



UvA-DARE (Digital Academic Repository)

The Logic of AGM Learning from Partial Observations

Baltag, A.; Özgün, A.; Vargas-Sandoval, A.L.

DOI

[10.1007/978-3-030-38808-9_3](https://doi.org/10.1007/978-3-030-38808-9_3)

Publication date

2020

Document Version

Final published version

Published in

Dynamic Logic : New Trends and Applications

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Baltag, A., Özgün, A., & Vargas-Sandoval, A. L. (2020). The Logic of AGM Learning from Partial Observations. In L. Soares Barbosa, & A. Baltag (Eds.), *Dynamic Logic : New Trends and Applications: Second International Workshop, DaLi 2019, Porto, Portugal, October 7–11, 2019 : proceedings* (pp. 35-52). (Lecture Notes in Computer Science; Vol. 12005). Springer. https://doi.org/10.1007/978-3-030-38808-9_3

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.



The Logic of AGM Learning from Partial Observations

Alexandru Baltag¹, Aybüke Özgün^{1,2}, and Ana Lucia Vargas-Sandoval¹(✉)

¹ ILLC, University of Amsterdam, Amsterdam, The Netherlands
ana.varsa@gmail.com

² Arché, University of St. Andrews, St Andrews, UK

Abstract. We present a dynamic logic for inductive learning from partial observations by a “rational” learner, that obeys AGM postulates for belief revision. We apply our logic to an example, showing how various concrete properties can be learnt with certainty or inductively by such an AGM learner. We present a sound and complete axiomatization, based on a combination of relational and neighbourhood version of the canonical model method.

1 Introduction

In this paper, we extend our previous work [3], presented at DaLi 2017, in which we introduced a dynamic logic for learning theory, building on our past work [4, 5] (that bridged Formal Learning Theory and Dynamic Epistemic Logic in a topological setting): a learner forms conjectures based on a continuous stream of observations, with the goal of inductively converging to a true conjecture. To reason about this framework, we added to Subset Space Logics [12, 16] dynamic *observation modalities* $[o]\varphi$, as well as a *learning operator* $L(\vec{o})$, which encodes the learner’s *conjecture* after observing a finite sequence of data \vec{o} . In [3], we completely axiomatized this logic, and used it to characterize various epistemological and learning-theoretic notions.

However, the learner in [3] was assumed to satisfy only very few rationality constraints (essentially, only consistency of conjectures, and the Success postulate requiring that the conjectures fit the evidence). In contrast, in this paper we focus on fully rational learners, whose conjectures obey all the AGM postulates for belief revision [1]. Semantically, such an “AGM learner” comes with a family of nested Grove spheres (encoding the agent’s defaults and her belief-revision policy), or equivalently with a total plausibility (pre)order on the set of possible worlds. After observing some evidence, the learner forms a conjecture by applying “AGM conditioning”: essentially, her conjecture encompasses the most plausible worlds that fit the evidence. This belief dynamics is non-monotonic, but only minimally so: it respects the principle of Rational Monotonicity (equivalent to some of the AGM postulates in [1]), requiring that the dynamics is just monotonic logical updating (the so-called “expansion”, putting together the old

conjecture with the new evidence) whenever the old conjecture is consistent with the new evidence.

Our aim is to realize the same program for such AGM learners as the one we achieved in [3]. There are many reasons to focus on AGM learners: first, AGM postulates seem inherently plausible, or at least strongly desirable as constraints on rational learners’ belief dynamics. Second, imposing such constraints does not lead to any loss in learning power: as shown in [4], AGM conditioning is a “universal learning method”: any questions that can be inductively solved (or solved with certainty) by some learner can also be solved by an AGM learner. Third, the additional constraints posed by the AGM postulates make the logic of inductive AGM learning more interesting, and its completeness more challenging, than the logic of unconstrained learners.

And indeed, as it turns out, forming conjectures *only based on a sequence of direct, complete observations* (as is standard in Learning Theory, and as we also assumed in [3]) does not seem to be enough to allow us to characterize AGM learning! In order to obtain our completeness result, we had to extend the domain of our learning functions to *partial observations*: incomplete reports of a full-fledged observation, equivalent to finite disjunctions of observations. Technically, we had to move from the framework of *intersection spaces* adopted in [3] (in which the observable properties were closed under finite intersections, to capture the effect of successive observations) to the one of *lattice spaces* (in which closure under finite unions is also required, to capture the effect of partial observations). At the syntactic level, this led us to replace simple observations by *observational events*: like PDL programs, these are built from simple observations $!o$, using sequential composition $e; e'$ (to represent successive observations) and epistemic non-determinism $e \sqcup e'$ (to capture the receipt of partial information, after which the agent is not sure which of the two observations e, e' has been made). After an observational event e , the learner forms a conjecture $L(e)$, obtained by applying AGM conditioning (with respect to her plausibility order \leq) to the event’s informational content $pre(e)$ (its “precondition”, defined recursively by taking conjunctions of the preconditions in a sequential composition $e; e'$, and disjunctions of the preconditions in an epistemic non-determinism $e \sqcup e'$).

As in Subset Space Logics [12, 16], our language features an S5-type ‘*infallible knowledge*’ modality, capturing the learner’s *hard information*, as well as the so-called ‘*effort*’ modality, which we interpret as ‘stable truth’ (i.e., truth that will resist further observations). As in [3], we add *dynamic modalities* $[e]\varphi$, this time capturing *updates after observational events* (“ φ becomes true after event e ”), similarly to the role of dynamic modalities in Propositional Dynamic Logic (PDL) and especially in Dynamic Epistemic Logic (DEL).¹ Finally, we have an *AGM learning operator* $L(e)$, which encodes the AGM learner’s conjecture (her “strongest belief”) given an observational event e . As in [3], these can be used to give natural definitions of *belief*, *stable (undefeated) belief*, *inductive*

¹ Indeed, our observational events can be seen as corresponding to a special type of (single-agent) epistemic events in the so-called BMS style.

knowledge and *inductive learnability*. We begin the study of the expressivity of this language, and we apply it to an example, showing how these notions work on specific AGM learners.

Though our completeness proof uses some standard techniques in non-monotonic and conditional logics, there are some important differences. First, since we don't allow conditioning on arbitrary formulas, but only on those corresponding to (preconditions of) observational events, the proof is more subtle. In particular, it shows that AGM has no need for conditioning on negated formulas. Second, the completeness proof uses a mixture of relational and neighbourhood versions of the standard canonical model construction, with further complications due to the presence of the “effort” modality. As in [3], its connection with dynamic updates is embodied by our Effort Axiom and Effort Rule, which together say that a proposition φ is “stably true” iff its truth is preserved by every correct observational event. The presence of fresh observational variables as “witnesses” of stability of φ in the Effort Rule requires the restriction of the canonical model to “witnessed” theories (rather than all maximally consistent ones).

Due to space limitations, proofs are omitted from the main body and presented in the Appendix of the longer version, available online at <https://analuciavargassan.com/page/>.

2 Syntax and Semantics

Let $\text{Prop} = \{p, q, \dots\}$ be a countable set of *propositional variables*, denoting arbitrary ‘ontic’ (i.e., non-epistemic) facts and $\text{Prop}_\theta = \{o, u, v, \dots\}$ a countable set of *observational variables*, denoting ‘observable facts’.

