



A Dynamic Logic for Learning Theory

Baltag, Alexandru; Gierasimczuk, Nina; Özgün, Aybüke; Vargas Sandoval, Ana Lucia; Smets, Sonja

Published in:
Dynamic Logic. New Trends and Applications.

Link to article, DOI:
[10.1007/978-3-319-73579-5_3](https://doi.org/10.1007/978-3-319-73579-5_3)

Publication date:
2018

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Baltag, A., Gierasimczuk, N., Özgün, A., Vargas Sandoval, A. L., & Smets, S. (2018). A Dynamic Logic for Learning Theory. In A. Madeira, & M. Benevides (Eds.), *Dynamic Logic. New Trends and Applications. : First International Workshop, DALI 2017, Brasilia, Brazil, September 23-24, 2017, Proceedings* (Vol. 10669, pp. 35-54). Springer. Lecture Notes in Computer Science https://doi.org/10.1007/978-3-319-73579-5_3

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A Dynamic Logic for Learning Theory

Alexandru Baltag¹, Nina Gierasimczuk², Aybüke Özgün^{1,3}, Ana Lucia Vargas Sandoval¹, and Sonja Smets¹

¹ ILLC, University of Amsterdam, Amsterdam, The Netherlands

² DTU Compute, Technical University of Denmark, Copenhagen, Denmark

³ LORIA, CNRS-Université de Lorraine, Nancy, France

Abstract. Building on previous work [4, 5] that bridged Formal Learning Theory and Dynamic Epistemic Logic in a topological setting, we introduce a *Dynamic Logic for Learning Theory* (DLLT), extending Subset Space Logics [17, 9] with dynamic *observation modalities* $[o]\varphi$, as well as with a *learning operator* $L(\vec{o})$, which encodes the learner’s *conjecture* after observing a finite sequence of data \vec{o} . We completely axiomatise DLLT, study its expressivity and use it to characterise various notions of knowledge, belief, and learning.

Keywords: learning theory, dynamic epistemic logic, modal logic, subset space semantics, inductive knowledge, epistemology

1 Introduction

The process of learning consists of incorporating new information into one’s prior information state. Dynamic epistemic logic (DEL) studies such one-step information changes from a logical perspective [6, 19, 23]. But the general concept of *learning* encompasses not only one-step revisions, but also their *long-term horizon*. In the long run, learning should lead to *knowledge*—an epistemic state of a particular value. Examples include language learning (inferring the underlying grammar from examples of correct sentences), and scientific inquiry (inferring a theory of a phenomenon on the basis of observations). Our goal in this paper is to provide a simple logic for reasoning about this process of *inductive learning* from successful observations. Understanding inductive inference is of course an infamously difficult open problem, and there are many different approaches in the literature.⁴ However, in this paper we *do not try to solve* the problem of induction, but only to *reason about* a (rational) inductive learner. For this, we adopt the more flexible and open-ended approach of Formal Learning Theory (FLT). While most other approaches adopt a normative stance, aimed at prescribing ‘the’ correct algorithm for forming and changing rational beliefs from observations (e.g., Bayesian conditioning), or at least at prescribing some general rational constraints that any such algorithm should obey (e.g., the AGM postulates for belief revision), FLT gives the

⁴ From probabilistic and statistical formalisms based on Bayesian reasoning, Popper-style measures of corroboration, through default and non-monotonic logics, Carnap-style ‘inductive logic’, to AGM-style rational belief revision and theory change.

learner a high degree of freedom, allowing the choice of *any learner* that produces conjectures *based on the data* (no matter how ‘crazy’ or unjustified are these conjectures, or how erratic is the process of belief change). In FLT the only criterion of success is... success: tracking the truth in the limit. In other words, the only thing that matters is whether or not the iterated belief revision process will eventually stabilise on a conjecture which matches the truth (about some given issue). Of course, we are not interested in cases of convergence to the truth ‘by accident’, but in determining whether or not a given learner is *guaranteed* to eventually track the truth; hence, the focus on ‘The Logic of *Reliable Inquiry*’.⁵

We propose a formalism that combines ideas from: Subset Space Logics, as introduced by Moss and Parikh [17], investigated further by Dabrowski et al. [9] and already merged with the DEL tradition in prior work [25, 3, 20, 22, 8, 7]; the topological approach to FLT in [16, 5]; and the general agenda of bridging DEL and FLT in [13]. Semantically, we take *intersection spaces* (a type of subset spaces that are closed under finite non-empty intersections), with points interpreted as possible worlds and neighbourhoods interpreted as *observations* (or *information states*) (see, e.g., [18] for a survey on subset space logics). We enhance these structures with a *learner* \mathbb{L} , mapping every information state to a *conjecture*, representing the learner’s strongest belief in this state. As in Subset Space logics, our language features an S5-type ‘*knowledge-with-certainty*’ modality, capturing the learner’s *hard information*, as well as the so-called ‘*effort*’ modality, which we interpret as ‘stable truth’ (i.e., truth immune to further observations). We add to this *observation modalities* $[o]\varphi$, analogous to the dynamic modalities in Public Announcement Logic (PAL), as well as a *learning operator* $L(\vec{\sigma})$, which encodes the learner’s conjecture after observing a finite sequence of pieces of evidence $\vec{\sigma}$. This can be used to give a natural definition of *belief*: a learner believes P iff she knows that P is entailed by her current conjecture.

We present a sound and complete axiomatisation of DLLT with respect to our learning models. The completeness uses a neighbourhood version of the standard canonical model construction. We use this logic to characterise various learnability notions. In particular, we are able to model inductive learning as *coming to stably believe a true fact after observing an incoming sequence of true data*. The possibility of such learning corresponds to a key concept in FLT, namely *identifiability in the limit* first introduced and studied by Gold in [15]. Finally, we discuss the expressivity of DLLT, showing that the dynamic observation modalities are in principle eliminable via reduction laws.

1.1 Effort modality and knowledge

In [24], Vickers reconstructed general topology as a logic of observation, in which the points of the space represent possible states of the world, while basic open neighborhoods of a point are interpreted as *information states* produced by accumulating finitely many observations. Moss and Parikh [17] gave an account of learning in terms of *observational effort*. Making the epistemic effort to obtain more information about a possible world has a natural topological interpretation—it can be seen as *shrinking* the open

⁵ ‘The Logic of Reliable Inquiry’ is the title of a classic text in FLT-based epistemology [16].

neighborhood (representing the current information state), thus providing a more accurate approximation of the actual state of the world [17, 9, 11, 12, 18]. A similar line was proposed in Formal Epistemology [16, 5], where it was combined with more sophisticated notions of learning borrowed from FLT. The following example relates the effort modality with knowledge.

Example 1 ([18]). Let us consider some *measurement*, say of a vehicle’s velocity. Suppose a policeman uses radar to determine whether a car is speeding in a 50-mile speed-limit zone. The property *speeding* can be identified with the interval $(50, \infty)$. Suppose the radar shows 51 mph, but the radar’s accuracy is ± 2 mph. The intuitive meaning of a speed measurement of 51 ± 2 is that the car’s true speed v is in the *open interval* $(49, 53)$. According to [18], “anything which we *know* about v must hold not only of v itself, but also of any v' in the same interval” [18, p. 300]. Since the interval $(49, 53)$ is not fully included in the ‘speeding’ interval $(50, \infty)$, the policeman does *not know* that the car is speeding. But suppose that he does another measurement, using a more accurate radar with an accuracy of ± 1 mph, which shows 51.5 mph. Then he will *come to know* that the car is speeding: the open interval $(50.5, 52.5)$ is included in $(50, \infty)$.




Fig. 1. Example 1; $P :=$ “the car is speeding”, $Q :=$ “the reading of the radar is 51 km/h”

Infallible Knowledge versus Inductive Knowledge Let us now extend this picture with learning as understood in FLT. We start by setting the stage—briefly introducing *learning frame*, the underlying structure of learning.⁶ Using them we will be able to explain and model various epistemic notions.

First, consider a pair (X, \mathcal{O}) , where X is a non-empty set of *possible worlds*; $\mathcal{O} \subseteq \mathcal{P}(X)$ is a non-empty set of *information states* (or ‘observables’, or ‘evidence’). We take \mathcal{O} to be closed on intersections, i.e., for any $O_1, O_2 \in \mathcal{O}$, we have $O_1 \cap O_2 \in \mathcal{O}$, the resulting (X, \mathcal{O}) is called an intersection space. A *learning frame* is a triplet $(X, \mathcal{O}, \mathbb{L})$, where $\mathbb{L} : \mathcal{O} \rightarrow \mathcal{P}(X)$ is a *learner*, i.e., a map associating to every finite sequence of observations $O \in \mathcal{O}$ some ‘conjecture’ $\mathbb{L}(O) \subseteq X$.

Let us now reconstruct Example 1 as a learning frame. We take $X = (0, \infty)$ as the set of possible worlds (representing possible velocities of the car, where we assume the car is known to be *moving*); $\mathcal{O} = \{(a, b) \in Q \times Q : 0 < a < b < \infty\}$ is the set of all open

⁶ We will return to it, with complete definitions, later in the paper. Our DLLT is interpreted over such frames.

intervals with positive rational endpoints (representing possible measurement results by arbitrarily accurate radars). The pair (X, \mathcal{O}) is an intersection frame, and the topology generated by \mathcal{O} is the standard topology on real numbers (restricted to X).

Certain (Infallible) Knowledge In an information state $U \in \mathcal{O}$, the learner is said to *infallibly know a proposition* $P \subseteq X$ *conditional on observation* O if her conditional information state entails P , i.e. if $U \cap O \subseteq P$. The learner (*unconditionally*) *knows* P if $U \subseteq P$. The possibility of achieving certain knowledge about a proposition $P \subseteq X$ in a possible world $x \in X$ by a learner \mathbb{L} if given enough evidence (true at x) is called *learnability with certainty*. In other worlds P is learnable with certainty if there exists some observable property $O \in \mathcal{O}$ (with $x \in O$) such that the learner infallibly knows P in information state O . Learnability can be used to define verifiability and falsifiability: a proposition $P \subseteq X$ is *verifiable (resp. falsifiable) with certainty (by \mathbb{L})* if it is learnable with certainty by \mathbb{L} whenever it is true (resp. false); i.e. if P is learnable with certainty at all worlds $x \in P$ (resp. $x \notin P$). Finally, a proposition $P \subseteq X$ is *decidable with certainty (by \mathbb{L})* if it is both verifiable and falsifiable with certainty (by \mathbb{L}).

In the context of Example 1, let us consider the certain knowledge of the policeman. In the information state $U = (49, 53)$, the learner/policeman does not know the proposition $P = (50, \infty)$, so he cannot be certain that the car is speeding. However, the speeding property P is verifiable with certainty: whenever P is actually true, he could perform a more accurate speed measurement, by which he can get to an information state in which P is infallibly known. In our example, the policeman refined his measurement getting to the information state $O = (50.5, 52.5)$, thus coming to know P . In contrast, the property $X - P = (0, 50]$ ('not speeding') is not verifiable with certainty: if by some kind of miraculous coincidence, the speed of the car is exactly 50 mph, then the car is not speeding, but the policeman will never know that for certain (since every speed measurement, of any degree of accuracy, will be consistent both with P and with $X - P$). Nevertheless, $X - P$ is always falsifiable with certainty: if false (i.e. if the speed is in P , so that car is speeding), then as we saw the policeman will come to infallibly know that (by some more accurate measurement).

Inductive (Defeasible) Knowledge

Before we proceed to Inductive Knowledge let us consider epistemic states weaker than certainty, *belief*. In an information state $U \in \mathcal{O}$, the learner \mathbb{L} is said to:

- *un-conditionally believe* $P \subseteq X$ if $\mathbb{L}(U) \subseteq P$.⁷
- *believe a proposition* $P \subseteq X$ *conditional on observation* O if $\mathbb{L}(U \cap O) \subseteq P$;
- *have undefeated belief in a proposition* $P \subseteq X$ *at world* x if she believes P in every information state $O \in \mathcal{O}$ that is true at x (i.e., $x \in O$) and is at least as strong as U (i.e., $O \subseteq U$). This means that, once she reaches information state U , no further evidence can defeat the learner's belief in P .

One of the central problems in epistemology is to define a realistic notion of knowledge that fits the needs of empirical sciences. It should allow fallibility, while requiring a higher standards of evidence and robustness than simple belief. One of the main contenders is the so-called Defeasibility Theory of Knowledge, which defines *defeasible*

⁷ In the tautological information state X , the learner believes P iff $\mathbb{L}(X) \subseteq P$.

(fallible) knowledge as true undefeated belief. In the learning-theoretic context, this gives us an evidence-based notion of ‘inductive knowledge’: in an information state U , P is *inductively known* at world x if it is true at x (i.e., $x \in P$) and it is undefeated belief (in the sense defined above). This is the kind of knowledge that can be gained by empirical (incomplete) induction, based on experimental evidence.

As in the case of learnability with certainty, achieving inductive knowledge is defined as learnability. A proposition $P \subseteq X$ is *inductively learnable* (or ‘learnable in the limit’) by the learner \mathbb{L} at world x if \mathbb{L} will come to inductively know P if given enough evidence (true at x); i.e. if there exists some observable property $O \in \mathcal{O}$ of world x (i.e., with $x \in O$) such that \mathbb{L} inductively knows P in information state O . Inductive verifiability and falsifiability are defined in terms of learnability. A proposition $P \subseteq X$ is *inductively verifiable* (resp. *falsifiable*) by the learner \mathbb{L} , if it is inductively learnable whenever it is true (resp. false); i.e., if P is inductively learnable at all worlds $x \in P$ (resp. $x \notin P$). A proposition $P \subseteq X$ is *inductively decidable* by \mathbb{L} if it is both inductively verifiable and inductively falsifiable by \mathbb{L} .

In the context of Example 1, let us now turn to inductive knowledge of the policeman. Both speeding (P) and non-speeding ($X - P$) are inductively decidable (and thus both inductively verifiable and inductively falsifiable): for instance, they are inductively decidable by the learner \mathbb{L} , defined by putting $\mathbb{L}(O) := O \cap P$ for every open interval $O = (a, b) \in \mathcal{O}$ s.t. $O \cap P \neq \emptyset$, and putting $\mathbb{L}(O) := O \subseteq X - P$ otherwise.

