 13. *Lexicographic revision generates a standardly universal belief-revision-based learning method for fair streams on the class of negation-closed epistemic spaces.*

PROOF. First let us recall that lexicographic revision, `Lex`, is standardly universal for sound and complete streams on negation-closed spaces. It is left to show that `Lex` retains its power on fair streams. It is sufficient to show that lexicographic revision is 'error-correcting', i.e., that the effect of revising with the stream $(p * \sigma * \overline{p})$ is exactly the same as with the stream $(\sigma * \overline{p})$, where σ is any sequence of observables. 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

$$(\neg p, p)(\neg\sigma_1, \sigma_1) \dots (\neg\sigma_n, \sigma_n)(p, \neg p)$$

is equivalent to

$$(\neg\sigma_1, \sigma_1) \dots (\neg\sigma_n, \sigma_n)(p, \neg p).$$

From Lemma 5 we know that propositional upgrade is closed under sequential composition. Hence, in the equivalence to be shown, we can replace the composition $(\neg\sigma_1, \sigma_1) \dots (\neg\sigma_n, \sigma_n)$ by only one generalized upgrade, which we denote by $(x_1, \dots, x_m)$. Now, we have to show that: $(\neg p, p)(x_1, \dots, x_m)(p, \neg p)$ is equivalent to: $(x_1, \dots, x_m)(p, \neg p)$.

By the proof of Lemma 5, the composition $(x_1, \dots, x_n)(p, \neg p)$ has the following form:

$$(x_1 \wedge p, \dots, x_n \wedge p, x_1 \wedge \neg p, \dots, x_n \wedge \neg p).$$

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

$$(\neg p \wedge x_1 \wedge p, p \wedge x_1 \wedge p, \dots, \neg p \wedge x_n \wedge p, p \wedge x_n \wedge p, \neg p \wedge x_1 \wedge \neg p, p \wedge x_1 \wedge \neg p, \dots,$$
$$\neg p \wedge x_n \wedge \neg p, p \wedge x_n \wedge \neg p).$$

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 plausibility space. We obtain:

$$(x_1 \wedge p, \dots, x_n \wedge p, x_1 \wedge \neg p, \dots, x_n \wedge \neg p).$$

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

7. Conclusions and perspectives

We have considered iterated belief-revision policies of conditioning, lexicographic, and minimal belief revision. We have identified certain features of those methods relevant in the context of iterated revision, especially data-retention and conservativity turned out to be very important. We defined learning methods based on those revision policies and we have shown how the aforementioned properties influence the learning process. Throughout the paper 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 and on-going work consists of multi-level investigation of the relationship between formal learning theory, belief revision theory, and DEL. There surely are many links still to be found. 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.

Acknowledgements. The research of Nina Gierasimczuk is funded by an Innovational Research Incentives Scheme Veni grant 275-20-043, Netherlands Organisation for Scientic Research (NWO). Nina Gierasimczuk would also like to express her thanks to Jakub Szymanik and Umberto Grandi for their useful comments.

The contribution of Sonja Smets is funded in part by an Innovational Research Incentives Scheme Vidi grant from the Netherlands Organisation for Scientic Research (NWO) and by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement no. 283963.

References

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

[2] Angluin, D., 'Inductive inference of formal languages from positive data', *Information and Control*, 45(2) (1980), 117–135.

[3] Aucher, G., 'A combined system for update logic and belief revision', in M. W. Barley, and N. Kasabov, (eds.), *Intelligent Agents and Multi-Agent Systems*, vol. 3371 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2005, pp. 1–17.

[4] Baltag, A., N. Gierasimczuk, and S. Smets, 'Belief revision as a truth-tracking process', in K. Apt, (ed.), *TARK'11: Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge, Groningen, The Netherlands, July 12-14, 2011*, ACM, New York, NY, USA, 2011, pp. 187–190.

[5] Baltag, A., and L. Moss, 'Logics for epistemic programs', *Synthese*, 139 (2004), 2, 165–224.

[6] Baltag, A., L. S. Moss, and S. Solecki, 'The logic of public announcements, common knowledge, and private suspicions', in I. Gilboa, (ed.), *TARK'98: Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1998, pp. 43–56.

[7] Baltag, A., and S. Smets, 'Dynamic belief revision over multi-agent plausibility models', in G. Bonanno, W. van der Hoek, and M. Wooldridge, (eds.), *LOFT'06: Proceedings of 7th Conference on Logic and the Foundations of Game and Decision Theory*, University of Liverpool, 2006, pp. 11–24.

[8] Baltag, A., and S. Smets, 'The logic of conditional doxastic actions', in R. van Rooij, and K. Apt, (eds.), *New Perspectives on Games and Interaction. Texts in Logic and Games*, Amsterdam University Press, 2008, pp. 9–31.

[9] Baltag, A., and S. Smets, 'A qualitative theory of dynamic interactive belief revision', in G. Bonanno, W. van der Hoek, and M. Wooldridge, (eds.), *LOFT'08: Proceedings of 8th Conference on Logic and the Foundations of Game and Decision Theory*, no. 3 in Texts in Logic and Games, Amsterdam University Press, 2008, pp. 9–58.