Observational Events. We consider *observational events* e (or, in short, *observations*) by which the agent acquires some evidence about the world. We denote the set of all observational events by Π_{Ob} and define it by the following recursive clauses:

$$e := !\top \mid !o \mid e; e' \mid e \sqcup e'$$

where $o \in \text{Prop}_\theta$. Intuitively: for every observational variable o , we have a primitive observational event, denoted by $!o$, corresponding to *the event of observing variable o* . We also denote by $!\top$ the *null event* (in which *no new observation has taken place yet*). Observational events are naturally closed under regular operations on programs, of which we consider only two: $e; e'$ represents *sequential composition* of observational events (first observation e is made then observation e' is made); while $e \sqcup e'$ captures *epistemic non-determinism*: one of the two observational events e or e' happens, but the observing agent is uncertain which of the two. The last construct can be used to represent *partial observations*, including indirect evidence obtained from other agents' reports: the agent observes (or is told) only some feature of the evidence, so her information is compatible with multiple fully-determined events.

The Language of AGM Learning. The dynamic language \mathcal{L} of AGM learning from partial observations is defined recursively as

$$\varphi := p \mid o \mid \neg\varphi \mid \varphi \wedge \varphi \mid L(e) \mid K\varphi \mid [e]\varphi \mid \Box\varphi$$

where $p \in \text{Prop}$, $o \in \text{Prop}_\theta$, and $e \in \Pi_{Ob}$. We employ the usual abbreviations for propositional connectives $\top, \perp, \vee, \rightarrow, \leftrightarrow$, and $\langle K \rangle\varphi$, $\langle e \rangle\varphi$ and $\diamond\varphi$ denote $\neg K\neg\varphi$, $\neg[e]\neg\varphi$, and $\neg\Box\neg\varphi$, respectively. Given a formula $\varphi \in \mathcal{L}$, we denote by O_φ and O_e the set of all observational variables occurring in φ and e , respectively.

Intuitively, $L(e)$ denotes the learner’s *conjecture* given observation e ; i.e., her “strongest belief” after having performed observational event e . We read $K\varphi$ as ‘the learner *knows* φ (with absolute certainty)’. The operator $[e]\varphi$ is similar to the update operator in Public Announcement Logic: we read $[e]\varphi$ as ‘after event e is observed, φ holds’. Finally, \Box is the so-called ‘effort modality’ from Subset Space Logic [12, 16]; we read $\Box\varphi$ as ‘ φ is *stably true*’ (i.e. it is true and will stay true under any further observations).

We interpret \mathcal{L} on *plausibility learning models* in the style of subset space semantics, as given in turn.

Definition 1 (Plausibility Learning Frame/Model). A plausibility learning frame is a triple (X, \mathcal{O}, \leq) , where: X is a non-empty set of possible worlds (or ‘ontic states’); $\mathcal{O} \subseteq \mathcal{P}(X)$ is a non-empty set of subsets, called information states (or ‘partial observations’, or ‘evidence’), which is assumed to be closed under finite intersections and finite unions: if $\mathcal{F} \subseteq \mathcal{O}$ is finite then $\bigcap \mathcal{F} \in \mathcal{O}$ and $\bigcup \mathcal{F} \in \mathcal{O}$; and \leq is a total preorder² on X , called plausibility order and satisfying the observational version of Lewis’ ‘Limit Condition’: every non-empty information state O has maximal elements. More precisely, if for any evidence $O \in \mathcal{O}$, we put³

$$\text{Max}_{\leq}(O) := \{x \in O : y \leq x \text{ for all } y \in O\}$$

for the set of maximal (“most plausible”) worlds compatible with the evidence, then the Limit Condition requires that $\text{Max}_{\leq}(O) \neq \emptyset$ whenever $O \neq \emptyset$. The pair (X, \mathcal{O}) is known in the literature as a “lattice frame” [12, 16], while $x \leq y$ is read as ‘world y is at least as plausible as world x ’.

A plausibility learning model $M = (X, \mathcal{O}, \leq, \|\cdot\|)$ consists of a plausibility learning frame (X, \mathcal{O}, \leq) , together with a valuation map $\|\cdot\| : \text{Prop} \cup \text{Prop}_\theta \rightarrow \mathcal{P}(X)$ that maps propositional variables p into arbitrary sets $\|p\| \subseteq X$ and observational variables o into information states $\|o\| \in \mathcal{O}$.

A learner $\mathbb{L}_{\leq} : \mathcal{O} \rightarrow \mathcal{P}(X)$ on a plausibility lattice frame (X, \mathcal{O}, \leq) is a function that maps to every information state $O \in \mathcal{O}$ some ‘conjecture’ $\mathbb{L}_{\leq}(O) \subseteq X$. An *AGM-learner* is a learner who, upon having observed $O \in \mathcal{O}$, always

² A total preorder \leq on X is a reflexive and transitive binary relation such that every two points are comparable: for all $x, y \in X$, either $x \leq y$ or $y \leq x$ (or both).

³ Since \leq is a total preorder, this definition coincides with the standard definition of maximal elements as $\text{Max}_{\leq}(O) := \{x \in O : \forall y \in O (x \leq y \text{ implies } x \leq y)\}$.

conjectures the set of most plausible O -states. That is, $\mathbb{L}_{\leq} : \mathcal{O} \rightarrow \mathcal{P}(X)$ is an AGM-learner on (X, \mathcal{O}, \leq) if $\mathbb{L}_{\leq}(O) = \text{Max}_{\leq}(O)$ for all $O \in \mathcal{O}$. By the observational Limit Condition given in Definition 1, it is then guaranteed that $\mathbb{L}_{\leq}(O) \neq \emptyset$ for all $O \in \mathcal{O}$ with $O \neq \emptyset$. This means that an AGM-learner makes consistent conjectures whenever her information state is consistent.

Epistemic Scenarios. As in Subset Space Semantics, the formulas of our logic are interpreted *not* at possible worlds, but at so-called *epistemic scenarios*: pairs (x, U) of an ontic state $x \in X$ and an information state $U \in \mathcal{O}$ such that $x \in U$. Therefore, only the *truthful* observations about the actual state play a role in the evaluation of formulas. Intuitively, x represents the *actual state* of the world, while U represents the *agent's current evidence* (based on her previous observations). We denote by $ES(M) := \{(x, U) \mid x \in U \in \mathcal{O}\}$ the set of all epistemic scenarios of model M .

Dynamics: Observational Updates. Each observational event $e \in \Pi_{Ob}$ induces a dynamic “update” of the agent’s information state. This is encoded in an *update function* (also denoted by) $e : \mathcal{O} \rightarrow \mathcal{O}$, that maps any information state $U \in \mathcal{O}$ to an updated information state $e(U) \in \mathcal{O}$. The map is given by recursion:

$$\begin{aligned} !\top(U) &= U, & !o(U) &= U \cap \|o\|, \\ (e; e')(U) &= e'(e(U)), & (e \sqcup e')(U) &= e(U) \cup e'(U). \end{aligned}$$

The meaning of these clauses should be obvious: the null event $!\top$ does not change the agent’s information state; the single observation of variable o simply adds $\|o\|$ to the current evidence U (so that the agent will know the world is in $U \cap \|o\|$); the information state after a sequential composition $e; e'$ is the same as the one obtained by updating first with e then with e' ; while the information state produced by a partial observation $e \sqcup e'$ is the disjunction of the information states produced by the two events (since the agent doesn’t know which of the two happened).

It is easy to see that the update map is appropriately defined:

Lemma 1. *Let $M = (X, \mathcal{O}, \leq, \|\cdot\|)$ be a plausibility learning model and $U \in \mathcal{O}$ be an information state. Then, for all $e \in \Pi_{Ob}$ we have $e(U) \in \mathcal{O}$.*

Proof. The proof follows easily by induction on the structure of e . For the base cases $!\top$ and $!o$, we have $!\top(U) = U \in \mathcal{O}$ and $!o(U) = \|o\| \cap U \in \mathcal{O}$ by the closure of \mathcal{O} under finite intersections. In the inductive case $e; e'$, we apply the inductive hypothesis to e and U , yielding that $e(U) \in \mathcal{O}$, then we obtain that $(e; e')(U) = e'(e(U))$ (by applying again the inductive hypothesis to e' and $e(U)$). Finally, in the inductive case $e \sqcup e'$, we use the inductive hypothesis for e and U , as well as for e' and U , together with the closure of \mathcal{O} under finite unions, to conclude that $(e \sqcup e')(U) = e(U) \cup e'(U) \in \mathcal{O}$.