Intuitively, this learner is the ‘suspicious cop’, who believes the car to be *speeding* whenever the available evidence cannot settle the issue, and keeps this conjecture until it is disproven by some more accurate measurement. Regardless of the car’s speed, this policeman will be right ‘in the limit’: after doing enough accurate measurements, he will *eventually settle on the correct belief* (about speeding or not); though of course (in case the car’s speed is exactly 50 mph) he may still never be certain. Obviously, the dual learner (the ‘judge’, who assumes innocence until proven guilty) will also inductively decide the speeding issue. An example of property which is *inductively decidable but neither verifiable with certainty nor falsifiable with certainty* is the proposition $S = [50, 51)$. It is not verifiable with certainty, since if the car’s speed is exactly 50 mph, then S is true but the learner will never be certain of this; and it is not falsifiable with certainty, since if the car’s speed is exactly 51 mph, then S is false but the learner will never be certain of that. Nevertheless, S is *inductively decidable*, e.g. by the learner defined by: $\mathbb{L}(a, b) := (a, b) \cap S$ for open intervals with $a < 50 < b$; $\mathbb{L}(a, b) := (a, b)$ for open intervals (a, b) s.t. either $(a, b) \subseteq S$ or $(a, b) \cap S = \emptyset$; and $\mathbb{L}(a, b) := [51, b)$ whenever $50 < a < 51 < b$.

Dependence on the Learner It is easy to see that learnability (verifiability, falsifiability, decidability) with certainty are *learner-independent* notions (since they are directed towards achieving infallible knowledge), so they do not depend on \mathbb{L} but only on the underlying intersection model. In contrast, *the corresponding inductive notions above are learner-dependent*. As a consequence, the interesting concepts in Learning Theory are obtained from them by *quantifying existentially over learners*: a proposition P is *inductively learnable* (verifiable, falsifiable, decidable) if there exists *some* learner \mathbb{L} s.t. P is respectively inductively learnable verifiable falsifiable decidable by \mathbb{L} . This property of a learning frame is called *generic inductive learnability*.

Topological Characterizations As it is well-known in learning theory and formal epistemology [24, 16], the above notions are *topological* in nature: P is learnable with certainty at world x iff x is in the *interior* of P with respect to the topology generated by \mathcal{O} ; P is verifiable with certainty iff it is *open* in the same topology; P is falsifiable with certainty iff it is *closed* in this topology; finally, P is decidable with certainty iff it is *clopen*. The corresponding inductive notions can be easily characterized [16] in the case that the topology generated by \mathcal{O} satisfies the separation condition⁸ $T1$: in this case, P is inductively verifiable iff it is Σ_2 in the Borel hierarchy for this topology (i.e. a countable union of closed sets); in the same conditions, P is inductively falsifiable iff it is Π_2 (a countable intersection of open sets), and it is inductively decidable iff it is Δ_2 (i.e. Σ_2 and Π_2). More recently, in work by three of this paper’s coauthors [5], these characterisations were generalised to arbitrary topologies satisfying the weaker separation condition⁹ $T0$; in particular, P is inductively verifiable iff it is a countable union of locally closed sets.¹⁰

2 Dynamic Logic for Learning Theory

In this section we introduce our ‘dynamic logic for learning theory’ DLLT. As already mentioned, this is obtained by adding two ingredients to the language of Subset Space Logics: *dynamic observation modalities* $[o]\varphi$ and a *learning operator* $L(\vec{o})$.

2.1 Syntax and Semantics of DLLT

Let $\text{Prop} = \{p, q, \dots\}$ be a countable set of *propositional variables*, denoting arbitrary ‘ontic’ (i.e., non-epistemic) facts that might hold in a world (even if they might never be observed), and let $\text{Prop}_{\mathcal{O}} = \{o, u, v, \dots\}$ be a countable set of *observational variables*, denoting ‘observable facts’ (which, if true, will eventually be observed).

Definition 1. *The syntax of our language \mathcal{L} is defined by the grammar:*

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

where $p \in \text{Prop}$ and $o \in \text{Prop}_{\mathcal{O}}$, while $\vec{o} = (o_1, \dots, o_n) \in \text{Prop}_{\mathcal{O}}^*$ is a finite sequence of observational variables. (In particular, empty sequence is denoted by λ .) We employ the usual abbreviations for propositional connectives $\top, \perp, \vee, \rightarrow, \leftrightarrow$ and for the dual modalities $\langle K \rangle, \diamond, \langle o \rangle$.

The informal meaning of our formulas is as follows. Propositional variables denote *ontic facts* (i.e. factual, non-epistemic features of a world), while observational variables o denote *observable facts* (i.e. facts that, if true, will eventually be observed).

⁸ This topology is $T1$ iff for every two distinct points $x \neq y$ there exist an observation $O \in \mathcal{O}$ with $x \in O$ and $y \notin O$.

⁹ The observational topology is $T0$ iff points can be distinguished by observations; i.e. if x and y satisfy the same observable properties in \mathcal{O} , then $x = y$. Obviously, $T0$ is a minimally necessary condition for any kind of learnability of the real world from observations.

¹⁰ A set is locally closed if it is the intersection of a closed and an open set.

We read $K\varphi$ as ‘the learner *knows* φ (with absolute certainty)’. $\Box\varphi$ is the so-called ‘effort modality’ from Subset Space Logic; we read $\Box\varphi$ as ‘ φ is *stably true*’. Indeed, $\Box\varphi$ holds iff φ is true and will stay true *no matter what new (true) evidence is observed*. The operator $[o]\varphi$ is similar to the operator $[\psi]\varphi$ in Public Announcement Logic, but it is restricted to the cases when ψ is a particular kind of atomic formula, namely an observational variable $o \in \text{Prop}_\mathcal{O}$. So we read $[o]\varphi$ as ‘after o is observed, φ will hold’. Finally, $L(\vec{o})$ denotes the learner’s *conjecture* given observations \vec{o} ; i.e. her strongest belief (i.e., the set of worlds considered to be most plausible) after observing \vec{o} .

Definition 2 (Intersection Frame/Model and Learning Frame/Model). An *intersection frame* [17, 9] is a pair (X, \mathcal{O}) , 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 ‘observables’, or ‘evidence’), which is assumed to be *closed under finite intersections*: if $\mathcal{F} \subseteq \mathcal{O}$ is finite then $(\bigcap \mathcal{F}) \in \mathcal{O}$. An *intersection model* $(X, \mathcal{O}, \|\cdot\|)$ is an intersection frame (X, \mathcal{O}) together with a valuation map $\|\cdot\| : \text{Prop} \cup \text{Prop}_\mathcal{O} \rightarrow \mathcal{P}(X)$, that maps propositional variables p into arbitrary sets $\|p\| \subseteq X$ and observational variables o into observable properties $\|o\| \in \mathcal{O}$.

A *learning frame* is a triplet $(X, \mathcal{O}, \mathbb{L})$, where (X, \mathcal{O}) is an intersection frame and $\mathbb{L} : \mathcal{O} \rightarrow \mathcal{P}(X)$ is a *learner*, i.e. a map associating to every information state $O \in \mathcal{O}$ some ‘conjecture’ $\mathbb{L}(O) \subseteq X$, and satisfying two properties: (1) $\mathbb{L}(O) \subseteq O$ (*conjectures fit the evidence*), and (2) if $O \neq \emptyset$ then $\mathbb{L}(O) \neq \emptyset$ (*consistency of conjectures* based on consistent evidence). We can extend \mathbb{L} to range over *strings* of information states $\vec{O} = (O_1, \dots, O_n) \in \mathcal{O}^*$ in a natural way, by putting $\mathbb{L}(\vec{O}) := \mathbb{L}(\bigcap \vec{O})$, where $\bigcap \vec{O} := O_1 \cap \dots \cap O_n$. A *learning model* $\mathcal{M} = (X, \mathcal{O}, \mathbb{L}, \|\cdot\|)$ is a learning frame $(X, \mathcal{O}, \mathbb{L})$ together with a valuation map $\|\cdot\| : \text{Prop} \cup \text{Prop}_\mathcal{O} \rightarrow \mathcal{P}(X)$ as above; equivalently, it consists of an intersection model $(X, \mathcal{O}, \|\cdot\|)$ together with a learner, as defined above.

Intuitively, the states in X represent possible worlds. The tautological evidence $X = \bigcap \emptyset$ represents the state of ‘no information’ (before anything is observed), while the contradictory evidence \emptyset represents inconsistent information. Finally, $\mathbb{L}(O)$ represents the learner’s conjecture after observing O , while $\mathbb{L}(O_1, \dots, O_n) = \mathbb{L}(O_1 \cap \dots \cap O_n)$ represents the conjecture after observing a finite sequence of observations O_1, \dots, O_n . (The fact that \mathcal{O} is closed under finite intersections is important here for identifying any finite sequence O_1, \dots, O_n with a single observation $O = O_1 \cap \dots \cap O_n \in \mathcal{O}$.)

Epistemic Scenarios. As in Subset Space Semantics, the formulas of our logic are not interpreted at possible worlds, but at so-called *epistemic scenarios*, i.e. 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. We denote by $ES(\mathcal{M}) := \{(x, U) : x \in U \in \mathcal{O}\}$ the set of all epistemic scenarios.

Definition 3 (Semantics). Given a learning model $\mathcal{M} = (X, \mathcal{O}, \mathbb{L}, \|\cdot\|)$ and an epistemic scenario (x, U) , the semantics of the language \mathcal{L} is given by a binary relation $(x, U) \models_{\mathcal{M}} \varphi$ between epistemic scenario and formulas, called the *satisfaction relation*, as well as a *truth set* (interpretation) $\llbracket \varphi \rrbracket_{\mathcal{M}}^U := \{x \in U \mid (x, U) \models_{\mathcal{M}} \varphi\}$, for all formulas φ . We typically omit the subscript, simply writing $(x, U) \models \varphi$ and $\llbracket \varphi \rrbracket^U$, whenever the model \mathcal{M} is understood. The satisfaction relation is defined by the following recursive clauses:

$(x, U) \models p$	iff $x \in \ p\ $
$(x, U) \models o$	iff $x \in \ o\ $
$(x, U) \models L(o_1, \dots, o_n)$	iff $x \in \mathbb{L}(U, \ o_1\ , \dots, \ o_n\)$
$(x, U) \models \neg\varphi$	iff $(x, U) \not\models \varphi$
$(x, U) \models \varphi \wedge \psi$	iff $(x, U) \models \varphi$ and $(x, U) \models \psi$
$(x, U) \models K\varphi$	iff $(\forall y \in U) ((y, U) \models \varphi)$
$(x, U) \models \Box\varphi$	iff $(\forall O \in \mathcal{O}) (x \in O \subseteq U \text{ implies } (x, O) \models \varphi)$
$(x, U) \models [o]\varphi$	iff $(\forall O \in \mathcal{O}) (x \in O \text{ implies } (x, U \cap O) \models \varphi)$
$(x, U) \models \langle o \rangle \varphi$	iff $(x \in \ o\ \text{ implies } (x, U \cap \ o\) \models \varphi)$

where $p \in \text{Prop}$, $o, o_1, \dots, o_n \in \text{Prop}_{\mathcal{O}}$, $\vec{o} \in \text{Prop}_{\mathcal{O}}^*$, and where we used the notation $L(O_1, \dots, O_n) := L(O_1 \cap \dots \cap O_n)$ introduced above. We say that a formula φ is *valid in a learning model* \mathcal{M} , and write $\mathcal{M} \models \varphi$, if $(x, U) \models_{\mathcal{M}} \varphi$ for all epistemic scenarios $(x, U) \in ES(\mathcal{M})$. We say φ is *validable in an intersection model* $(X, \mathcal{O}, \|\cdot\|)$, and write $(X, \mathcal{O}, \|\cdot\|) \models \varphi$, if there exists some learner $\mathbb{L} : \mathcal{O} \rightarrow \mathcal{P}(X)$ such that φ is valid in the learning model $(X, \mathcal{O}, \mathbb{L}, \|\cdot\|)$. We say φ is *valid*, and write $\models \varphi$, if it is valid in *all* learning models.

Abbreviations: For any string $\vec{o} = (o_1, \dots, o_n) \in \text{Prop}_{\mathcal{O}}^*$ of observational variables, and any formula φ we set:

$$\begin{aligned} \bigwedge \vec{o} &:= o_1 \wedge \dots \wedge o_n \text{ (with the convention that } \bigwedge \lambda := \top) \\ \vec{o} \Leftrightarrow \vec{u} &:= K((\bigwedge \vec{o}) \leftrightarrow (\bigwedge \vec{u})) \text{ (extensional equivalence of observations)} \\ [\vec{o}]\varphi &:= [o_1] \dots [o_n]\varphi \text{ (with the convention that } [\lambda]\varphi := \varphi); \text{ similarly for } \langle \vec{o} \rangle \\ B^{\vec{o}}\varphi &:= K(L(\vec{o}) \rightarrow \varphi) \\ B\varphi &:= B^\lambda\varphi \end{aligned}$$

(where λ is the empty string). We read $B\varphi$ as the ‘observer *believes* φ ’ (given no observations), and $B^{\vec{o}}\varphi$ as ‘the observer *believes* φ conditional on evidence \vec{o} ’.

2.2 Axiomatization and proof system

We will now provide the formal definition of our proposed system \mathbf{L} of the Dynamic Logic for Learning Theory (DLLT) by listing the axioms and derivation rules, see Table 1 below. Given a formula $\varphi \in \mathcal{L}$, we denote by P_φ and O_φ the set of all propositional variables and observational variables respectively occurring in φ (we will use the same notation for the necessity and possibility forms defined below).