[10] Baltag, A., and S. Smets, 'Learning by questions and answers: From belief-revision cycles to doxastic fixed points', in H. Ono, M. Kanazawa, and R. de Queiroz, (eds.), *Logic, Language, Information and Computation*, vol. 5514 of *Lecture Notes in Computer Science*, Springer, 2009, pp. 124–139.

[11] van Benthem, J., 'Dynamic logic for belief revision', *Journal of Applied Non-Classical Logics*, 2 (2007), 129–155.

[12] van Benthem, J., *Logical Dynamics of Information and Interaction*, Cambridge University Press, Cambridge, MA, USA, 2011.

[13] Blum, L., and M. Blum, 'Toward a mathematical theory of inductive inference', *Information and Control*, 28 (1975), 125–155.

[14] Boutilier, C., 'Iterated revision and minimal change of conditional beliefs', *Journal of Philosophical Logic*, 25 (1996), 263–305.

[15] Dégremont, C., and N. Gierasimczuk, 'Can doxastic agents learn? On the temporal structure of learning', in X. He, J. F. Horty, and E. Pacuit, (eds.), *LORI'09: Proceedings of 2nd International Workshop on Logic, Rationality, and Interaction, Chongqing, China, October 8-11, 2009*, vol. 5834 of *Lecture Notes in Computer Science*, Springer, 2009, pp. 90–104.

[16] Dégremont, C., and N. Gierasimczuk, 'Finite identification from the viewpoint of epistemic update', *Information and Computation*, 209 (2011), 3, 383–396.

[17] van Ditmarsch, H., 'Prolegomena to dynamic logic for belief revision', *Knowledge, Rationality & Action (Synthese)*, 147 (2005), 229–275 (41–87).

[18] van Ditmarsch, H., W. Van der Hoek, and B. Kooi, *Dynamic Epistemic Logic*, Springer, The Netherlands, 2007.

[19] Gierasimczuk, N., 'Bridging learning theory and dynamic epistemic logic', *Synthese*, 169 (2009), 2, 371–384.

[20] Gierasimczuk, N., 'Learning by erasing in dynamic epistemic logic', in A. H. Dediu, A. M. Ionescu, and C. Martin-Vide, (eds.), *LATA'09: Proceedings of 3rd International Conference on Language and Automata Theory and Applications, Tarragona, Spain, April 2-8, 2009*, vol. 5457 of *Lecture Notes in Computer Science*, Springer, The Netherlands, 2009, pp. 362–373.

[21] Gierasimczuk, N., *Knowing One's Limits. Logical Analysis of Inductive Inference*, Ph.D. thesis, Universiteit van Amsterdam, The Netherlands, 2010.

[22] Gierasimczuk, N., D. de Jongh, and V. F. Hendricks, 'Logic and learning', in A. Baltag, and S. Smets, (eds.), *Johan van Benthem on Logical and Informational Dynamics*, Springer, 2014.

[23] Gierasimczuk, N., and D. de Jongh, 'On the complexity of conclusive update', *The Computer Journal*, 56 (2013), 3, 365–377.

[24] Gold, E. M., 'Language identification in the limit', *Information and Control*, 10 (1967), 447–474.

[25] Jain, S., D. Osherson, J. S. Royer, and A. Sharma, *Systems that Learn*, MIT Press, Chicago, 1999.

[26] Kelly, K. T., 'Iterated belief revision, reliability, and inductive amnesia', *Erkenntnis*, 50 (1998), 11–58.

[27] Kelly, K. T., 'The learning power of belief revision', in *TARK'98: Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1998, pp. 111–124.

[28] Kelly, K. T., 'Learning theory and epistemology', in I. Niiniluoto, M. Sintonen, and J. Smolenski, (eds.), *Handbook of Epistemology*, Dordrecht: Kluwer, 2004, pp. 183–203.

[29] Kelly, K. T., 'Ockham's razor, truth, and information', in P. Adriaans, and J. van Benthem, (eds.), *Handbook of the Philosophy of Information*, Elsevier, 2008, pp. 321–359.

[30] Kelly, K. T., O. Schulte, and V. Hendricks, 'Reliable belief revision', in *Proceedings of the 10th International Congress of Logic, Methodology, and Philosophy of Science*, Kluwer Academic Publishers, 1995, pp. 383–398.

[31] Martin, E., and D. Osherson, 'Scientific discovery based on belief revision', *The Journal of Symbolic Logic*, 62 (1997), 4, 1352–1370.

[32] Martin, E., and D. Osherson, *Elements of Scientific Inquiry*, MIT Press, Cambridge, MA, USA, 1998.

[33] Osherson, D., M. Stob, and S. Weinstein, *Systems that Learn*, M.I.T. Press, Cambridge, MA, USA, 1986.

[34] Rott, H., 'Conditionals and theory change: Revisions, expansions and additions', *Synthese*, 81 (1989), 91–113.

[35] Spohn, W., 'Ordinal conditional functions: A dynamic theory of epistemic states', in B. Skyrms, and W. L. Harper, (eds.), *Causation in Decision, Belief Change, and Statistics*, vol. II, Dordrecht: Kluwer, 1988, pp. 105–134.

Alexandru Baltag
Nina Gierasimczuk
Sonja Smets
Institute for Logic, Language, and Computation
University of Amsterdam
Science Park 107
Amsterdam, The Netherlands
nina.gierasimczuk@gmail.com