Definition 2 (Semantics). Given a plausibility learning model $M = (X, \mathcal{O}, \leq, \|\cdot\|)$ and an epistemic scenario (x, U) , the semantics of the language \mathcal{L} is given by a binary relation $(x, U) \models_M \varphi$ between epistemic scenario and formulas,

called the *satisfaction relation*, as well as a *truth set* (interpretation) $\llbracket \varphi \rrbracket_M^U := \{x \in U \mid (x, U) \models_M \varphi\}$, for all formulas $\varphi \in \mathcal{L}$. We typically omit the subscript, simply writing $(x, U) \models \varphi$ and $\llbracket \varphi \rrbracket^U$, whenever the model M is understood. The satisfaction relation is defined by the following recursive clauses:

$$\begin{aligned}
(x, U) \models p & \quad \text{iff } x \in \llbracket p \rrbracket \\
(x, U) \models o & \quad \text{iff } x \in \llbracket o \rrbracket \\
(x, U) \models \neg \varphi & \quad \text{iff } (x, U) \not\models \varphi \\
(x, U) \models \varphi \wedge \psi & \quad \text{iff } (x, U) \models \varphi \text{ and } (x, U) \models \psi \\
(x, U) \models L(e) & \quad \text{iff } x \in \text{Max}_{\leq} e(U) \\
(x, U) \models K\varphi & \quad \text{iff } (\forall y \in U) ((y, U) \models \varphi) \\
(x, U) \models [e]\varphi & \quad \text{iff } x \in e(U) \text{ implies } (x, e(U)) \models \varphi \\
(x, U) \models \Box\varphi & \quad \text{iff } (\forall O \in \mathcal{O}) (x \in O \subseteq U \text{ implies } (x, O) \models \varphi) \\
& \quad \text{i.e. } (\forall O \in \mathcal{O}) (x \in O \text{ implies } (x, U \cap O) \models \varphi)
\end{aligned}$$

where $p \in \text{Prop}$, $o \in \text{Prop}_{\mathcal{O}}$, and $e \in \Pi_{Ob}$.

We say that a formula φ is *valid in a plausibility learning model* M , and write $M \models \varphi$, if $(x, U) \models_M \varphi$ for all epistemic scenarios $(x, U) \in ES(M)$. We say φ is *valid*, and write $\models \varphi$, if it is valid in *all* plausibility learning models.

Precondition (Informational Content). To each observational event $e \in \Pi_{Ob}$, we can associate a formula $pre(e) \in \mathcal{L}$, called the *precondition* of event e . The definition is by recursion: $pre(!\top) = \top$, $pre(!o) = o$, $pre(e; e') = pre(e) \wedge pre(e')$, and $pre(e \sqcup e') = pre(e) \vee pre(e')$. The precondition formula $pre(e)$ captures the “condition of possibility” of the event e (i.e. e can happen in a world x iff $pre(e)$ is true at (x, U) , for any $U \in \mathcal{O}$ with $x \in U$), as well as its *informational content* (the learner’s new information after e). Both these interpretations are justified by the following result:

Lemma 2. *Let $M = (X, \mathcal{O}, \leq, \|\cdot\|)$ be a plausibility learning model and $U \in \mathcal{O}$ be an information state. Then, for all $e \in \Pi_{Ob}$ we have:*

$$\llbracket pre(e) \rrbracket^U = e(U) = \llbracket \langle e \rangle \top \rrbracket^U.$$

2.1 Expressive Power of \mathcal{L} and Its Fragments

In this brief subsection, we compare the expressive power of \mathcal{L} to those of its fragments of interest. Let $\mathcal{L}_{KL\Box}$ denote the fragment of \mathcal{L} obtained by removing only the operators $[e]\varphi$. The fragment obtained by further removing the effort modality $\Box\varphi$ is called the *static fragment* and denoted by \mathcal{L}_{KL} . Finally, we denote the epistemic fragment having only the knowledge modality K by \mathcal{L}_K .

Theorem 1 (Expressivity). *\mathcal{L} is equally expressive as $\mathcal{L}_{KL\Box}$, and they are strictly more expressive than the static fragment \mathcal{L}_{KL} with respect to plausibility learning models. Moreover, \mathcal{L}_{KL} is strictly more expressive than the epistemic fragment \mathcal{L}_K .*

Proof. \mathcal{L} is equally expressive as $\mathcal{L}_{KL\Box}$: use step-by-step the reduction axioms in Table 1 as a rewriting process and prove termination by defining a strict partial order on \mathcal{L} that satisfies similar properties as in [3, Lemma 11].⁴ For the second claim, consider the following two-state models $M_1 = (X, \mathcal{O}_1, \leq, \|\cdot\|)$ and $M_2 = (X, \mathcal{O}_2, \leq, \|\cdot\|)$ where $X = \{x, y\}$, $\leq = \{(x, x), (y, y), (x, y)\}$ and the valuation $\|p\| = \{y\}$. And, take $\mathcal{O}_1 = \{X, \emptyset\}$ (the trivial topology on X) and $\mathcal{O}_2 = \mathcal{P}(X)$ (the discrete topology on X). It is then easy to see that $M_1, (x, \{x, y\})$ and $M_2, (x, \{x, y\})$ are modally equivalent with respect to the language \mathcal{L}_{KL} . However, $M_2, (x, \{x, y\}) \models \Diamond K\neg p$ (since $\{x\}$ is an open set of M_2) whereas $M_1, (x, \{x, y\}) \not\models \Diamond K\neg p$, since the only open including x is $\{x, y\}$ and $x \notin \|p\| = \{y\}$. To prove that \mathcal{L}_{KL} is strictly more expressive than the epistemic fragment \mathcal{L}_K , consider the models $M'_1 = (X, \mathcal{O}_1, \leq_1, \|\cdot\|)$ and $M'_2 = (X, \mathcal{O}_2, \leq_2, \|\cdot\|)$, where X, \mathcal{O}_1 , and \mathcal{O}_2 are as above but $\leq_1 = \leq$ and $\leq_2 = \{(x, x), (y, y), (y, x)\}$. It is then easy to see that $M'_1, (x, \{x, y\})$ and $M'_2, (x, \{x, y\})$ are modally equivalent with respect to the language \mathcal{L}_K whereas $M'_1, (x, \{x, y\}) \not\models L(!\top)$ (since $x \notin \text{Max}_{\leq_1}(!\top(\{x, y\})) = \{y\}$) but $M'_2, (x, \{x, y\}) \models L(!\top)$ (since $x \in \text{Max}_{\leq_2}(!\top(\{x, y\})) = \{x\}$).

3 Expressing Belief and Notions of Learnability

Having presented the Dynamic Logic of AGM Learning, we now explore how various notions of belief and learnability can be expressed within this framework. We first recall the definitions of these notions given in [3].

Certain (Infallible) Knowledge and Learnability with Certainty. The notion of infallible knowledge is in our logic directly represented by the modality K , whose semantic clause mimics the following definition. The AGM learner is said to *infallibly know* a proposition $P \subseteq X$ in an information state $U \in \mathcal{O}$ if her information state U entails P , i.e., $U \subseteq P$. The possibility of learning a proposition with such certainty in a possible world $x \in X$ by a learner \mathbb{L}_{\leq} if given enough evidence (true at x) is called *learnability with certainty*.⁵ In other words, P is learnable with certainty at world x if there exists some truthful information state $O \in \mathcal{O}$ (i.e., $x \in O$) such that the learner infallibly knows P in information state O . As anticipated in [16], the notion of learnability with certainty is syntactically characterised in our language by $\Diamond Kp$, as shown in the following proposition.