The intuitive reading of the $S5$ axioms for epistemic modality K expresses the fact that K is factive and (positively and negatively) introspective. The intuitive nature of the reduction axioms should be as in Public Announcement Logic [2], when we take into account the natural atomic behaviour of observables. The learning axioms (CC), (EC) and (SP) express pre-conditions in formal learning theory on *observations*, namely that: they are truthful observations about the world (CC); that the history of observations is irrelevant for the learner, except for the extensional evidence provided by observations

Table 1. The axiom schemas for the Dynamic Logic of Learning Theory, **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 _[o])	$[o](\psi \rightarrow \chi) \rightarrow ([o]\psi \rightarrow [o]\chi)$	
Basic rules:		
(MP)	From $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$, infer $\vdash \psi$	
(Nec _K)	From $\vdash \varphi$, infer $\vdash K\varphi$	
(Nec _[o])	From $\vdash \varphi$, infer $\vdash [o]\varphi$	
Learning axioms:		
(CC)	$(\bigwedge \vec{\sigma}) \rightarrow \langle K \rangle L(\vec{\sigma})$	Consistency of Conjectures
(EC)	$(\vec{\sigma} \leftrightarrow \vec{u}) \rightarrow (L(\vec{\sigma}) \leftrightarrow L(\vec{u}))$	Extensionality of Conjectures
(SP)	$L(\vec{\sigma}) \rightarrow \bigwedge \vec{\sigma}$	Success Postulate
Reduction axioms:		
(R _p)	$[o]p \leftrightarrow (o \rightarrow p)$	
(R _u)	$[o]u \leftrightarrow (o \rightarrow u)$	
(R _L)	$[o]L(\vec{u}) \leftrightarrow (o \rightarrow L(o, \vec{u}))$	
(R _¬)	$[o]\neg\psi \leftrightarrow (o \rightarrow \neg[o]\psi)$	
(R _K)	$[o]K_i\psi \leftrightarrow (o \rightarrow K[o]\psi)$	
(R _□)	$[o]\Box\psi \leftrightarrow \Box[o]\psi$	
Effort axiom and rule:		
(□-Ax)	$\Box\varphi \rightarrow [\vec{\sigma}]\varphi$ for all $\vec{\sigma} \in \text{Prop}_{\mathcal{O}}^*$ arbitrary	
(□-Rule)	From $\vdash \psi \rightarrow [o]\varphi$, infer $\vdash \psi \rightarrow \Box\varphi$, where $o \notin \mathcal{O}_\psi \cup \mathcal{O}_\varphi$	

(EC) and that conjectures fit what is observed (SP). Since the effort modality \Box quantifies over possible observations, we could think of the Effort axiom and the Effort rule as elimination and introduction rules for \Box . The former one expresses the fact that: if a property is stably true then it holds after any observations. Finally, the latter says that, if a property holds after any arbitrary observation, it is stably true.

So each of our axioms is simple and readable and has a transparent intuitive interpretation, in contrast to other axiomatizations of (the less expressive) Subset Space Logic over intersection spaces (i.e., the L -free analogues of our models). Having such a simple axiomatization is one of the advantages brought by the addition of dynamic observation modalities. See more discussion of this issue in the Conclusions section.

We now give reduction laws for strings of observational variables in $\text{Prop}_{\mathcal{O}}^*$.

Proposition 1 (Reduction laws for strings of observational variables). *The following reduction laws are provable in **L** for all $\varphi \in \mathcal{L}$:*

1. $[\vec{u}]p \leftrightarrow (\bigwedge \vec{u} \rightarrow p)$
2. $[\vec{u}]o \leftrightarrow (\bigwedge \vec{u} \rightarrow o)$
3. $[\vec{u}]L(\vec{\sigma}) \leftrightarrow (\bigwedge \vec{u} \rightarrow L(\vec{u}, \vec{\sigma}))$
4. $[\vec{u}]\neg\varphi \leftrightarrow (\bigwedge \vec{u} \rightarrow \neg[\vec{u}]\varphi)$

5. $[\vec{u}]K\varphi \leftrightarrow (\bigwedge \vec{u} \rightarrow K[\vec{u}]\varphi)$
6. $[\vec{u}]\Box\varphi \leftrightarrow \Box[\vec{u}]\varphi$

Proposition 2. *The following reduction laws are provable in \mathbf{L} for all formulas $\varphi \in \mathcal{Q}$:*

$$\begin{array}{ll}
(R_{\wedge}) & [u](\varphi \wedge \psi) \leftrightarrow ([u]\varphi \wedge [u]\psi) \\
(R_{\rightarrow \wedge}) & [\vec{u}](\varphi \wedge \psi) \leftrightarrow ([\vec{u}]\varphi \wedge [\vec{u}]\psi) \\
\langle o \rangle & \langle o \rangle \psi \leftrightarrow (o \wedge [o]\psi) \\
\langle \vec{o} \rangle & \langle \vec{o} \rangle \psi \leftrightarrow (\bigwedge \vec{o} \wedge [\vec{o}]\psi)
\end{array}$$

In our framework, belief (B) and conditional beliefs ($B^{\vec{o}}\varphi$) are defined in terms of the operators K and L . The axiomatic system \mathbf{L} given in Table 1 over the language \mathcal{Q} can therefore derive the properties describing the type of belief and conditional belief modalities we intend to formalize in this paper. More precisely, as stated in Proposition 3, the system \mathbf{L} yields the standard belief system KD45 for B . More generally, if we replace the D axiom for a ‘weaker’ version $D' := \langle \langle K \rangle \vec{o} \rightarrow \neg B^{\vec{o}} \perp$ then we have a weak version of KD45 system for conditional belief $B^{\vec{o}}$, namely wKD45.

Proposition 3 (wKD45 axioms and rules for Conditional Belief). *The standard axioms and rules of the doxastic logic KD45 are derivable for our belief operator B in the system \mathbf{L} . More generally, the axioms of rules of the weaker system wKD45 are derivable for our conditional-belief operator $B^{\vec{o}}$ in the system \mathbf{L} .*

Proposition 4. *The S4 axioms for the effort modality \Box are derivable in \mathbf{L} .*

3 Soundness and Completeness

In this section we prove soundness and completeness. Note that, although our logic is more expressive than Subset Space Logic (interpreted on intersection spaces), our completeness proof is *much simpler*, via a *canonical construction*: this is one of the advantages of having the (expressively redundant) dynamic observation modalities!

Soundness We first prove soundness, for which we need the following lemma. Note that by the definition of the valuation $\|\cdot\|$ in a learning model \mathcal{M} , we have that for all $U \in \mathcal{O}$, $U \cap \|p\| = \llbracket p \rrbracket_{\mathcal{M}}^U$ and $U \cap \|o\| = \llbracket o \rrbracket_{\mathcal{M}}^U$ for all $p \in \text{Prop}$ and $o \in \text{Prop}_{\mathcal{O}}$ in \mathcal{Q} .

Lemma 1. *Let $\mathcal{M} = (X, \mathcal{O}, \mathbb{L}, \|\cdot\|)$ and $\mathcal{M}' = (X, \mathcal{O}, \mathbb{L}, \|\cdot\|')$ be two learning models and $\varphi \in \mathcal{Q}$ such that \mathcal{M} and \mathcal{M}' differ only in the valuation of some $o \notin O_{\varphi}$. Then, for all $U \in \mathcal{O}$, we have $\llbracket \varphi \rrbracket_{\mathcal{M}}^U = \llbracket \varphi \rrbracket_{\mathcal{M}'}^U$.*

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

Completeness We now move to the completeness proof for our logic \mathbf{L} , which will be shown via a ‘simple’ canonical model construction. But its simplicity is deceiving, due to two main technical differences between our construction and the standard canonical model from Modal Logic. First, this is *not* a relational (Kripke) model, but a *neighborhood model*; so the closest analogue is the type of canonical construction used in Topological Modal Logic or Neighborhood Semantics [1]. Second, the standard notion of maximally consistent theory is *not* very useful for our logic, since such theories do not

‘internalize’ the \Box -Rule. 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*, in the following sense:

Definition 4 (‘Pseudo-modalities’: necessity and possibility forms). *The set of necessity-form expressions of our language is given by $NF_{\mathfrak{L}} := (\{\varphi \rightarrow \mid \varphi \in \mathfrak{L}\} \cup \{K\} \cup \text{Prop}_{\mathcal{O}})^*$. For any finite string $s \in NF_{\mathfrak{L}}$, we define ‘pseudo-modalities’ $[s]$ (called necessity form) and $\langle s \rangle$ (called possibility form), that generalize our dynamic modalities $[o]$ and $\langle o \rangle$. These pseudo-modalities are just functions mapping any formula $\varphi \in \mathfrak{L}$ to another formula $[s]\varphi \in \mathfrak{L}$, and respectively $\langle s \rangle\varphi \in \mathfrak{L}$. Necessity forms are defined recursively, by putting: $[\lambda]\varphi := \varphi$, $[\varphi \rightarrow, s]\varphi := \varphi \rightarrow [s]\varphi$, $[K, s]\varphi := K[s]\varphi$, $[o, s]\varphi := [o][s]\varphi$. As for possibility forms, we put $\langle s \rangle\varphi := \neg[s]\neg\varphi$.*

Lemma 2. *For every necessity form $[s]$, there exist an observational variable $o \in \mathcal{O}$ and a formula $\psi \in \mathfrak{L}$ such that for all $\varphi \in \mathfrak{L}$, we have: $\vdash [s]\varphi$ iff $\vdash \psi \rightarrow [o]\varphi$.*

Proof. The proof follows similarly as in [2, Lemma 4.8], but is even simpler since our dynamic modalities only involve observational variables which are atomic formulas.

Lemma 3. *The following rule is admissible in \mathbf{L} :*

$$\text{if } \vdash [s][o]\varphi \text{ then } \vdash [s]\Box\varphi, \text{ where } o \notin \mathcal{O}_s \cup \mathcal{O}_\varphi.$$

Proof. Suppose $\vdash [s][o]\varphi$ where $o \notin \mathcal{O}_s \cup \mathcal{O}_\varphi$. Then, by Lemma 2, there exist $u \in \text{Prop}_{\mathcal{O}}$ and $\psi \in \mathfrak{L}$ such that $\vdash \psi \rightarrow [u][o]\varphi$. Thus we get $\vdash \psi \rightarrow [u, o]\varphi$. It is not hard to see that $\vdash [u, o]\varphi \leftrightarrow [o, u]\varphi$ (which follows by subformula induction on φ , using the corresponding reduction axiom given in Proposition 1, and the fact that $\vdash u \wedge o \leftrightarrow o \wedge u$). Therefore, $\vdash \psi \rightarrow [u, o]\varphi$ iff $\vdash \psi \rightarrow [o, u]\varphi$. Hence we obtain $\vdash \psi \rightarrow [o, u]\varphi$, i.e., $\vdash \psi \rightarrow [o][u]\varphi$. By the construction of the formulas ψ and u , we know that $\mathcal{O}_\psi \cup \mathcal{O}_u \subseteq \mathcal{O}_s$, and so $o \notin \mathcal{O}_\psi \cup \{u\} \cup \mathcal{O}_\varphi$. Therefore, by the Effort rule (\Box -Rule) we have $\vdash \psi \rightarrow \Box[u]\varphi$, implying, by the reduction axiom (R_{\Box}), that $\vdash \psi \rightarrow [u]\Box\varphi$. Applying again Lemma 2, we obtain $\vdash [s]\Box\varphi$.

Definition 5. *For every countable set \mathcal{O} , let $\mathfrak{L}^{\mathcal{O}}$ be the language of the logic $\mathbf{L}^{\mathcal{O}}$ based only on the observational variables in \mathcal{O} (i.e. having as set of observational variables $\text{Prop}_{\mathcal{O}} := \mathcal{O}$). Let $NF_{\mathfrak{L}^{\mathcal{O}}}$ denote the set of necessity-form expressions of $\mathbf{L}^{\mathcal{O}}$ (i.e. necessity forms involving only observational variables in \mathcal{O}). An \mathcal{O} -theory is a consistent set of formulas in $\mathfrak{L}^{\mathcal{O}}$. Here, ‘consistent’ means consistent with respect to the axiomatization \mathbf{L} formulated for $\mathfrak{L}^{\mathcal{O}}$. A maximal \mathcal{O} -theory is an \mathcal{O} -theory Γ that is maximal with respect to \subseteq among all \mathcal{O} -theories; in other words, Γ cannot be extended to another \mathcal{O} -theory. An \mathcal{O} -witnessed theory is an \mathcal{O} -theory Γ such that, for every $s \in NF_{\mathfrak{L}^{\mathcal{O}}}$ and $\varphi \in \mathfrak{L}^{\mathcal{O}}$, if $\langle s \rangle\Diamond\varphi$ is consistent with Γ then there is $o \in \mathcal{O}$ such that $\langle s \rangle\langle o \rangle\varphi$ is consistent with Γ . A maximal \mathcal{O} -witnessed theory Γ is an \mathcal{O} -witnessed theory that is not a proper subset of any \mathcal{O} -witnessed theory.*

Lemma 4. *For every $\Gamma \subseteq \mathfrak{L}^{\mathcal{O}}$, if Γ is an \mathcal{O} -theory and $\Gamma \not\vdash \neg\varphi$ for some $\varphi \in \mathfrak{L}^{\mathcal{O}}$, then $\Gamma \cup \{\varphi\}$ is an \mathcal{O} -theory. Moreover, if Γ is \mathcal{O} -witnessed, then $\Gamma \cup \{\varphi\}$ is also \mathcal{O} -witnessed.*

Lemma 5 (Lindenbaum’s Lemma). *Every O-witnessed theory Γ can be extended to a maximal O-witnessed theory T_Γ .*

Lemma 6 (Extension Lemma). *Let O be a set of observational variables and O' be a countable set of fresh observational variables, i.e., $O \cap O' = \emptyset$. Let $\tilde{O} = O \cup O'$. Then, every O-theory Γ can be extended to a \tilde{O} -witnessed theory $\tilde{\Gamma} \supseteq \Gamma$, and hence to a maximal \tilde{O} -witnessed theory $T_{\tilde{\Gamma}} \supseteq \Gamma$.*

We are now ready to build the canonical model.