Proposition 1. *Given a plausibility learning model $M = (X, \mathcal{O}, \leq, \|\cdot\|)$ and $(x, U) \in ES(M)$, $(x, U) \models \Diamond Kp$ iff $\|p\|$ is learnable with certainty at x .*

⁴ This is a standard method in Dynamic Epistemic Logic and we refer the reader to [18, Chap. 7.4] for further details.

⁵ When we quantify over learners, learnability with certainty (by *some* learners) matches the standard concept of “finite identifiability” from Formal Learning Theory.

Belief, Inductive Knowledge, and Inductive Learnability. The notion of infallible knowledge is obviously very strong: we know very few things with such certainty (maybe some logical or mathematical truths that require only hard thinking and no empirical evidence). One needs weaker notions of knowledge if one desires to model the type of knowledge we can acquire from experimental evidence that is typically partial and incomplete. This type of knowledge is taken to be fallible, yet resistant to truthful evidence gain and stronger than plain belief. In this learning theoretical context, it is captured by an evidence-based notion of *inductive knowledge* defined as *true undefeated belief*.

In an information state U , we say that the AGM learner *believes* a proposition $P \subseteq X$ if her conjecture given U entails P , that is, $\mathbb{L}_{\leq}(U) \subseteq P$. This gives us the standard interpretation of belief on plausibility models (see, e.g., [9, 10, 17]):

$$(x, U) \models B\varphi \text{ iff } \text{Max}_{\leq} e(U) \subseteq \llbracket \varphi \rrbracket^U.$$

In our formal language, belief is not a primitive notion, but can be defined as an abbreviation:

$$B\varphi := K(L(!\top) \rightarrow \varphi).$$

Indeed, it is easy to check that this defined notion satisfies the semantic clause above.

We say that, in information state U and ontic state x , the AGM learner *has undefeated belief* in a proposition $P \subseteq X$ if she believes P and will continue to believe P no matter what new true observations will be made; i.e. iff $(x, O) \models BP$ for every $O \in \mathcal{O}$ with $x \in O$. We then say, in an information state U , the AGM learner *inductively knows* P at world x if P is true and the learner has undefeated belief in P . Finally, P is *inductively learnable* by the AGM learner \mathbb{L}_{\leq} at world x if there exists some truthful information state $O \in \mathcal{O}$ (i.e., $x \in O$) such that \mathbb{L}_{\leq} inductively knows P in information state O at x .⁶ The following proposition shows that the Dynamic Logic of AGM Learning can capture these notions:

Proposition 2. *Given a plausibility learning model $M = (X, \mathcal{O}, \leq, \|\cdot\|)$ and $(x, U) \in ES(M)$,*

1. $(x, U) \models \Box Bp$ iff the learner \mathbb{L}_{\leq} has undefeated belief in $\llbracket p \rrbracket$ (at world x in information state U).
2. $(x, U) \models p \wedge \Box Bp$ iff the learner \mathbb{L}_{\leq} inductively knows $\llbracket p \rrbracket$ (at world x in information state U).
3. $(x, U) \models p \wedge \Diamond \Box Bp$ iff $\llbracket p \rrbracket$ is inductively knowable by \mathbb{L}_{\leq} (at world x in information state U).

Example: The alcohol inspector. An alcohol inspector needs to randomly check cars that pass through a security point in a perimetrical highway of

⁶ When we quantify over learners, inductive learnability (by *some* learners) matches the standard concept of “identifiability in the limit” from Formal Learning Theory, see e.g. [13].

Munich during the October fest to check the driver’s alcohol levels. The maximum alcoholic-level allowed is 30 points (which corresponds to two small beers). His alcohol-measuring tool, known as *breathalyser*, has an accuracy of ± 20 . At some point, a young woman gets the stop sign in order to get inspected. The breathalyser outputs a reading of 40 points. Given the accuracy of the tool, this first measurement can be represented by the interval $(20, 60) \subseteq \mathbb{R}$. At this point, the inspector cannot know for sure that the driver has drunk more beers than allowed. The inspector then borrows a more advanced and accurate breathalyser from one of his colleagues, with an accuracy of ± 5 . The more accurate breathalyser outputs a reading of 35 points. So the measurement of the second breathalyser can be represented by the interval $(30, 40)$. Therefore, after the reading of the second device, the inspector can know with certainty that the woman has exceeded the levels of alcohol, so she needs to wait for a couple of hours before driving again and to pay a costly fine. Moreover, let us assume that inspector obeys the legal principle of “*believing in innocence until proven guilty beyond doubt*”: so, whenever he is in doubt (because his measurements do not prove either case), he believes the driver is *not drunk*.

This situation can be represented in a plausibility learning frame⁷ $(X, \mathcal{O}, \preceq)$, where (X, \mathcal{O}) is a lattice frame with $X = [0, \infty) \subseteq \mathbb{R}$ is the set of “possible worlds” (=possible alcohol levels), while the family of partial observations \mathcal{O} is the closure under finite intersections and finite unions of the family of breathalyser measurements (=single-step total observations) $\mathcal{B} = \{[0, b) \subseteq \mathbb{R} : 0 < b \in \mathbb{Q}\} \cup \{(a, b) \subseteq \mathbb{R} : 0 < a, b \in \mathbb{Q}\}$. The sets in \mathcal{B} represent all possible readings of arbitrarily accurate breathalysers, while the sets in \mathcal{O} represent all possible information states of the inspector, based on iterated (and possibly) partial reports of such readings. Finally, the policy of believing in “innocence until proven guilty” is captured by assuming that (in the absence of any evidence) the inspector considers all non-drunken states to be a priori *more plausible* than all drunken states: i.e. $x \preceq y$ for $x > 30$ and $y \leq 30$. This policy is not enough to fully determine the plausibility relation; to make it precise, let us assume for now that the inspector has *no other strong belief on the matter*, i.e. he considers all the *drunken states to be equally plausible* (and similarly for the non-drunken states). So the relation is given by putting: $x \preceq y$ iff either $y \leq 30$ or else $30 < x, y$. It is easy to check that \preceq is indeed a total preorder.

Consider the propositions *drunk* $D = (30, \infty)$ and *not drunk* $ND = [0, 30]$ in the context of this example. We can then ask if the inspector knows with certainty that the woman is outside the permitted alcohol levels, namely if the inspector knows proposition D . After the second reading, the inspector knows with certainty that the woman has drunk more than allowed. Thus, given enough more accurate measurements, the inspector can infallibly know D (whenever D is actually the case); i.e. D is always learnable with certainty. However proposition ND is *not* always learnable with certainty: if the real level of alcohol happens to be exactly 30, then the driver is not drunk (ND) but the inspector will never

⁷ We use \preceq to denote the plausibility order in this frame, to distinguish it from the natural order on $X \subseteq \mathbb{R}$.

come to infallibly know ND . (This is simply because any interval containing 30 has non-empty intersection with D .) Still, ND is “falsifiable” with certainty (since its negation is learnable with certainty whenever true). A property that is *neither learnable with certainty nor falsifiable with certainty* is having alcohol level *barely-above-permitted* $BAP = (30, 31]$.

Inductive learnability is of course a weaker, more general form of knowledge: *both* properties *drunk* $D := (30, \infty)$ and *not-drunk* $ND := [0, 30]$ are inductively learnable by the inspector, if endowed with the above plausibility order \preceq . Indeed, if the true alcohol level is some $w \in ND = [0, 30]$, then the inspector (in the absence of any evidence), starts by believing ND (since $L(X) = \text{Max}_{\preceq} X = [0, 30]$); and no matter what further direct evidence (a, b) she gets, with $a < w < b$, she will still believe ND (since in this case $L(a, b) = \text{Max}_{\preceq}(a, b) \subseteq [0, 30]$). So in this case the inspector inductively knows ND from the start! While if $w \in D = (30, \infty)$, then after doing an accurate enough measurement, the inspector will obtain some evidence (a, b) , with $30 < a < w < b$. For any further refinement $(a', b') \subseteq (a, b)$ of this evidence, we will have $(a', b') \subseteq (a, b) \subseteq (30, \infty) = D$, hence $L(a', b') = \text{Max}_{\preceq}(a', b') = (a', b') \subseteq D$. Which means that, after reading (a, b) , the inspector achieves inductive knowledge of D : he will believe D no matter what further observations might be made.

What about the property $BAP = (30, 31]$ of having a *barely-above-permitted* alcohol level? This property is in principle also inductively learnable (by *some* learners), but *not* by the above AGM learner! To design an AGM learner who can inductively learn it, we need to change the plausibility relation, using a different refinement of the general “innocent until proven guilty” policy. The inspector still believes *all the non-drunk states to be more plausible than all the drunk ones*; but now, within the drunk-world zone, he has a similarly generous attitude: “if guilty then barely guilty”. In other words, he considers *the barely-above-permitted levels in $BAP = (30, 31]$ to be more plausible than the way-above-permitted ones in $WAV = (31, \infty)$* ; and in the rest, he is indifferent, as before. This amounts to adopting a plausibility order \ll , given by putting $x \ll y$ iff: either we have $y \leq 30$, or else we have both $30 < x$ and $y \leq 31$, or otherwise we have $31 < x, y < \infty$. It is easy to check that \ll is a linear pre-order, and moreover that properties D , ND , BAP , $NBAP = X - BAP = (0, 30] \cup (31, \infty)$, $WAV = (31, \infty)$ and $NWAV = [0, 31]$ are all inductively learnable by an inspector endowed with this plausibility order.

4 A Complete Proof System

In this section, we present a sound and complete proof system for our logic.

4.1 Axiomatization

Table 1 presents the axioms and inference rules of the Logic of AGM Learning (**L**).

Table 1. The axiom schemas for the Dynamic Logic of AGM Learning (**L**)

Basic axioms:	
(P)	All instantiations of propositional tautologies
(K _K)	$K(\varphi \rightarrow \psi) \rightarrow (K\varphi \rightarrow K\psi)$
(T _K)	$K\varphi \rightarrow \varphi$
(4 _K)	$K\varphi \rightarrow KK\varphi$
(5 _K)	$\neg K\varphi \rightarrow K\neg K\varphi$
(K _[e])	$[e](\psi \rightarrow \chi) \rightarrow ([e]\psi \rightarrow [e]\chi)$
Basic rules:	
(MP)	From $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$, infer $\vdash \psi$
(Nec _K)	From $\vdash \varphi$, infer $\vdash K\varphi$
(Nec _[e])	From $\vdash \varphi$, infer $\vdash [e]\varphi$
Learning axioms:	
(CC)	$pre(e) \rightarrow \langle K \rangle L(e)$ Consistency of Conjecture
(EC)	$K(pre(e) \leftrightarrow pre(e')) \rightarrow (L(e) \leftrightarrow L(e'))$ Extensionality of Conjecture
(SP)	$L(e) \rightarrow pre(e)$ Success Postulate
(Inc)	$(pre(e) \wedge L(e')) \rightarrow L(e; e')$ Inclusion
(RMon)	$\langle K \rangle (L(e') \wedge pre(e)) \rightarrow (L(e; e') \rightarrow (pre(e) \wedge L(e')))$ Rational Monotonicity
Reduction axioms:	
(R _p)	$[e]p \leftrightarrow (pre(e) \rightarrow p)$
(R _o)	$[e]o \leftrightarrow (pre(e) \rightarrow o)$
(R _L)	$[e]L(e') \leftrightarrow (pre(e) \rightarrow L(e; e'))$
(R _¬)	$[e]\neg\psi \leftrightarrow (pre(e) \rightarrow \neg[e]\psi)$
(R _K)	$[e]K\psi \leftrightarrow (pre(e) \rightarrow K[e]\psi)$
(R _e)	$[e][e']\psi \leftrightarrow [e; e']\psi$
(R _□)	$[e]\Box\psi \leftrightarrow \Box[e]\psi$
Effort axiom and rule:	
(□Ax)	$\Box\varphi \rightarrow [e]\varphi$, for $e \in \Pi_{Ob}$
(□Ru)	From $\vdash \psi \rightarrow [e; !o]\varphi$, infer $\vdash \psi \rightarrow [e]\Box\varphi$, where $o \notin O_\psi \cup O_e \cup O_\varphi$

Proposition 3. *The following formulas are derivable in **L** for all $\varphi \in \mathcal{L}$ and $e \in \Pi_{Ob}$:*

1. $\langle K \rangle (L(e') \wedge pre(e)) \rightarrow (L(e; e') \leftrightarrow (pre(e) \wedge L(e')))$
2. $[e](\varphi \wedge \psi) \leftrightarrow ([e]\varphi \wedge [e]\psi)$
3. $\langle e \rangle \psi \leftrightarrow (pre(e) \wedge [e]\psi)$
4. *from $\vdash \varphi \leftrightarrow \psi$, infer $\vdash [e]\varphi \leftrightarrow [e]\psi$*
5. $\langle e \rangle pre(e') \leftrightarrow pre(e; e')$
6. *from $\vdash pre(e) \leftrightarrow pre(e')$, infer $\vdash [e]\varphi \leftrightarrow [e']\varphi$*
7. $[!T]\varphi \leftrightarrow \varphi$ (we denote it R[**T**])
8. *from $\vdash \psi \rightarrow [!o]\varphi$ infer $\vdash \psi \rightarrow \Box\varphi$ (where $o \notin O_\psi \cup O_\varphi$)*

Intuitive Reading of the Axioms and Rules. The axiomatization of the Dynamic Logic of AGM Learning, roughly speaking, extends that of Dynamic

Logic for Learning Theory (DLLT) presented in [3] with axioms capturing AGM-type learning from partial observations. Group Basic Axioms and rules are quite standard: $S5$ axioms and rules for K says that the notion of knowledge with absolute certainty we study in this paper is factive and fully (both positively and negatively) introspective. $(K_{[e]})$ and $(Nec_{[e]})$ together show that dynamic modalities $[e]\varphi$ behave like normal modal operators. The reduction axioms are as in Epistemic Action Logic (EAL) [6, 8] (a.k.a., Action Model Logic [18]), where the precondition of an observational event e is captured by $pre(e)$, that is, the informational content of the event e being true. The first three learning axioms - in slightly different forms - are also part of DLLT. To recap, (CC) states that the learner conjectures consistent propositions upon having received truthful information; (EC) says that the form of the observational event (primitive, sequential, or non-deterministic) is irrelevant for learning, what is important is the informational content of the observation: observing informationally equivalent events gives rise to equivalent conjectures. Moreover, (SP) states that what the learner conjectures fits what is observed, that is, the learner conjectures propositions that support what she has observed. The last two learning axioms (Inc) and (RMon) are novel to the current system and corresponds to the AGM postulates Inclusion and Rational Monotonicity in [1], respectively. These are better understood in terms of belief. (Inc) states that the agent believes a proposition P after having observed e only if she initially believes that e entails P . (RMon) on the other hand says that the agent revises her beliefs in a monotonic way as long as the newly observed event is consistent with her previous conjecture. Finally, we have the Effort rule ($\Box Ru$) and axiom ($\Box Ax$) which together explain the dynamic behavior of the effort modality. While the former expresses that if φ is stably true then it holds after any observational event has taken place, the latter states that if φ holds after any more informative event has taken place ($[e; o]$), φ is stably true after e has taken place.