Canonical Model for T_0 . For any consistent set of formulas Φ , consider a maximally consistent O-witnessed extension $T_0 \supseteq \Phi$. As our canonical set of worlds, we take the set $\mathcal{X}^c := \{T : T \text{ maximally consistent O-witnessed theory with } T \sim_K T_0\}$, where we put $T \sim_K T'$ iff $\forall \varphi \in \mathcal{L}_O ((K\varphi) \in T \rightarrow \varphi \in T')$. As usual, 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 $\widehat{\varphi} := \{T \in \mathcal{X}^c : \varphi \in T\}$. In particular, for any observational variable $o \in O$, we have $\widehat{o} = \{T \in \mathcal{X}^c : o \in T\}$. We can generalize this notation to *finite sequences* $\vec{o} = (o_1, \dots, o_n) \in O^*$ of observational variables, by putting: $\widehat{\vec{o}} := \{T \in \mathcal{X}^c : o_1, \dots, o_n \in T\}$.

As canonical set of information states, we take $\mathcal{O}^c := \{\widehat{\vec{o}} : \vec{o} \in O^*\}$. Finally, our canonical learner is given by $\mathbb{L}^c(\widehat{\vec{o}}) := \widehat{L(\vec{o})}$, and the canonical valuation $\|\cdot\|_c$ is given as $\|p\|_c = \widehat{p}$ and $\|o\|_c = \widehat{o}$. The learning model $\mathcal{M}^c = (\mathcal{X}^c, \mathcal{O}^c, \mathbb{L}^c, \|\cdot\|_c)$ is called the *canonical model*. Note that we use c as a subindex instead of a superindex for the canonical valuation $\|\cdot\|_c$, this is in order to avoid confusion with our ‘open-restriction’ notation for the truth set of a formula $\llbracket \varphi \rrbracket^U$.

Before proving that the canonical model is well-defined, we need the following.

Lemma 7. *For every maximal O-witnessed theory T , the set $\{\theta : K\theta \in T\}$ is an O-witnessed theory.*

Lemma 8. *Let $T \in \mathcal{X}^c$. Then, $K\varphi \in T$ iff $\varphi \in S$ for all $S \in \mathcal{X}^c$.*

Corollary 1. *Let $T \in \mathcal{X}^c$. Then, $\langle K \rangle \varphi \in T$ iff there is $S \in \mathcal{X}^c$, such that $\varphi \in S$.*

Proposition 5. *The canonical model is well-defined.*

Proof. We need to show that the following properties hold:

1. *If $F = \{\widehat{o}_1, \dots, \widehat{o}_m\} \subseteq \mathcal{O}^c$ is finite then $\bigcap F \in \mathcal{O}^c$:* Let $F = \{\widehat{o}_1, \dots, \widehat{o}_m\} \subseteq \mathcal{O}^c$. It is easy to see that $\bigcap \{\widehat{o}_1, \dots, \widehat{o}_m\} = \widehat{\vec{o}}$, where \vec{o} is the concatenation of all the \vec{o}_i ’s with $1 \leq i \leq m$. Since each \vec{o}_i is finite, \vec{o} is finite. Therefore, by the definition of \mathcal{O}^c , we obtain $\bigcap \{\widehat{o}_1, \dots, \widehat{o}_m\} = \widehat{\vec{o}} \in \mathcal{O}^c$.
2. *\mathbb{L}^c is a well-defined function and a learner:* For this, note that $\mathbb{L}^c(\widehat{\vec{o}}) := \widehat{L(\vec{o})} \subseteq \mathcal{X}^c$. We will first prove that:
 - (2a) *if $\widehat{\vec{o}} = \widehat{\vec{u}}$ then $\mathbb{L}^c(\widehat{\vec{o}}) = \mathbb{L}^c(\widehat{\vec{u}})$:* Suppose $\widehat{\vec{o}} = \widehat{\vec{u}}$. This means that $(\forall T \in \mathcal{X}^c)(\bigwedge \vec{o} \in T \text{ iff } \bigwedge \vec{u} \in T)$. Therefore, we obtain $\vdash \bigwedge \vec{o} \leftrightarrow \bigwedge \vec{u}$. Then, by

(Nec_K), we have $\vdash K(\bigwedge \vec{\sigma} \leftrightarrow \bigwedge \vec{u})$, i.e., $\vdash \vec{\sigma} \Leftrightarrow \vec{u}$. Since $\mathbb{L}^c(\widehat{\vec{\sigma}}) := \widehat{L(\vec{\sigma})}$, showing $\mathbb{L}^c(\widehat{\vec{\sigma}}) = \mathbb{L}^c(\widehat{\vec{u}})$ boils down to showing that $\widehat{L(\vec{\sigma})} = \widehat{L(\vec{u})}$, i.e., that $\vdash L(\vec{\sigma}) \leftrightarrow L(\vec{u})$, which follows from axiom (EC) and the assumption that $\vdash \vec{\sigma} \Leftrightarrow \vec{u}$.

Next, we must prove that

(2b) \mathbb{L}^c is a learner, i.e., \mathbb{L}^c satisfies the properties of a learner given in Definition 2. To show this, we first check that $\mathbb{L}^c(\widehat{\vec{\sigma}}) \subseteq \widehat{\vec{\sigma}}$ holds. Let $T \in \mathbb{L}^c(\widehat{\vec{\sigma}})$. This means, by the definition of $\mathbb{L}^c(\widehat{\vec{\sigma}})$, that $L(\vec{\sigma}) \in T$. Since $(L(\vec{\sigma}) \rightarrow \bigwedge \vec{\sigma}) \in T$ (by the axiom (SP)), we have that $\bigwedge \vec{\sigma} \in T$. Therefore, as T is maximally consistent, we obtain $o_1, \dots, o_m \in T$ for $\vec{\sigma} = (o_1, \dots, o_m)$, meaning that $\vec{\sigma} \in T$. Thus, $T \in \widehat{\vec{\sigma}}$. Finally we show that if $\widehat{\vec{\sigma}} \neq \emptyset$ then $\mathbb{L}^c(\widehat{\vec{\sigma}}) \neq \emptyset$. Suppose $\widehat{\vec{\sigma}} \neq \emptyset$, i.e., there is $T \in X^c$ with $T \in \widehat{\vec{\sigma}}$. This means, by the definition of $\widehat{\vec{\sigma}}$, that $\vec{\sigma} \in T$. Then, since T is a maximal consistent theory, we have $\bigwedge \vec{\sigma} \in T$, and $((\bigwedge \vec{\sigma}) \rightarrow \langle K \rangle L(\vec{\sigma})) \in T$ (by the axiom (CC)). Thus we obtain $\langle K \rangle L(\vec{\sigma}) \in T$. Then, by Corollary 1, there is $S \in X^c$ such that $L(\vec{\sigma}) \in S$. Thus, by the definition of $\widehat{L(\vec{\sigma})}$, we have $S \in \widehat{L(\vec{\sigma})}$, and therefore, $\widehat{L(\vec{\sigma})} = \mathbb{L}^c(\widehat{\vec{\sigma}}) \neq \emptyset$.

Our aim is to prove a Truth Lemma for the canonical model, that will immediately imply completeness, as usual. But for this we first need the following result.

Lemma 9. *Let $T \in X^c$. Then, $\Box\varphi \in T$ iff $[\vec{u}]\varphi \in T$ for all $\vec{u} \in \text{Prop}_O^*$.*

We now proceed to our key result:

Lemma 10 (Truth Lemma). *For all formulas φ , all $T \in X^c$ and all $\widehat{\vec{\sigma}} \in \mathcal{O}^c$, we have:*

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

Proof. The proof is by induction over subformulas. The cases for propositional and observational variables, as well as for Boolean connectives are as usual. So we only check the remaining cases. At each step of the proof, $\bigwedge \vec{\sigma} \in T$ guarantees that the pair $(T, \widehat{\vec{\sigma}})$ is a well-defined epistemic scenario of the canonical model since $\widehat{\bigwedge \vec{\sigma}} = \widehat{\vec{\sigma}}$.

– Case $\varphi := L(\vec{u})$.

$$\begin{aligned} \langle \vec{\sigma} \rangle L(\vec{u}) \in T &\text{ iff } (\bigwedge \vec{\sigma} \wedge [\vec{\sigma}]L(\vec{u})) \in T && \text{(Proposition 2-}(\langle \vec{\sigma} \rangle)) \\ &\text{ iff } (\bigwedge \vec{\sigma} \wedge L(\vec{\sigma}, \vec{u})) \in T && \text{(Proposition 1.3)} \\ &\text{ iff } \bigwedge \vec{\sigma} \in T \text{ and } L(\vec{\sigma}, \vec{u}) \in T \\ &\text{ iff } T \in \widehat{\vec{\sigma}} \text{ and } T \in \widehat{L(\vec{\sigma}, \vec{u})} = \mathbb{L}^c(\widehat{\vec{\sigma}}, \widehat{\vec{u}}) && \text{(since } \widehat{\bigwedge \vec{\sigma}} = \widehat{\vec{\sigma}}) \\ &\text{ iff } (T, \widehat{\vec{\sigma}}) \models_{\mathcal{M}^c} L(\vec{u}) && \text{(by the semantics of } L) \end{aligned}$$

– Case $\varphi := K\psi$.

$$\begin{aligned} \langle \vec{\sigma} \rangle K\psi \in T &\text{ iff } (\bigwedge \vec{\sigma} \wedge K[\vec{\sigma}]\psi) \in T && \text{(Propositions 2-}(\langle \vec{\sigma} \rangle) \text{ and 1.5)} \\ &\text{ iff } \bigwedge \vec{\sigma} \in T \text{ and } K[\vec{\sigma}]\psi \in T \end{aligned}$$

$$\begin{aligned}
& \text{iff } \bigwedge \vec{\sigma} \in T \text{ and } (\forall S \sim_K T)([\vec{\sigma}] \psi \in S) && \text{(by Lemma 8)} \\
& \text{iff } \bigwedge \vec{\sigma} \in T \text{ and } (\forall S \in \widehat{\vec{\sigma}})(\langle \vec{\sigma} \rangle \psi \in S) && \text{(Propositions 2-}\langle \vec{\sigma} \rangle\text{)} \\
& \text{iff } \vec{\sigma} \in T \text{ and } (\forall S \in \widehat{\vec{\sigma}})((S, \vec{\sigma}) \models \psi) && \text{(by } \widehat{\bigwedge \vec{\sigma}} = \widehat{\vec{\sigma}} \text{ and I.H)} \\
& \text{iff } (T, \widehat{\vec{\sigma}}) \models_{\mathcal{M}^c} K\psi && \text{(by the semantics of } K\text{)}
\end{aligned}$$

– Case $\varphi := \langle \vec{u} \rangle \psi$.

$$\begin{aligned}
\langle \vec{\sigma} \rangle \langle \vec{u} \rangle \psi \in T & \text{ iff } \langle \vec{\sigma}, \vec{u} \rangle \psi \in T && \text{(by the abbreviation for } \langle \vec{\sigma} \rangle \psi\text{)} \\
& \text{iff } (\bigwedge \vec{\sigma} \wedge \bigwedge \vec{u}) \wedge \langle \vec{\sigma}, \vec{u} \rangle \psi \in T && \text{(Propositions 2-}\langle \vec{\sigma} \rangle\text{)} \\
& \text{iff } (\bigwedge \vec{\sigma} \wedge \bigwedge \vec{u}) \in T \text{ and } \langle \vec{\sigma}, \vec{u} \rangle \psi \in T \\
& \text{iff } T \in \widehat{\vec{\sigma}} \cap \widehat{\vec{u}} \text{ and } \langle \vec{\sigma}, \vec{u} \rangle \psi \in T && \text{(since } \widehat{\bigwedge \vec{\sigma} \wedge \bigwedge \vec{u}} = \widehat{\vec{\sigma}} \cap \widehat{\vec{u}}\text{)} \\
& \text{iff } T \in (\widehat{\vec{\sigma}}, \vec{u}) \text{ and } (T, (\widehat{\vec{\sigma}}, \vec{u})) \models \psi && \text{(since } (\widehat{\vec{\sigma}}, \vec{u}) = \widehat{\vec{\sigma}} \cap \widehat{\vec{u}}\text{)} \\
& \text{iff } T \in \|\vec{\sigma}, \vec{u}\|_c \text{ and } (T, \|\vec{\sigma}, \vec{u}\|_c) \models \psi && \text{(by the definition of } \|\cdot\|_c\text{)} \\
& \text{iff } (T, \|\vec{\sigma}, \vec{u}\|_c) \models \langle \vec{\sigma}, \vec{u} \rangle \psi && \text{(by the semantics)} \\
& \text{iff } (T, \|\vec{\sigma}, \vec{u}\|_c) \models \langle \vec{\sigma} \rangle \langle \vec{u} \rangle \psi && \text{(by the abbreviation for } \langle \vec{\sigma} \rangle \psi\text{)}
\end{aligned}$$

– Case $\varphi := \Box\psi$.

(\Leftarrow) Suppose $\langle \vec{\sigma} \rangle \Box\psi \in T$. Then, by Propositions 2- $\langle \vec{\sigma} \rangle$ and Proposition 1.6, we obtain that (1) $\bigwedge \vec{\sigma} \in T$, i.e., $T \in \widehat{\vec{\sigma}}$, and (2) $\Box[\vec{\sigma}]\psi \in T$. Thus, by Lemma 9 and (2), we have $[\vec{u}][\vec{\sigma}]\psi \in T$, i.e., $[\vec{u}, \vec{\sigma}]\psi \in T$, for all $\vec{u} \in \text{Prop}_O^*$. Now let $O \in \mathcal{O}^c$ such that $T \in O$. By the construction of \mathcal{O}^c , we know that $O = \widehat{\vec{v}}$ for some $\vec{v} \in \text{Prop}_O^*$. We want to show that $(T, \widehat{\vec{\sigma}} \cap \widehat{\vec{v}}) \models \psi$. Since $T \in \widehat{\vec{\sigma}} \cap \widehat{\vec{v}}$ and \mathcal{M}^c is an intersection space, we know that $(T, \widehat{\vec{\sigma}} \cap \widehat{\vec{v}})$ is a well-defined epistemic scenario. $T \in \widehat{\vec{\sigma}} \cap \widehat{\vec{v}}$ also implies that $(\bigwedge \vec{\sigma} \wedge \bigwedge \vec{v}) \in T$ as in the above case. Hence, by Proposition 2- $\langle \vec{\sigma} \rangle$ and the fact that $[\vec{v}, \vec{\sigma}]\psi \in T$, we obtain $\langle \vec{v}, \vec{\sigma} \rangle \psi \in T$. Then, by I.H, we obtain $(T, \widehat{\vec{v}} \cap \widehat{\vec{\sigma}}) \models \psi$ as in the previous case. Therefore, by the semantics of \Box , we obtain $(T, \widehat{\vec{\sigma}}) \models \Box\psi$.