4.2 Soundness and Completeness

The soundness of the axiomatization \mathbf{L} is not entirely straightforward due to the non-standard inference rule $\Box Ru$. We present validity proofs for $\Box Ax$ and $\Box Ru$ and the completeness proof in full detail in the longer online version of this paper. We here provide sketches of the aforementioned proofs by listing the crucial lemmas.

The following lemma plays an important role in the soundness of $\Box Ru$.

Lemma 3. *Let $M = (X, \mathcal{O}, \leq, \|\cdot\|)$ and $M' = (X, \mathcal{O}, \leq, \|\cdot\|')$ be two plausibility learning models and $\varphi \in \mathcal{L}$ such that M and M' differ only in the valuation of some $o \notin O_\varphi$. Then, for all $U \in \mathcal{O}$, we have $\llbracket \varphi \rrbracket_M^U = \llbracket \varphi \rrbracket_{M'}^U$.*

Theorem 2. *The system \mathbf{L} in Table 1 is sound wrt the class of plausibility learning models.*

Canonical Model Construction. The standard notion of maximally consistent theory is *not* very useful for our logic, since such theories do not ‘internalize’ the Effort rule $\Box\text{Ru}$. To do this, we need instead to consider ‘witnessed’ (*maximally consistent*) theories, in which every occurrence of a $\Diamond\varphi$ in any ‘existential context’ is ‘witnessed’ by some $\langle!o\rangle\varphi$ (with o observational variable). The appropriate notion of ‘existential contexts’ is represented by *possibility forms*, as in e.g., [2, 3, 7], given in Definition 3.

Definition 3 (‘Pseudo-modalities’: Necessity and Possibility Forms).

The set of necessity-form expressions of our language is given by $NF_{\mathcal{L}} := (\{\varphi \rightarrow \mid \varphi \in \mathcal{L}\} \cup \{K\} \cup \{e : e \in \Pi_{Ob}\})^*$. For any finite string $s \in NF_{\mathcal{L}}$, we define pseudo-modalities $[s]$ (called necessity form) and $\langle s \rangle$ (called possibility form) that generalize our dynamic modalities $[e]$ and $\langle e \rangle$. These pseudo-modalities are functions mapping any formula $\varphi \in \mathcal{L}$ to another formula $[s]\varphi \in \mathcal{L}$, and respectively $\langle s \rangle\varphi \in \mathcal{L}$. Necessity forms are defined recursively, by putting: $[e]\varphi := \varphi$ (where e is the empty string), $[s, \varphi \rightarrow]\varphi := [s](\varphi \rightarrow \varphi)$, $[s, K]\varphi := [s]K\varphi$, $[s, e]\varphi := [s][e]\varphi$. As for possibility forms, we put $\langle s \rangle\varphi := \neg[s]\neg\varphi$.

Lemma 4. For every necessity form $[s]$, there exist an observational event $e \in \Pi_{Ob}$ and a formula $\psi \in \mathcal{L}$, with $O_\psi \cup O_e \subseteq O_s$, such that for all $\varphi \in \mathcal{L}$, we have

$$\vdash [s]\varphi \text{ iff } \vdash \psi \rightarrow [e]\varphi.$$

Lemma 5. The following rule is admissible in \mathbf{L} :

$$\text{if } \vdash [s][!o]\varphi \text{ then } \vdash [s]\Box\varphi, \text{ where } o \notin O_s \cup O_\varphi.$$

Proof. Suppose $\vdash [s][!o]\varphi$ where $o \notin O_s \cup O_\varphi$. Then, by Lemma 4, there exist $e \in \Pi_{Ob}$ and $\psi \in \mathcal{L}$ with $O_\psi \cup O_e \subseteq O_s$ such that $\vdash \psi \rightarrow [e][!o]\varphi$. Thus we get $\vdash \psi \rightarrow [e; !o]\varphi$ by an instance of R_e . Therefore, by the Effort rule ($\Box\text{Ru}$) we have $\vdash \psi \rightarrow [e]\Box\varphi$. Then, again by Lemma 4, we obtain $\vdash [s]\Box\varphi$.

Definition 4. For every countable set O , let \mathcal{L}^O be the language of the logic \mathbf{L}^O based only on the observational variables in O (i.e., having as set of observational variables $\text{Prop}_{\mathcal{L}^O} := O$). Let $NF_{\mathcal{L}^O}$ denote the set of necessity-form expressions of \mathcal{L}^O (i.e., necessity forms involving only observational variables in O). An O -theory is a consistent set of formulas in \mathcal{L}^O . Here, ‘consistent’ means consistent with respect to the axiomatization \mathbf{L} formulated for \mathcal{L}^O . A maximal O -theory is an O -theory Γ that is maximal with respect to \subseteq among all O -theories; in other words, Γ cannot be extended to another O -theory. An O -witnessed theory is an O -theory Γ such that, for every $s \in NF_{\mathcal{L}^O}$ and $\varphi \in \mathcal{L}^O$, if $\langle s \rangle\Diamond\varphi$ is consistent with Γ then there is $o \in O$ such that $\langle s \rangle\langle!o\rangle\varphi$ is consistent with Γ . A maximal O -witnessed theory Γ is an O -witnessed theory that is not a proper subset of any O -witnessed theory.

The proofs of the following lemmas are exactly as in the corresponding proofs in [3], taking into account that (maximally) O -(witnessed) theories here are defined using primitive observational events $\langle!o\rangle$ (rather than observational variables o).

Lemma 6. For every maximal \mathcal{O} -witnessed theory Γ , and any $\varphi, \psi \in \mathcal{L}^{\mathcal{O}}$,

1. either $\varphi \in \Gamma$ or $\neg\varphi \in \Gamma$,
2. $\varphi \wedge \psi \in \Gamma$ iff $\varphi \in \Gamma$ and $\psi \in \Gamma$,
3. $\varphi \in \Gamma$ and $\varphi \rightarrow \psi \in \Gamma$ implies $\psi \in \Gamma$.

Lemma 7 (Lindenbaum's Lemma). Every \mathcal{O} -witnessed theory Γ can be extended to a maximal \mathcal{O} -witnessed theory T_{Γ} .

Lemma 8 (Extension Lemma). Let \mathcal{O} be a set of observational variables and \mathcal{O}' be a countable set of fresh observational variables, i.e., $\mathcal{O} \cap \mathcal{O}' = \emptyset$. Let $\tilde{\mathcal{O}} = \mathcal{O} \cup \mathcal{O}'$. Then, every \mathcal{O} -theory Γ can be extended to an $\tilde{\mathcal{O}}$ -witnessed theory $\tilde{\Gamma} \supseteq \Gamma$, and hence to a maximal $\tilde{\mathcal{O}}$ -witnessed theory $T_{\tilde{\Gamma}} \supseteq \tilde{\Gamma}$.

Canonical Model for T_0 . For any consistent set of formulas Φ , consider a maximally consistent \mathcal{O} -witnessed extension $T_0 \supseteq \Phi$. As our canonical set of worlds, we take the set $X^c := \{T : T \text{ maximally consistent } \mathcal{O}\text{-witnessed theory with } T \sim_K T_0\}$, where we put

$$T \sim_K T' \text{ iff } \forall \varphi \in \mathcal{L}^{\mathcal{O}} (K\varphi \in T \text{ implies } \varphi \in T').$$

It is easy to see (given the $S5$ axioms for K) that \sim_K is an equivalence relation. For any formula φ , we use the notation $\widehat{\varphi} := \{T \in X^c : \varphi \in T\}$. As the canonical set of information states, we take $\mathcal{O}^c := \{\widehat{pre}(e) : e \in \Pi_{\mathcal{O}b}^{\mathcal{O}}\}$. Toward defining the canonical plausibility relation \leq^c , let

$$S_e = \bigcup \{\widehat{L}(e') : \widehat{pre}(e) \subseteq \widehat{pre}(e') \text{ and } e' \in \Pi_{\mathcal{O}b}^{\mathcal{O}}\},$$

and $\$ = \{S_e : e \in \Pi_{\mathcal{O}b}^{\mathcal{O}}\} \cup \{X^c\}$. The canonical plausibility order \leq^c on X^c is given by, for any $T, T' \in X^c$:

$$T \leq^c T' \text{ iff } \forall S \in \$ (T \in S \text{ implies } T' \in S).$$

The definition of \leq^c is inspired by the construction of the so-called *order models* from *sphere* and *selection* models presented in [14]. Roughly speaking, while $\widehat{L}(e')$ plays the role of a selection function that picks out a set of maximally consistent \mathcal{O} -witnessed theories given e' (see, e.g., [11, 14] for selection models), the collection of sets $\$$ forms a sphere system (see, e.g., [15] for sphere models).

The canonical valuation $\|\cdot\|_c$ is given as $\|p\|_c = \widehat{p}$ and $\|o\|_c = \widehat{o}$. The tuple $M^c = (X^c, \mathcal{O}^c, \leq^c, \|\cdot\|_c)$ is called the *canonical model*.

Theorem 3. $M^c = (X^c, \mathcal{O}^c, \leq^c, \|\cdot\|_c)$ is a plausibility learning model.

The following lemmas will be useful for proving the Truth Lemma.

Lemma 9. For all $e \in \Pi_{\mathcal{O}b}^{\mathcal{O}}$, $Max_{\leq^c}(\widehat{pre}(e)) = \widehat{L}(e)$.

Lemma 10. *For every maximal O-witnessed theory T , the set $\{\theta : K\theta \in T\}$ is O-witnessed.*

Lemma 11. *Let $T \in X^c$. Then, $K\varphi \in T$ iff $\varphi \in T'$ for all $T' \in X^c$.*

Lemma 12. *Let $T \in X^c$. Then, $\Box\varphi \in T$ iff $[e]\varphi \in T$ for all $e \in \Pi_{Ob}^O$.*

Proof. The direction from left-to-right follows by the axiom (\Box Ax). For the direction from right-to-left, suppose, toward a contradiction, that for all $e \in \Pi_{Ob}^O$, $[e]\varphi \in T$ and $\Box\varphi \notin T$. Then, since T is a maximally consistent theory, $\Diamond\neg\varphi \in T$. Since T is an O-witnessed theory, there is $o \in O$ such that $\langle!o\rangle\neg\varphi$ is consistent with T . Since T is also maximally consistent, we obtain that $\langle!o\rangle\neg\varphi \in T$, i.e., that $\neg[!o]\varphi \in T$, contradicting our initial assumption.

Lemma 13 (Truth Lemma). *Let $M^c = (X^c, \mathcal{O}^c, \leq^c, \|\cdot\|_c)$ be the canonical model for some T_0 . For all formulas $\varphi \in \mathcal{L}^O$, all $T \in X^c$ and all $e \in \Pi_{Ob}^O$, we have:*

$$\langle e \rangle \varphi \in T \text{ iff } (T, \widehat{pre}(e)) \models_{M^c} \varphi.$$

Proof. The proof is by induction on the structure of φ and uses the following *induction hypothesis (IH)*: for all ψ subformula of φ , and $e \in \Pi_{Ob}^O$, $\langle e \rangle \psi \in T$ iff $(T, \widehat{pre}(e)) \models_{M^c} \psi$. The base case for propositional and observational variables, as well as Boolean formulas are straightforward. We only verify the remaining inductive cases.

Observe that at this point of the proof we have that: $\forall e, e' \in \Pi_{Ob}^O$, $\langle e \rangle \widehat{pre}(e') = \llbracket pre(e') \rrbracket_{M^c}^{\widehat{pre}(e)}$ since $pre(e)$ is a Boolean formula.

- Case $\varphi := L(e')$. We have the following sequence of equivalencies: $\langle e \rangle L(e') \in T$ iff $(pre(e) \wedge [e]L(e')) \in T$ (by Proposition 3.3) iff $(pre(e) \wedge L(e; e')) \in T$ (by R_L and CPL) iff $pre(e) \in T \wedge L(e; e') \in T$ iff $T \in \widehat{pre}(e) \wedge T \in \widehat{L}(e; e')$ (by Def. of $\widehat{\varphi}$) iff $T \in \widehat{pre}(e) \wedge T \in Max_{\leq^c}(pre(e; e'))$ (by Lemma 9) iff $T \in \widehat{pre}(e) \wedge T \in Max_{\leq^c}(\langle e \rangle \widehat{pre}(e'))$ (by Proposition 3.5) iff $T \in \widehat{pre}(e) \wedge T \in Max_{\leq^c}(\llbracket pre(e') \rrbracket_{M^c}^{\widehat{pre}(e)})$ (by the above observation) iff $(T, \widehat{pre}(e)) \models_{M^c} L(e')$ (by the semantics).
- Case $\varphi := K\psi$. We have the sequence of equivalencies: $\langle e \rangle K\psi \in T$ iff $(pre(e) \wedge K[e]\psi) \in T$ (by Proposition 3.3 and R_K) iff $pre(e) \in T \wedge K[e]\psi \in T$ iff $pre(e) \in T \wedge (\forall T' \sim_K T)([e]\psi \in T')$ (by Lemma 11) iff $pre(e) \in T \wedge (\forall T' \sim_K T \text{ s.t. } pre(e) \in T')(\langle e \rangle \psi \in T')$ (by Proposition 3.3) iff $T \in \widehat{pre}(e) \wedge (\forall T' \in \widehat{pre}(e))(T', \widehat{pre}(e)) \models_{M^c} \psi$ (by I.H.) iff $(T, \widehat{pre}(e)) \models_{M^c} K\psi$ (by the semantics).
- Case $\varphi := \langle e' \rangle \psi$. We have the sequence of equivalencies: $\langle e \rangle \langle e' \rangle \psi \in T$ iff $\langle e; e' \rangle \psi \in T$ (by R_e) iff $pre(e; e') \wedge \langle e; e' \rangle \psi \in T$ (by Proposition 3.3) iff $pre(e; e') \in T \wedge \langle e; e' \rangle \psi \in T$ iff $T \in \widehat{pre}(e) \cap \widehat{pre}(e') \wedge T \in \langle e; e' \rangle \psi$ iff $(T, \widehat{pre}(e) \cap \widehat{pre}(e')) \models_{M^c} \psi$ (by I.H.) iff $(T, \widehat{pre}(e)) \models_{M^c} \langle e' \rangle \psi$ (by the semantics).

- Case $\varphi := \Box\psi$. First observe the following: $\langle e \rangle \Box\psi \in T$ iff $e \wedge [e] \Box\psi \in T$ (by Proposition 3.3) iff $\widehat{pre}(e) \in T$ and $\Box[e]\psi \in T$ (by R_{\Box}) iff $\widehat{pre}(e) \in T$ and $[e'] [e]\psi \in T, \forall e' \in \Pi_{Ob}^O$ (by Lemma 12) iff $\widehat{pre}(e) \in T$ and $[e'; e]\psi \in T, \forall e' \in \Pi_{Ob}^O$ (by R_e) iff $\widehat{pre}(e) \in T$ and $[e; e']\psi \in T, \forall e' \in \Pi_{Ob}^O$ (by Proposition 3.6).