(\Rightarrow) Suppose $(T, \widehat{\vec{\sigma}}) \models \Box\psi$. This means, by the definition of \mathcal{O}^c , that for all $\vec{u} \in \text{Prop}_O^*$, if $T \in \widehat{\vec{u}}$ then $(T, \widehat{\vec{\sigma}} \cap \widehat{\vec{u}}) \models \psi$. Now let $\vec{v} \in \text{Prop}_O^*$ such that $T \in \widehat{\vec{v}}$. Therefore, $T \in \widehat{\vec{v}} \cap \widehat{\vec{\sigma}}$. Since $(\vec{v}, \vec{\sigma}) \in \text{Prop}_O^*$ and $\widehat{\vec{v}} \cap \widehat{\vec{\sigma}} = \widehat{(\vec{v}, \vec{\sigma})}$, we obtain by the assumption that $(T, \widehat{\vec{v}} \cap \widehat{\vec{\sigma}}) \models \psi$. Thus, by I.H., we have $\langle \vec{v}, \vec{\sigma} \rangle \psi \in T$. As $\vdash \langle \vec{v}, \vec{\sigma} \rangle \psi \rightarrow [\vec{v}, \vec{\sigma}]\psi$ and \mathbf{T} is maximal, we obtain $[\vec{v}, \vec{\sigma}]\psi \in T$, i.e., $[\vec{v}][\vec{\sigma}]\psi \in T$. Hence, by Lemma 9, we have $\Box[\vec{\sigma}]\psi \in T$. Then, by Proposition 1.6, the fact that $\bigwedge \vec{\sigma} \in T$ and Propositions 2- $\langle \vec{\sigma} \rangle$ and, we obtain $\langle \vec{\sigma} \rangle \Box\psi \in T$.

Theorem 2. \mathbf{L} is complete with respect to the class of all learning models.

Proof. Let φ be an \mathbf{L} -consistent formula, i.e., it is an \mathbf{O}_φ -theory. Then, by Lemma 6, it can be extended to a maximal \mathbf{O}_φ -witnessed theory T . Then, we have $\langle \lambda \rangle \varphi \in T$ where λ is the empty string, i.e., $T \in \widehat{\langle \lambda \rangle \varphi}$. Note that $\widehat{\langle \lambda \rangle \varphi} = \bigcap \emptyset = X^c$. Then, by Truth Lemma (Lemma 10), we obtain that $(T, X^c) \models_{\mathcal{M}^c} \varphi$, where $\mathcal{M}^c = (X^c, \mathcal{O}^c, \mathbb{L}^c, \|\cdot\|_c)$ is the canonical model for T . This proves completeness.

4 Expressivity

We first investigate how various notions of learnability can be expressed in our language. In fact, the following result was already noticed in [17]:

Proposition 6. $\diamond Kp$ is true at (x, U) in a model \mathcal{M} iff $\|p\|$ is learnable with certainty at state x . Similarly, $p \rightarrow \diamond Kp$ is valid (i.e. true at all epistemic scenarios) in a model \mathcal{M} iff $\|p\|$ is verifiable with certainty (i.e. ‘finitely identifiable’ in the sense of FLT [10, 14]). A similar statement holds for falsifiability with certainty.

Proof. As we know from Section 1.1, $\|p\|$ is learnable with certainty x iff $x \in \text{Int}\|p\|$, and $\|p\|$ is verifiable with certainty iff it is open in the topology generated by \mathcal{O} . It is well-known [17] that these properties are expressible in SSL via the above validities.

In particular, the following validity of our logic expresses the fact that *all observable properties are verifiable with certainty*:

$$o \rightarrow \diamond Ko.$$

By adding the learning operator to subset space logic, DLLT can capture, not only belief, but also the various inductive notions of knowledge and learnability:

Proposition 7. [Inductive notions of knowledge and learnability]

- $\Box Bp$ holds at (x, U) in a model \mathcal{M} iff the learner \mathbb{L} has undefeated belief in $\|p\|$ (at world x in information state U). Hence, $p \wedge \Box Bp$ captures inductive knowledge of p , and so $p \wedge \diamond \Box Bp$ captures inductive learnability of p by learner \mathbb{L} .
- Similarly, $p \rightarrow \diamond \Box Bp$ is valid in a model \mathcal{M} iff $\|p\|$ is inductively verifiable by L . For the corresponding generic notion: $\|p\|$ is inductively verifiable (by some learner) iff $p \rightarrow \diamond \Box Bp$ is valid in the intersection space (X, \mathcal{O}) . Similar statements hold for inductive falsifiability.
- Finally, $\diamond L(\lambda)$ is true if (given enough observations) the observer will eventually reach a true conjecture (though he might later fall again into false ones); and similarly, $\diamond \Box L(\lambda)$ is true if (given enough observations) the observer will eventually produce only true conjectures thereafter.

Proof. This is an easy verification, given the relevant definitions and our semantics.

As usual in Dynamic Epistemic Logic, the dynamic ‘observation’ modalities $[u]\varphi$ are only a convenient way to express complex properties in a succinct manner, but they can in principle be eliminated. To show this, we first need the following lemma.

Lemma 11. There is a well-founded strict partial order $<$ on formulas (called ‘complexity order’), satisfying the following conditions:

- if φ is a (proper) subformula of ψ then $\varphi < \psi$
- $(u \rightarrow p) < [u]p$
- $(u \rightarrow o) < [u]o$
- $L(u, \vec{o}) < [u]L(\vec{o})$

- $([u]\varphi \wedge [u]\psi) < [u](\varphi \wedge \psi)$
- $(u \rightarrow K[u]\varphi) < [u]K\varphi$
- $\Box[u]\varphi < [u]\Box\varphi$

Proposition 8. (*Expressivity*) *The above language is co-expressive with the one obtained by removing all dynamic modalities $[u]\varphi$. Moreover, this can be proved in the above proof system: for every formula φ there exists some formula φ' free of any dynamic modalities, such that $\varphi \leftrightarrow \varphi'$ is a theorem in the above proof system. Furthermore, if φ contains dynamic modalities then φ' can be chosen such that $\varphi' < \varphi$.*

5 Conclusion and Comparison with Other Work

In this paper we proposed a dynamic logic which allows reasoning about inductive inference. Our Dynamic Logic of Learning Theory (DLLT) is an extension of previously studied Subset Space Logics, and a natural continuation of the work bridging Dynamic Epistemic Logic and Formal Learning Theory. Together with a syntax, featuring dynamic observation operators, and a topological semantics, we give a sound and complete axiomatization of this logic. We show how natural learnability properties, as learnability in the limit and learnability with certainty, can be expressed in DLLT.

Our technical results (the complete axiomatization and expressivity results), as well as the methods used to prove them (the canonical neighborhood model and the reduction laws), may look deceptively simple. But in fact, achieving this simplicity is one of the major contributions of our paper! The most well-known relative to our logic is *Subset Space Logic (SSL) over intersection spaces*, completely axiomatized by Weiss and Parikh [26] (- and indeed our operator \Box originates in the ‘effort modality’ of the SSL formalism introduced in [17, 9]). Although less expressive than our logic (since it has no notion of belief B or conjecture L), the Weiss-Parikh axiomatization of SSL over intersection spaces is in a sense more complex and less transparent (as is their completeness proof, which is non-canonical). That axiomatization consists of the following list:

$S5_K$	The $S5$ axioms and rules for K
$S4_\Box$	The $S4$ axioms and rules for \Box
Cross Axiom	$K\Box\varphi \rightarrow \Box K\varphi$
Weak Directedness	$\Diamond\Box\varphi \rightarrow \Box\Diamond\varphi$
M_n (for all n)	$(\Box\langle K \rangle\varphi \wedge \Diamond K\psi_1 \wedge \dots \wedge \Diamond K\psi_n) \rightarrow \langle K \rangle(\Diamond\varphi \wedge \Diamond K\psi_1 \wedge \dots \wedge \Diamond K\psi_n)$

Though this list looks shorter than our list in Table 1, each of our axioms is simple and readable and has a transparent intuitive interpretation. In contrast, note the complexity and opaqueness of the last axiom schemata M_n above (having one schema for each natural number n)! Our completeness result implies that all these complex validities are provable in our simple system (and in fact in the even simpler system obtained by deleting from ours all the axioms that refer to the learner L). This shows the usefulness of adding the (expressively redundant) dynamic observation modalities: they help to describe the behavior of the effort modality \Box in a much simpler and natural manner, via the combination of the Effort axiom and the Effort rule (which together capture the meaning of \Box as universally quantifying over observation modalities).

Moreover, our completeness proof is also much simpler (though with some technical twists). Traditionally, the use of canonical models has been considered impossible for

Subset Space Logics, and so authors had to use other, more ad-hoc methods (e.g. step-by-step constructions). The fact that in this paper we can get away with a canonical construction is again due to the addition of the dynamic modalities.

More recent papers, closely related to our logic, are Bjorndahl [8], van Ditmarsch et alia [20, 21], and Baltag et alia [7]. Bjorndahl [8] introduces dynamic modalities $[\varphi]$ for arbitrary formulas (rather than restricting to observational variables $[o]$, as we do), though with a different semantics (according to which $[\varphi]$ restricts the space to the interior of φ , in contrast to our simpler semantics, that follows the standard definition of update or “public announcement”). His syntax does *not* contain the effort modality, or any other form of quantifying over observations. The work of van Ditmarsch et alia [20, 21] uses Bjorndahl-style dynamic modalities in combination with a topological version of the so-called “arbitrary public announcement” operator, which is a more syntactic-driven relative of the effort modality. This syntactic nature comes with a price: the logic of arbitrary public announcements is much less well-behaved than SSL (or our logic), in particular it has non-compositional features (-the meaning of a formula may depend on the meaning of *all* atomic variables, including the ones that do not occur in that formula!). As a consequence, the soundness of (the arbitrary-public announcement analogue of) our Effort Rule is not at all obvious for their logic, which instead relies on an infinitary inference rule. Since that rule makes use of infinitely many premisses, their complete axiomatization is truly infinitary, and impossible to automatize: indeed, it does not even necessarily imply that the set of their validities is recursively enumerable (in contrast with our finitary axiomatization, which immediately implies such a result). The recent, unpublished work by Baltag et alia [7] (due to a subset of the current authors, using techniques similar to the ones we used in this paper) fixes these problems by replacing the arbitrary announcement modality with the effort modality (or equivalently, extending SSL with Bjorndahl-style dynamic modalities). But note that, in contrast to the work presented here, *all the above papers are concerned only with axiomatizations over topological spaces* (rather than the wider class of intersection spaces), and that *none of them has any belief B or conjecture operators L* . Hence, none of them can be used to capture any learning-theoretic notions going beyond finite identifiability.

References

1. Aiello, M., Pratt-Hartmann, I., van Benthem, J., et al. (2007) *Handbook of spatial logics*, vol. 4. Springer.
2. Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., and Lima, T. D. (2008) ‘Knowable’ as ‘Known after an announcement’. *The Review of Symbolic Logic*, **1**, 305–334.
3. Balbiani, P., van Ditmarsch, H., and Kudinov, A. (2013) Subset space logic with arbitrary announcements. *Proc. of the 5th ICLA*, pp. 233–244, Springer.
4. Baltag, A., Gierasimczuk, N., and Smets, S. (2011) Belief revision as a truth-tracking process. *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 187–190, ACM.
5. Baltag, A., Gierasimczuk, N., and Smets, S. (2015) On the solvability of inductive problems: A study in epistemic topology. Ramanujam, R. (ed.), *Proceedings of the 15th conference TARK*, also available as a technical report in ILLC Prepublication Series PP-2015-13.

6. Baltag, A., Moss, L. S., and Solecki, S. (1998) The logic of public announcements, common knowledge, and private suspicions. *Proceedings of the 7th Conference TARK*, pp. 43–56, Morgan Kaufmann Publishers Inc.
7. Baltag, A., Özgün, A., and Vargas-Sandoval, A. L. (2017) Topo-logic as a dynamic-epistemic logic. *Proceedings of the 6th International Workshop on Logic, Rationality and Interaction (LORI 2017)*, to appear.
8. Bjorndahl, A. (2016) Topological subset space models for public announcements. To appear in *Trends in Logic, Outstanding Contributions: Jaakko Hintikka*.
9. Dabrowski, A., Moss, L. S., and Parikh, R. (1996) Topological reasoning and the logic of knowledge. *Annals of Pure and Applied Logic*, **78**, 73–110.
10. Dégrémont, C. and Gierasimczuk, N. (2011) Finite identification from the viewpoint of epistemic update. *Information and Computation*, **209**, 383–396.
11. Georgatos, K. (1994) Knowledge theoretic properties of topological spaces. Masuch, M. and Pólos, L. (eds.), *Knowledge Representation and Reasoning Under Uncertainty: Logic at Work*, vol. 808 of *Lecture Notes in Artificial Intelligence*, pp. 147–159, Springer.
12. Georgatos, K. (1997) Knowledge on treelike spaces. *Studia Logica*, **59**, 271–301.
13. Gierasimczuk, N. (2010) *Knowing One's Limits. Logical Analysis of Inductive Inference*. Ph.D. thesis, Universiteit van Amsterdam, The Netherlands.
14. Gierasimczuk, N. and de Jongh, D. (2013) On the complexity of conclusive update. *The Computer Journal*, **56**, 365–377.
15. Gold, E. M. (1967) Language identification in the limit. *Information and control*, **10**, 447–474.
16. Kelly, K. T. (1996) *The Logic of Reliable Inquiry*. Oxford University Press.
17. Moss, L. S. and Parikh, R. (1992) Topological reasoning and the logic of knowledge. *Proceedings of the 4th conference TARK*, pp. 95–105, Morgan Kaufmann.
18. Parikh, R., Moss, L. S., and Steinsvold, C. (2007) Topology and epistemic logic. *Handbook of Spatial Logics*, pp. 299–341.
19. van Benthem, J. (2011) *Logical Dynamics of Information and Interaction*. Cambridge University Press, New York, NY, USA.
20. van Ditmarsch, H., Knight, S., and Özgün, A. (2014) Arbitrary announcements on topological subset spaces. *Proceedings of the 12th European Conference on Multi-Agent Systems (EUMAS)*, pp. 252–266, Springer.
21. van Ditmarsch, H., Knight, S., and Özgün, A. (2015) Announcement as effort on topological spaces. *Proc. of the 15th TARK*, pp. 95–102.
22. van Ditmarsch, H., Knight, S., and Özgün, A. (2015) Announcement as effort on topological spaces. Extended version, Submitted. <http://www.lix.polytechnique.fr/sophia/papers/effort.pdf>.
23. van Ditmarsch, H., van der Hoek, W., and Kooi, B. (2007) *Dynamic Epistemic Logic*. Springer Publishing Company, 1st edn.
24. Vickers, S. (1989) *Topology via logic*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press.
25. Wang, Y. N. and Ågotnes, T. (2013) Multi-agent subset space logic. *Proceedings of the 23rd IJCAI*, pp. 1155–1161, IJCAI/AAAI.
26. Weiss, M. A. and Parikh, R. (2002) Completeness of certain bimodal logics for subset spaces. *Studia Logica*, **71**, 1–30.

6 Proofs of Section 2.2: Axiomatization and proof system

Proof (Proposition 1. Reduction laws for strings of observational variables). The proofs follow directly from instances of the corresponding reduction axiom for one observational variable, by $(K_{[o]})$ and the definition $[\vec{o}] := [o_1][o_2] \dots [o_m]$. To illustrate the proof, we provide the one corresponding to the base case: $[\vec{u}]p \leftrightarrow (\bigwedge \vec{u} \rightarrow p)$. Recall that $\vec{u} := u_k, \dots, u_2, u_1$. W.l.o.g suppose $k = 2$.

From left to right:

$$\begin{aligned}
\vdash [u_2, u_1]p &\rightarrow [u_2][u_1]p && \text{(by definition of } [\vec{u}]) \\
\vdash [u_2, u_1]p &\rightarrow [u_2](u_1 \rightarrow p) && \text{(by reducc.ax } (R_p)) \\
\vdash [u_2, u_1]p &\rightarrow ([u_2]u_1 \rightarrow [u_2]p) && \text{(by } K_{[o]} \text{ axiom)} \\
\vdash [u_2, u_1]p &\rightarrow ((u_2 \rightarrow u_1) \rightarrow (u_2 \rightarrow p)) && \text{(by reducc.ax } (R_o) \text{ and } (R_p)) \\
\vdash [u_2, u_1]p &\rightarrow (((u_2 \rightarrow u_1) \wedge u_2) \rightarrow p) && \text{(by CPL)} \\
\vdash [u_2, u_1]p &\rightarrow ((u_1 \wedge u_2) \rightarrow p) && \text{(by CPL)}
\end{aligned}$$

From right to left:

$$\begin{aligned}
\vdash ((u_2 \wedge u_1) \rightarrow p) &\rightarrow (u_2 \rightarrow (u_1 \rightarrow p)) && \text{(tautology in CPL)} \\
\vdash [u_2](((u_2 \wedge u_1) \rightarrow p) \rightarrow (u_2 \rightarrow (u_1 \rightarrow p))) &&& \text{(by } Nec_{[o]}) \\
\vdash [u_2]((u_2 \wedge u_1) \rightarrow p) &\rightarrow [u_2](u_2 \rightarrow (u_1 \rightarrow p)) && \text{(by } K_{[o]} \text{ and MP)} \\
\vdash ([u_2](u_2 \wedge u_1) \rightarrow [u_2]p) &\rightarrow ([u_2]u_2 \rightarrow [u_2][u_1]p) && \text{(by } K_{[o]} \text{ and reducc.ax } (R_p)) \\
\vdash ([u_2]u_2 \wedge [u_2]u_1 \rightarrow (u_2 \rightarrow p)) &\rightarrow ((u_2 \rightarrow u_2) \rightarrow [u_2][u_1]p) && \text{(by } R_\wedge, \text{ the reducc.ax } (R_p) \text{ and } (R_o)) \\
\vdash ([u_2]u_2 \wedge [u_2]u_1 \wedge u_2 \rightarrow p) &\rightarrow ([u_2][u_1]p) && \text{(by CPL)} \\
\vdash ((u_2 \wedge (u_2 \rightarrow u_1) \wedge u_2) \rightarrow p) &\rightarrow [u_2][u_1]p && \text{(by CPL)} \\
\vdash ((u_2 \wedge u_1) \rightarrow p) &\rightarrow [u_2][u_1]p && \text{(by CPL)}
\end{aligned}$$

We use the axiom (R_\wedge) that appears in the proposition that follows (in Proposition 2). Although the proof for (R_\wedge) is straightforward and independent of the axioms stated in Proposition 1.

Proof (Proposition 2). The proofs go straightforwardly, (R_\wedge) follows from the axioms $(K_{[o]})$ and $(Nec_{[o]})$ and $(R_{\rightarrow \wedge})$ follows from (R_\wedge) and the definition $[\vec{o}] := [o_1][o_2] \dots [o_m]$; the reduction law $(\langle o \rangle)$ follows from the definition $\langle o \rangle \psi := \neg[o]\neg\psi$ and the reduction axiom (R_\neg) . The proof of $(\langle \vec{o} \rangle)$ follows directly from $(\langle o \rangle)$ and the definitions $\langle o \rangle \psi := \neg[o]\neg\psi$ and $[\vec{o}] := [o_1][o_2] \dots [o_m]$.

Proof (Proposition 3. The wKD45 axioms and rules for Conditional Belief). Let $\vec{o} \in Prop_{\vec{o}}^*$. We prove Necessitation Rule for $B^{\vec{o}}$ and only the axioms K, D' and 4 since 5 follows similarly as 4.

- Necessitation Rule for $B^{\vec{\sigma}}$: From $\vdash \varphi$ infer $\vdash B^{\vec{\sigma}}\varphi$.

$$\begin{array}{ll}
\vdash \varphi & \text{(by our ininitial assumption, } \varphi \text{ is a theorem)} \\
\vdash L(\vec{\sigma}) \rightarrow \varphi & \text{(by CPL)} \\
\vdash K(L(\vec{\sigma}) \rightarrow \varphi) & \text{(by rule (Nec}_K \text{))} \\
\vdash B^{\vec{\sigma}}\varphi &
\end{array}$$

- Axiom K for $B^{\vec{\sigma}}$: For this we need to show that $\vdash K(L(\vec{\sigma}) \rightarrow (\varphi \rightarrow \psi)) \rightarrow K(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow K(L(\vec{\sigma}) \rightarrow \psi)$:

$$\begin{array}{ll}
\vdash K(L(\vec{\sigma}) \rightarrow (\varphi \rightarrow \psi)) & \\
\vdash K((L(\vec{\sigma}) \rightarrow \varphi) \rightarrow (L(\vec{\sigma}) \rightarrow \psi)) & \text{(by CPL)} \\
\vdash K(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow K(L(\vec{\sigma}) \rightarrow \psi) & \text{(by (K}_K \text{))}
\end{array}$$

- Axiom D' for $B^{\vec{\sigma}}$: We need to prove $\vdash \langle K \rangle \vec{\sigma} \rightarrow \neg B^{\vec{\sigma}} \perp$, i.e., to prove $\vdash \langle K \rangle \vec{\sigma} \rightarrow \neg K(\neg L(\vec{\sigma}) \vee \perp)$. Moreover, It is easy to see that the consequent of this implication is equivalent to $\langle K \rangle L(\vec{\sigma})$. Thus it suffices to be proved $\vdash \langle K \rangle \vec{\sigma} \rightarrow \langle K \rangle L(\vec{\sigma})$:

$$\begin{array}{ll}
\vdash L(\vec{\sigma}) \rightarrow \bigwedge \vec{\sigma} & \text{(is (SP) axiom)} \\
\vdash K(L(\vec{\sigma}) \rightarrow \bigwedge \vec{\sigma}) & \text{(by (Nec}_K \text{) rule)} \\
\vdash K(L(\vec{\sigma})) \rightarrow K(\bigwedge \vec{\sigma}) & \text{(by (K}_K \text{) axiom and MP with previous step)} \\
\vdash \langle K \rangle \vec{\sigma} \rightarrow \langle K \rangle L(\vec{\sigma}) & \text{(the dual of the formula in previous step)}
\end{array}$$

- Axiom 4 for $B^{\vec{\sigma}}$: For this we need to show that $\vdash K(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow K(L(\vec{\sigma}) \rightarrow B^{\vec{\sigma}}\varphi)$. It is easy to see that this is equivalent to showing that $\vdash K(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow K(L(\vec{\sigma}) \rightarrow K(K(L(\vec{\sigma})) \rightarrow K\varphi)$:

$$\begin{array}{ll}
1) \vdash K(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow KK(L(\vec{\sigma}) \rightarrow \varphi) & \text{(by axiom (4}_K \text{))} \\
2) \vdash KK(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow K(K(L(\vec{\sigma})) \rightarrow K\varphi) & \text{(by axiom (K}_K \text{))} \\
3) \vdash KK(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow (K(L(\vec{\sigma})) \rightarrow K(K(L(\vec{\sigma})) \rightarrow K\varphi)) & \text{(by CPL)} \\
4) \vdash K(L(\vec{\sigma}) \rightarrow \varphi) \rightarrow (K(L(\vec{\sigma})) \rightarrow K(K(L(\vec{\sigma})) \rightarrow K\varphi)) & \text{(by 1,3 and CPL)}
\end{array}$$

Proposition 9. *The S4 axioms for the effort modality \square are derivable in \mathbf{L} .*

Proof (Proposition 9. The S4 laws for \square). The proofs of the S4 laws for \square are very straightforward simply by using classical propositional logic (CPL), the axiom (\square -Ax) and/or (\square -Rule). The most engaged proof is the one for Kripke's axiom (\mathbf{K}_{\square}) for the effort modality \square , so we show it here. By CPL, it is equivalent to showing $(\square(\varphi \rightarrow \psi) \wedge \square\varphi) \rightarrow \square\psi$, which proceeds as follows:

$$1) \vdash (\square(\varphi \rightarrow \psi) \wedge \square\varphi) \rightarrow ([o](\varphi \rightarrow \psi) \wedge [o]\varphi) \quad \text{(by } (\square\text{-Ax}) \text{ for some } o \notin O_{\varphi} \cup O_{\psi})$$

- 2) $\vdash ([o](\varphi \rightarrow \psi) \wedge [o]\varphi) \rightarrow ([o]\varphi \rightarrow [o]\psi) \wedge [o]\varphi$ (by axiom $(K_{[o]})$)
- 3) $\vdash (([o]\varphi \rightarrow [o]\psi) \wedge [o]\varphi) \rightarrow [o]\psi$ (by CPL)
- 4) $\vdash ([o](\varphi \rightarrow \psi) \wedge [o]\varphi) \rightarrow [o]\psi$ (by CPL and steps 2 and 3)
- 5) $\vdash (\Box(\varphi \rightarrow \psi) \wedge \Box\varphi) \rightarrow [o]\psi$ (by CPL and step 1)
- 6) $\vdash (\Box(\varphi \rightarrow \psi) \wedge \Box\varphi) \rightarrow \Box\psi$ (by $o \notin O_\varphi \cup O_\psi$ and $(\Box\text{-Rule})$)

7 Proofs of Section 3: Soundness and Completeness

7.1 Proof of Lemma 1

The proof follows by subformula induction on φ .

Let $\mathcal{M} = (X, \mathcal{O}, \mathbb{L}, \|\cdot\|)$ and $\mathcal{M}' = (X, \mathcal{O}, \mathbb{L}, \|\cdot\|')$ be two learning models such that \mathcal{M} and \mathcal{M}' differ only in the valuation of some $o \notin O_\varphi$, and let $U \in \mathcal{O}$.

– The base cases:

Case $\varphi := q$ where $q \in \text{Prop}$ follows simply because q is not an observational variable, thus for all $q \in P_\varphi$ we have that $\|q\|_{\mathcal{M}} = \|q\|_{\mathcal{M}'}$. Since \mathcal{M} and \mathcal{M}' have the same set of opens \mathcal{O} , for all $U \in \mathcal{O}$ we have that $\llbracket q \rrbracket_{\mathcal{M}}^U = U \cap \|q\|_{\mathcal{M}} = U \cap \|q\|_{\mathcal{M}'} = \llbracket q \rrbracket_{\mathcal{M}'}^U$.

Case $\varphi := v$ where $v \in \text{Prop}_\varphi$, follows simply because since $v \in O_\varphi$ we have that $\|v\|_{\mathcal{M}} = \|v\|_{\mathcal{M}'}$. By a similar reasoning as before, $\llbracket v \rrbracket_{\mathcal{M}}^U = \llbracket v \rrbracket_{\mathcal{M}'}^U$.