(\Rightarrow) Suppose $\langle e \rangle \Box\psi \in T$. Now let $U \in \mathcal{O}^c$ such that $T \in U$. By the definition of \mathcal{O}^c we know that $U = \widehat{pre}(e'')$ for some $e'' \in \Pi_{Ob}^O$. Since $T \in \widehat{pre}(e'')$ and by the observation above we obtain, $\widehat{pre}(e; e'') \in T$ and $[e; e'']\psi \in T$. Thus, $T \in \widehat{pre}(e; e'')$ and $T \in [e; e'']\psi$. By Proposition 3.3 we have $T \in \widehat{pre}(e; e'')$ and $T \in \langle e; e'' \rangle \psi$. Since $\widehat{pre}(e; e'') = \widehat{pre}(e) \cap \widehat{pre}(e'')$ and $\langle e; e'' \rangle \psi \in T$, by I.H. we get $(T, \widehat{pre}(e) \cap \widehat{pre}(e'')) \models_{M^c} \psi$. Since $\widehat{pre}(e'') = U$ was taken arbitrarily in \mathcal{O}^c , by the semantics, we obtain $(T, \widehat{pre}(e)) \models_{M^c} \Box\psi$.

(\Leftarrow) Suppose $(T, \widehat{pre}(e)) \models_{M^c} \Box\psi$. By the semantics of \Box and the definition of \mathcal{O}^c , we have that for all $e' \in \Pi_{Ob}^O$, if $T \in \widehat{pre}(e) \cap \widehat{pre}(e')$ then $(T, \widehat{pre}(e) \cap \widehat{pre}(e')) \models_{M^c} \psi$. Let $e'' \in \Pi_{Ob}^O$ such that $T \in \widehat{pre}(e'')$, therefore $T \in \widehat{pre}(e) \cap \widehat{pre}(e'')$. Since $\widehat{pre}(e) \cap \widehat{pre}(e'') \in \mathcal{O}^c$, we obtain, by the assumption, that $(T, \widehat{pre}(e) \cap \widehat{pre}(e'')) \models_{M^c} \psi$. Thus, by I.H., we have $\widehat{pre}(e; e'') \in T$ and $\langle e; e'' \rangle \psi \in T$. By Propositions 3.3 and 3.6, $[e''; e]\psi \in T$. By (R_e) we have $[e''] [e]\psi \in T$. Since e'' was taken arbitrarily, by Lemma 12, we have $\Box[e]\psi \in T$. Then, by (R_{\Box}) , we obtain $[e] \Box\psi \in T$. Since $\widehat{pre}(e) \in T$ and $[e] \Box\psi \in T$, we have $\langle e \rangle \Box\psi \in T$ by Proposition 3.3.

Theorem 4. **L** is complete with respect to the class of plausibility learning models.

Proof. Let φ be an **L**-consistent formula, i.e., it is an O_{φ} -theory. Then, by Lemma 8, it can be extended to some maximal **O**-witnessed theory T . Then, we have $\langle !\top \rangle \varphi \in T$ i.e., $T \in \widehat{\langle !\top \rangle \varphi}$ (by Proposition 3.7). Then, by Truth Lemma (Lemma 13), we obtain that $(T, \widehat{pre}(\langle !\top \rangle)) \models_{M^c} \varphi$, where $M^c = (X^c, \mathcal{O}^c, \leq^c, \|\cdot\|_c)$ is the canonical model for T . This proves completeness.

5 Conclusions and Open Questions

In this paper, we enriched the dynamic logic for learning theory (DLLT) from [3] with additional structure in order to model learners whose conjectures satisfy standard rationality constraints (namely, the AGM postulates for belief revision). The standard model for such learners is provided by “AGM conditioning”: learners are endowed with a total preorder, describing their prior plausibility relation, and at each step they believe the set of most plausible states compatible with all the previous observations. To axiomatize the DLLT logic of AGM conditioning, we were lead to assume that the learner has access to a wider range of potential

information than in [3]: not only sequences of full observations, but also *partial* observations (finite unions of observations). Semantically, this required a technical shift from intersection spaces to lattice spaces, while on the syntactic side we needed to extend our dynamic modalities from simple observations to more complex PDL-like “observational events”. This leads to a rich evidential setting, with a more interesting logic and an elegant axiomatization.

Note that our move to partial observations and observational events still requires *less* information than the classical axiomatizations of AGM conditioning in the literature (which assumed full Boolean closure of the set of “conditions”, i.e. the observable sets formed a Boolean algebra). Still, while this move to partial observations seems general enough, as well as natural and desirable in itself, it does require a much wider access to information than the setting in [3]. So it is fair to ask the question: is there a way to axiomatize AGM learners *without* requiring them to access partial information? In other words, is AGM conditioning over intersection spaces axiomatizable in a simple, elegant way (similar to our axiomatization)? This problem is still open, though we conjecture that the answer is *no*. If we are right, this would be an argument for a deeper philosophical point: it may be that AGM postulates are best suited to “rich” evidential settings, in which both total and partial observations are available.

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: partial meet contraction and revision functions. *J. Symb. Log.* **50**, 510–530 (1985)
2. Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., de Lima, T.: ‘Knowable’ as ‘known after an announcement’. *Rev. Symb. Log.* **1**, 305–334 (2008)
3. Baltag, A., Gierasimczuk, N., Özgün, A., Vargas Sandoval, A.L., Smets, S.: A dynamic logic for learning theory. In: Madeira, A., Benevides, M. (eds.) DALI 2017. LNCS, vol. 10669, pp. 35–54. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73579-5_3
4. Baltag, A., Gierasimczuk, N., Smets, S.: Belief revision as a truth-tracking process. In: Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge, pp. 187–190. ACM (2011)
5. Baltag, A., Gierasimczuk, N., Smets, S.: On the solvability of inductive problems: a study in epistemic topology. In: Ramanujam, R. (ed.) Proceedings of the 15th Conference TARK, also Available as a Technical Report in ILLC Prepublication Series PP-2015-13 (2015)
6. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In: Proceedings of the 7th Conference TARK, pp. 43–56. Morgan Kaufmann Publishers Inc. (1998)
7. Baltag, A., Özgün, A., Vargas Sandoval, A.L.: Topo-logic as a dynamic-epistemic logic. In: Baltag, A., Seligman, J., Yamada, T. (eds.) LORI 2017. LNCS, vol. 10455, pp. 330–346. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-55665-8_23
8. Baltag, A., Renne, B.: Dynamic epistemic logic. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, Winter 2016 edn. (2016)

9. Baltag, A., Smets, S.: A qualitative theory of dynamic interactive belief revision. *Texts Log. Games* **3**, 9–58 (2008)
10. Board, O.: Dynamic interactive epistemology. *Games Econ. Behav.* **49**, 49–80 (2004)
11. Chellas, B.F.: Basic conditional logic. *J. Philos. Log.* **4**, 133–153 (1975)
12. Dabrowski, A., Moss, L.S., Parikh, R.: Topological reasoning and the logic of knowledge. *Ann. Pure Appl. Log.* **78**, 73–110 (1996)
13. Gold, E.M.: Language identification in the limit. *Inf. Control* **10**, 447–474 (1967)
14. Grahne, G.: Updates and counterfactuals. *J. Log. Comput.* **8**, 87–117 (1998)
15. Lewis, D.K.: *Counterfactuals*. Blackwell, Oxford (1973)
16. Moss, L.S., Parikh, R.: Topological reasoning and the logic of knowledge. In: *Proceedings of the 4th TARK*, pp. 95–105. Morgan Kaufmann (1992)
17. van Benthem, J.: Dynamic logic for belief revision. *J. Appl. Non-Class. Log.* **14**, 2004 (2004)
18. van Ditmarsch, H., van der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*, 1st edn. Springer, Heidelberg (2007). <https://doi.org/10.1007/978-1-4020-5839-4>