– Case $\varphi := L(\vec{v})$

Note that $O_{L(\vec{v})} = O_{\vec{v}} = \{v_1, \dots, v_n\}$ and $\|v_i\|_{\mathcal{M}} = \|v_i\|_{\mathcal{M}'}$ for every $1 \leq i \leq n$. Thus $\llbracket \vec{v} \rrbracket_{\mathcal{M}}^U = \llbracket \vec{v} \rrbracket_{\mathcal{M}'}^U$ for every $U \in \mathcal{O}$. Therefore, since \mathbb{L} and \mathcal{O} are the same in both models we have that $\llbracket L(\vec{v}) \rrbracket_{\mathcal{M}}^U = \llbracket L(\vec{v}) \rrbracket_{\mathcal{M}'}^U$.

– Booleans: The cases for Booleans are straightforward.

– Case $\varphi := \neg\psi$. It follows straightforwardly from $O_{\neg\psi} = O_\psi$ and simple algebra of sets.

– Case $\varphi := K\psi$

Note that $O_{K\psi} = O_\psi$. Then, by induction hypothesis (IH), we have that $\llbracket \psi \rrbracket_{\mathcal{M}}^U = \llbracket \psi \rrbracket_{\mathcal{M}'}^U$. We have two case (1) if $U = \llbracket \psi \rrbracket_{\mathcal{M}}^U = \llbracket \psi \rrbracket_{\mathcal{M}'}^U$, then $\llbracket K\psi \rrbracket_{\mathcal{M}}^U = \llbracket K\psi \rrbracket_{\mathcal{M}'}^U = U$, and (2) if $\llbracket \psi \rrbracket_{\mathcal{M}}^U = \llbracket \psi \rrbracket_{\mathcal{M}'}^U \neq U$, then we have $\llbracket K\psi \rrbracket_{\mathcal{M}}^U = \llbracket K\psi \rrbracket_{\mathcal{M}'}^U = \emptyset$.

– Case $\varphi := [v]\psi$

Note that $O_{[v]\psi} = \{v\} \cup O_\psi$. Suppose $x \in \llbracket [v]\psi \rrbracket_{\mathcal{M}}^U$ and $x \in \|v\|_{\mathcal{M}'}$. By IH on v , we have $\|v\|_{\mathcal{M}} = \|v\|_{\mathcal{M}'}$. Therefore, from $x \in \llbracket [v]\psi \rrbracket_{\mathcal{M}}^U$ and $x \in \|v\|_{\mathcal{M}}$, we have $(x, U \cap \|v\|_{\mathcal{M}}) \models_{\mathcal{M}} \psi$. Since $U \cap \|v\|_{\mathcal{M}} = U \cap \|v\|_{\mathcal{M}'}$ and $U \cap \|v\|_{\mathcal{M}'} \in \mathcal{O}$, by IH on ψ , we obtain $(x, U \cap \|v\|_{\mathcal{M}'}) \models_{\mathcal{M}'} \psi$. We then conclude $x \in \llbracket [v]\psi \rrbracket_{\mathcal{M}'}^U$. The other direction follows similarly.

– Case $\varphi := \Box\psi$

Suppose $x \in \llbracket \Box\psi \rrbracket_{\mathcal{M}}^U$. This means, by the semantics of \Box , that for all $O \in \mathcal{O}$ with $x \in O \subseteq U$ we have that $(x, O) \models_{\mathcal{M}} \psi$, i.e., that $x \in \llbracket \psi \rrbracket_{\mathcal{M}}^O$. Therefore, by IH and the fact that $O_{\Box\psi} = O_\psi$, we obtain $x \in \llbracket \psi \rrbracket_{\mathcal{M}'}^O$ for all $O \subseteq U$. Since \mathcal{M} and \mathcal{M}' carry exactly the same collection \mathcal{O} , we conclude that $x \in \llbracket \Box\psi \rrbracket_{\mathcal{M}'}^U$. The opposite direction follows similarly.

7.2 Proof of Theorem 1

The soundness proof follows simply by validity check. We here only present the validity proofs for the Learning axioms (all), the Reduction axioms R_L , R_{\square} , as well as the Effort-axiom (\square -Ax) and Effort-rule (\square -Rule). The other cases follow standardly.

Let $\mathcal{M} = (X, \mathcal{O}, \mathbb{L}, \|\cdot\|)$ be a learning model and $(x, U) \in ES(\mathcal{M})$.

– Learning axioms:

- (CC): Suppose $(x, U) \models \bigwedge \vec{\sigma}$. This means that $x \in \llbracket \bigwedge \vec{\sigma} \rrbracket^U$. Thus $x \in \llbracket \vec{\sigma} \rrbracket^U$. We want to show that $(x, U) \models \langle K \rangle L(\vec{\sigma})$, i.e., to show that $\exists y \in U$ s.t. $(y, U) \models L(\vec{\sigma})$, i.e., by the semantic definition of L , $\exists y \in U$ s.t. $y \in \mathbb{L}(U, \llbracket \vec{\sigma} \rrbracket)$. Recall that $\mathbb{L}(\vec{O}) := \mathbb{L}(\bigcap \vec{O})$, $\mathbb{L}(O) \subseteq O$ for every $O \in \mathcal{O}$ and $\mathbb{L}(O) \neq \emptyset$ if $O \neq \emptyset$. Since $x \in \llbracket o_i \rrbracket^U$ for every $o_i \in \vec{\sigma}$, we have that $\llbracket \vec{\sigma} \rrbracket^U \neq \emptyset$ thus $\mathbb{L}(\llbracket \vec{\sigma} \rrbracket^U) \neq \emptyset$. Therefore $\mathbb{L}(U, \llbracket \vec{\sigma} \rrbracket) \neq \emptyset$, so there is $y \in \mathbb{L}(U, \llbracket \vec{\sigma} \rrbracket)$. Then, by the semantics of L , we obtain that $(y, U) \models L(\vec{\sigma})$. Thus $(x, U) \models \langle K \rangle L(\vec{\sigma})$.
- (EC): Suppose $(x, U) \models \vec{\sigma} \Leftrightarrow \vec{\sigma}'$. We need to show that $(x, U) \models L(\vec{\sigma}) \leftrightarrow L(\vec{\sigma}')$. It suffices to show $(x, U) \models L(\vec{\sigma}) \rightarrow L(\vec{\sigma}')$, since the other direction of the bi-conditional is analogous. Suppose $(x, U) \models L(\vec{\sigma})$. Since $(x, U) \models \vec{\sigma} \Leftrightarrow \vec{\sigma}'$ we have that $x \in \llbracket K(\bigwedge \vec{\sigma} \leftrightarrow \bigwedge \vec{\sigma}') \rrbracket^U$. Therefore, by the semantics of K , we obtain that for all $y \in U$, it is the case that $y \in \llbracket (\bigwedge \vec{\sigma} \leftrightarrow \bigwedge \vec{\sigma}') \rrbracket^U$. This means that for all $y \in U$, $y \in \llbracket \bigwedge \vec{\sigma} \rrbracket^U$ iff $y \in \llbracket \bigwedge \vec{\sigma}' \rrbracket^U$. Since $\mathbb{L}(\vec{O}) := \mathbb{L}(\bigcap \vec{O})$, by our assumption we have that $x \in \mathbb{L}(U, \llbracket \vec{\sigma} \rrbracket) = \mathbb{L}(U, \llbracket \bigwedge \vec{\sigma} \rrbracket) = \mathbb{L}(U, \llbracket \bigwedge \vec{\sigma}' \rrbracket) = \mathbb{L}(U, \llbracket \vec{\sigma}' \rrbracket)$. By the semantic definition of L , we have $(x, U) \models L(\vec{\sigma}')$.
- (SP): Suppose $(x, U) \models L(\vec{\sigma})$. By semantic definition and the fact that $\mathbb{L}(\vec{O}) := \mathbb{L}(\bigcap \vec{O})$, we have $x \in \mathbb{L}(U, \llbracket \vec{\sigma} \rrbracket) = \mathbb{L}(U \cap \llbracket \vec{\sigma} \rrbracket) = \mathbb{L}(U \cap \llbracket \bigwedge \vec{\sigma} \rrbracket)$. Since we have that $\mathbb{L}(O) \subseteq O$ from Definition 2, then $x \in \mathbb{L}(U \cap \llbracket \bigwedge \vec{\sigma} \rrbracket) \subseteq U \cap \llbracket \bigwedge \vec{\sigma} \rrbracket$. Therefore $x \in \llbracket \bigwedge \vec{\sigma} \rrbracket^U$, it follows that $(x, U) \models \bigwedge \vec{\sigma}$.

– (R_L): We will only prove from left to right of the bi-conditional, the other direction is also straightforward by using the properties of \mathbb{L} and the semantic definition of L . Suppose $(x, U) \models [o']L(\vec{\sigma})$. This means that $x \in \llbracket o' \rrbracket$ implies $(x, U \cap \llbracket o' \rrbracket) \models L(\vec{\sigma})$, i.e., that $x \in \llbracket o' \rrbracket$ implies $x \in \mathbb{L}(U \cap \llbracket o' \rrbracket, \llbracket \vec{\sigma} \rrbracket) = \mathbb{L}(U, \llbracket o' \rrbracket, \llbracket \vec{\sigma} \rrbracket) = \mathbb{L}(U, \llbracket o', \vec{\sigma} \rrbracket)$ (where the last two equalities follows from $\mathbb{L}(\vec{O}) := \mathbb{L}(\bigcap \vec{O})$). Then, by the semantics of L , we have that $(x, U) \models o' \rightarrow L(o', \vec{\sigma})$.

– (R_{\square}):

$$\begin{aligned}
 & (x, U) \models [o]\square\varphi \\
 & \text{iff } x \in \llbracket o \rrbracket \text{ implies } (x, U \cap \llbracket o \rrbracket) \models \square\varphi && \text{(by the semantics of } [o]\text{)} \\
 & \text{iff } x \in \llbracket o \rrbracket \text{ implies } (\forall O \in \mathcal{O})(x \in O \text{ implies } (x, (U \cap \llbracket o \rrbracket) \cap O) \models \varphi) && \text{(by the semantics of } \square\text{)} \\
 & \text{iff } (\forall O \in \mathcal{O})(x \in \llbracket o \rrbracket \text{ and } x \in O \text{ implies } (x, (U \cap \llbracket o \rrbracket) \cap O) \models \varphi) \\
 & \text{iff } (\forall O \in \mathcal{O})(x \in O \text{ implies } (x \in \llbracket o \rrbracket \text{ implies } (x, (U \cap O) \cap \llbracket o \rrbracket) \models \varphi)) \\
 & \text{iff } (\forall O \in \mathcal{O})(x \in O \text{ implies } (x, U \cap O) \models [o]\varphi) && \text{(by the semantics of } [o]\text{)}
 \end{aligned}$$

iff $(x, U) \models \Box[o]\varphi$ (by the semantics of \Box)

- (\Box -Ax): Suppose $(x, U) \models \Box\varphi$. This means, by the semantics of \Box , that for all $O \in \mathcal{O}$ with $x \in O$ we have $(x, U \cap O) \models \varphi$. In particular, since $\|o\| \in \mathcal{O}$ for all $o \in \text{Prop}_{\mathcal{O}}$, we obtain that $x \in \|o\|$ implies $(x, U \cap \|o\|) \models \varphi$. Moreover, for every $\vec{o} \in \text{Prop}_{\mathcal{O}}^*$, $x \in \llbracket \bigwedge \vec{o} \rrbracket^U$ implies $(x, U \cap \llbracket \bigwedge \vec{o} \rrbracket) \models \varphi$. Therefore by the semantic definition we have $(x, U) \models \llbracket \vec{o} \rrbracket\varphi$ for every string \vec{o} of observational variables.
- (\Box -Rule): Suppose towards a contradiction that $\models \psi \rightarrow [o]\varphi$ and $\not\models \psi \rightarrow \Box\varphi$ where $o \notin O_\psi \cup O_\varphi$. The latter means that there's a learning model \mathcal{M} and (x, U) epistemic scenario such that $(x, U) \models_{\mathcal{M}} \psi \rightarrow [o]\varphi$ and $(x, U) \not\models_{\mathcal{M}} \psi \rightarrow \Box\varphi$, i.e., $(x, U) \models_{\mathcal{M}} \psi$ and $(x, U) \not\models_{\mathcal{M}} \Box\varphi$. Thus, $(x, U) \models_{\mathcal{M}} \psi$ and $(x, U) \models_{\mathcal{M}} \Diamond\neg\varphi$. By the semantics of \Diamond , $\exists U_0 \subseteq U$ s.t. $x \in U_0$ implies $(x, U_0) \models_{\mathcal{M}} \neg\varphi$. Now consider the model \mathcal{M}' with exactly the same frame as \mathcal{M} and the following valuation: $\|o\|_{\mathcal{M}'} = U_0$, $\|u\|_{\mathcal{M}'} = \|u\|_{\mathcal{M}}$ for all $u \neq o \in \text{Prop}_{\mathcal{O}}$ and $\|p\|_{\mathcal{M}'} = \|p\|_{\mathcal{M}}$ for all $p \in \text{Prop}$. By Lemma 1, we have that $\llbracket \psi \rrbracket_{\mathcal{M}'}^U = \llbracket \psi \rrbracket_{\mathcal{M}}^U$ and $\llbracket \neg\varphi \rrbracket_{\mathcal{M}'}^{U_0} = \llbracket \neg\varphi \rrbracket_{\mathcal{M}}^{U_0}$. Therefore $(x, U) \models_{\mathcal{M}'} \psi$ and $(x, U_0) \models_{\mathcal{M}'} \neg\varphi$. Since $\|o\|_{\mathcal{M}'} = U_0$, we have $(x, U) \models_{\mathcal{M}'} \psi$ and $(x, \|o\|_{\mathcal{M}'}) \models_{\mathcal{M}'} \neg\varphi$. By the semantics of $\langle o \rangle$ we obtain $(x, U) \models_{\mathcal{M}'} \psi \wedge \langle o \rangle\neg\varphi$. Thus $(x, U) \not\models_{\mathcal{M}'} \psi \rightarrow [o]\varphi$, contradicting the validity of $\models \psi \rightarrow [o]\varphi$.

7.3 Proof of Lemma 4

Let $\Gamma \subseteq \mathcal{Q}^O$ be an O-theory and $\varphi \in \mathcal{Q}^O$ such that $\Gamma \not\vdash \neg\varphi$. We first that $\Gamma \cup \{\varphi\}$ is an O-theory. Suppose by contradiction that $\Gamma \cup \{\varphi\}$ is not an O-theory, i.e., that $\Gamma \cup \{\varphi\} \vdash \perp$. Thus, there is a finite $\Delta \subseteq \Gamma$ such that $\Delta \vdash \neg\varphi$ and therefore $\Gamma \vdash \neg\varphi$, which contradicts the assumption that $\Gamma \not\vdash \neg\varphi$.

Now suppose that Γ is O-witnessed but $\Gamma \cup \{\varphi\}$ is not O-witnessed. By the previous statement, we know that $\Gamma \cup \{\varphi\}$ consistent. Therefore, the latter means that there is $s \in NF_{\mathcal{Q}}^O$ and $\psi \in \mathcal{Q}^O$ such that $\Gamma \cup \{\varphi\}$ is consistent with $\langle s \rangle \Diamond \psi$ but $\Gamma \cup \{\varphi\} \vdash \neg \langle s \rangle \langle o \rangle \psi$ for all $o \in \mathcal{O}$. This implies that $\Gamma \cup \{\varphi\} \vdash [s][o]\neg\psi$ for all $o \in \mathcal{O}$. Therefore, $\Gamma \vdash \varphi \rightarrow [s][o]\neg\psi$ for all $o \in \mathcal{O}$. Note that $\varphi \rightarrow [s][o]\neg\psi := [\varphi \rightarrow, s][o]\neg\psi$, and $[\varphi \rightarrow, s] \in NF_{\mathcal{Q}}^O$. We thus have $\Gamma \vdash [\varphi \rightarrow, s][o]\neg\psi$ for all $o \in \mathcal{O}$. Since Γ is O-witnessed, we obtain $\Gamma \vdash [\varphi \rightarrow, s]\Box\neg\psi$. By unraveling the necessity form $[\varphi \rightarrow, s]$, we get $\Gamma \vdash \varphi \rightarrow [s]\Box\neg\psi$, thus, $\Gamma \cup \{\varphi\} \vdash [s]\Box\neg\psi$, i.e., $\Gamma \cup \{\varphi\} \vdash \neg \langle s \rangle \Diamond \psi$, contradicting the assumption that $\Gamma \cup \{\varphi\}$ is consistent with $\langle s \rangle \Diamond \psi$.

7.4 Proof of Lemma 5: Lindenbaum's lemma

For the proof of Lemma 5 we need the following auxiliary lemmas.

Lemma 12. *For every maximal O-witnessed theory Γ , and any $\varphi, \psi \in \mathcal{Q}^O$,*

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 13. *If $\Gamma_0 \subseteq \Gamma_1 \subseteq \dots \subseteq \Gamma_n \subseteq \dots$ is an increasing chain of O-theories, then $\bigcup_{n \in \mathbb{N}} \Gamma_n$ is an O-theory.*

Proof. Let $\Gamma_0 \subseteq \Gamma_1 \subseteq \dots \subseteq \Gamma_n \subseteq \dots$ be an increasing chain of O-theories and suppose, toward contradiction, that $\bigcup_{n \in \mathbb{N}} \Gamma_n$ is not an O-theory, i.e., suppose that $\bigcup_{n \in \mathbb{N}} \Gamma_n \vdash \perp$. This means that there exists a finite $\Delta \subseteq \bigcup_{n \in \mathbb{N}} \Gamma_n$ such that $\Delta \vdash \perp$. Then, since $\bigcup_{n \in \mathbb{N}} \Gamma_n$ is a union of an increasing chain of O-theories, there is some $m \in \mathbb{N}$ such that $\Delta \subseteq \Gamma_m$. Therefore, $\Gamma_m \vdash \perp$ contradicting the fact that Γ_m is an O-theory.

Proof of Lemma 5 (Lindendaum's lemma): The proof follows by constructing an increasing chain of

$$\Gamma_0 \subseteq \Gamma_1 \subseteq \dots \subseteq \Gamma_n \subseteq \dots,$$

O-witnessed theories where $\Gamma_0 := \Gamma$, and each Γ_i will be recursively defined. We have to guarantee that each Γ_i is O-witnessed, in order to do so we follow a two-fold construction, where $\Gamma_0 = \Gamma'_0 := \Gamma$. Let $\gamma_n := (s_n, \varphi_n)$ be the n th-pair in the enumeration \mathcal{A} of all pairs of the form (s, φ) consisting of a possibility form s as described in Definition 4, and a formula φ , both in the language \mathcal{Q}^O . Note that all pairs of the form (λ, φ) , for every formula $\varphi \in \mathcal{Q}^O$ and the empty string λ of observational variables, are in \mathcal{A} . By the definition of necessity forms and possibility forms we have that every formula $\varphi \in \mathcal{Q}^O$ can be written as $\varphi := \langle \lambda \rangle \varphi = [\lambda] \varphi$. Therefore, all formulas in \mathcal{Q}^O are considered in the enumeration \mathcal{A} . We then set

$$\Gamma'_n = \begin{cases} \Gamma_n \cup \{ \langle s_n \rangle \varphi_n \} & \text{if } \Gamma_n \not\vdash \neg \langle s_n \rangle \varphi_n \\ \Gamma_n & \text{otherwise} \end{cases}$$

By Lemma 4, each Γ'_n is O-witnessed. Then, if φ_n is of the form $\varphi_n := \diamond \theta$ for some $\theta \in \mathcal{Q}^O$, there exists an $o \in \mathcal{O}$ such that Γ'_n is consistent with $\langle s \rangle \langle o \rangle \theta$ (since Γ'_n is O-witnessed). We then define

$$\Gamma_{n+1} = \begin{cases} \Gamma'_n & \text{if } \Gamma_n \not\vdash \neg \langle s_n \rangle \varphi_n \text{ and } \varphi_n \text{ is not of the form } \diamond \theta \\ \Gamma'_n \cup \{ \langle s_n \rangle \langle o \rangle \theta \} & \text{if } \Gamma_n \not\vdash \neg \langle s_n \rangle \varphi_n \text{ and } \varphi_n := \diamond \theta \text{ for some } \theta \in \mathcal{Q}^O \\ \Gamma_n & \text{otherwise} \end{cases}$$

where $o \in \mathcal{O}$ such that Γ'_n is consistent with $\langle s \rangle \langle o \rangle \theta$. Again by Lemma 4, it is guaranteed that each Γ_n is O-witnessed. Now consider the union $T_\Gamma = \bigcup_{n \in \mathbb{N}} \Gamma_n$. By Lemma 13, we know that T_Γ is an O-theory. To show that T_Γ is O-witnessed, let $s \in NF_{\mathcal{Q}}^O$ and $\theta \in \mathcal{Q}^O$ and suppose $\langle s \rangle \diamond \theta$ is consistent with T_Γ . The pair $(s, \diamond \theta)$ appears in the enumeration \mathcal{A} , thus $\gamma_m := (s_m, \varphi_m) = (s, \diamond \theta)$ with $s_m := s$ and $\varphi_m := \diamond \theta$, for some $\gamma_m \in \mathcal{A}$. Since $\langle s_m \rangle \varphi_m$ is consistent with T_Γ and $\Gamma_m \subseteq T_\Gamma$, we know that $\langle s_m \rangle \diamond \theta$ is in particular consistent with Γ_m . Therefore, by the above construction, $\langle s \rangle \langle o \rangle \theta \in \Gamma_{m+1}$ for some $o \in \mathcal{O}$ such that Γ'_m is consistent with $\langle s \rangle \langle o \rangle \theta$. Therefore, as T_Γ is consistent and $\Gamma_{m+1} \subseteq T_\Gamma$, we have that $\langle s \rangle \langle o \rangle \theta$ is also consistent with T_Γ . Hence, we conclude that T_Γ is O-witnessed. Finally, T_Γ is also maximal by construction: otherwise there would be an O-witnessed theory T such that $T_\Gamma \subset T$. This implies that there exists $\varphi \in \mathcal{Q}^O$ with $\varphi \in T$ but $\varphi \notin T_\Gamma$. Then, by the construction of T_Γ , we obtain $\Gamma_i \vdash \neg \langle \lambda \rangle \varphi$ for all $i \in \mathbb{N}$. Thus $\Gamma_i \vdash \neg \varphi$ for all $i \in \mathbb{N}$. Therefore, since $T_\Gamma \subseteq T$, we have $T \vdash \neg \varphi$. Hence, since $\varphi \in T$, we obtain $T \vdash \perp$ (contradicting T being consistent).

7.5 Proof of Lemma 6: Extension Lemma

Proof. Let $\mathcal{A} = \{\gamma_0, \gamma_1, \dots, \gamma_n, \dots\}$ be an enumeration of all pairs of the form $\gamma_i := (s_i, \varphi_i)$ consisting of any necessity form s_i as described in Definition 4, and every formula φ_i in the language $\mathfrak{L}^{\tilde{O}}$. We will recursively construct a chain of \tilde{O} -theories $\Gamma_0 \subseteq \dots \subseteq \Gamma_n \subseteq \dots$ such that:

1. $\Gamma_0 = \Gamma$,
2. $O'_n = \{o \in O' : o \text{ occurs in } \Gamma_n\}$ is finite for every $n \in \mathbb{N}$, and,
3. for every $\gamma_n := (s_n, \varphi_n)$ with $s_n \in NF_{\mathfrak{L}^{\tilde{O}}}^O$ and $\varphi_n \in \mathfrak{L}^O$:
 If $\Gamma_n \not\vdash \neg\langle s_n \rangle \diamond \varphi_n$ then there is o_m “fresh” such that $\langle s_n \rangle \langle o_m \rangle \varphi_n \in \Gamma_{n+1}$. Otherwise we will define $\Gamma_{n+1} = \Gamma_n$.

For every $\gamma_n \in \mathcal{A}$, let $O'(n) = \{o \in O' : o \text{ occurs either in } s_n \text{ or } \varphi_n\}$. Clearly every $O'(n)$ is always finite. We now construct an increasing chain of \tilde{O} -theories recursively: We fix $\Gamma_0 := \Gamma$ and let

$$\Gamma_{n+1} = \begin{cases} \Gamma_n \cup \{\langle s_n \rangle \langle o_m \rangle \varphi_n\} & \text{if } \Gamma_n \not\vdash \neg\langle s_n \rangle \diamond \varphi_n \\ \Gamma_n & \text{otherwise} \end{cases}$$

where m is the least natural number bigger than the indices in $O'_n \cup O'(n)$, i.e., o_m is fresh; We now need to show that $\tilde{\Gamma} := \bigcup_{n \in \mathbb{N}} \Gamma_n$ is an \tilde{O} -witnessed theory.

We first show that $\tilde{\Gamma}$ is an \tilde{O} -theory. By Lemma 13, it suffices to show by induction that every Γ_n is an \tilde{O} -theory. Clearly Γ_0 is an \tilde{O} -theory. For the inductive step suppose by contradiction that: $\Gamma_n \not\vdash \neg\langle s_n \rangle \diamond \varphi_n$, Γ_n is consistent but Γ_{n+1} is not. Hence $\Gamma \neq \Gamma_{n+1} \neq \Gamma_n$, and moreover $\Gamma_{n+1} \vdash \perp$. Then, since $\Gamma_{n+1} = \Gamma_n \cup \{\langle s_n \rangle \langle o_m \rangle \varphi_n\}$, we have $\Gamma_n \vdash [s_n][o_m] \neg \varphi_n$. Therefore there exists $\{\theta_1, \dots, \theta_k\} \subseteq \Gamma_n$ such that $\{\theta_1, \dots, \theta_k\} \vdash [s_n][o_m] \neg \varphi_n$. Let $\theta = \bigwedge_{1 \leq i \leq k} \theta_i$. Then $\vdash \theta \rightarrow [s_n][o_m] \neg \varphi_n$, so $\vdash [\theta \rightarrow, s_n][o_m] \neg \varphi_n$ with $o_m \notin O_{\Gamma_n} \cup O_{s_n} \cup O_{\varphi_n}$. Then, by the admissible rule in Lemma 3, we obtain the following,

$$\begin{aligned} & \vdash [\theta \rightarrow, s_n] \Box \neg \varphi_n && \text{(Lemma 3)} \\ & \vdash \theta \rightarrow [s_n] \Box \neg \varphi_n \\ & \theta \vdash [s_n] \Box \neg \varphi_n \\ & \theta \vdash \neg\langle s_n \rangle \diamond \varphi_n && \text{(Dual)} \end{aligned}$$

Since $\{\theta_1, \dots, \theta_k\} \subseteq \Gamma_n$, we therefore have $\Gamma_n \vdash \neg\langle s_n \rangle \diamond \varphi_n$ contradicting our assumption. Therefore Γ_{n+1} is consistent and thus an \tilde{O} -theory. Therefore, by Lemma 13, $\tilde{\Gamma}$ is an \tilde{O} -theory. Condition (3) above implies that $\tilde{\Gamma}$ is also \tilde{O} -witnessed. Then, by Lindenbaum’s Lemma, there is a maximal O -witnessed theory $T_{\tilde{\Gamma}}$ such that $T_{\tilde{\Gamma}} \supseteq \tilde{\Gamma} \supseteq \Gamma$.

7.6 Proofs of Lemmas 7, 8 and 9

Proof (Lemma 7). Observe that, by axiom (T_K) , $\{\theta : K\theta \in T\} \subseteq T$. Therefore, as T is consistent, the set $\{\theta : K\theta \in T\}$ is consistent. Let $s \in NF_{\mathfrak{L}^O}^O$ and $\varphi \in \mathfrak{L}^O$ such that

$\{\theta : K\theta \in T\} \vdash [s][o]\neg\varphi$ for all $o \in O$. Then, by normality of K , $T \vdash K[s][o]\neg\varphi$ for all $o \in O$. Since $K[s][o]\neg\varphi := [K, s][o]\neg\varphi$ is a necessity form and T is O-witnessed, we obtain $T \vdash [K, s]\Box\neg\varphi$, i.e., $T \vdash K[s]\Box\neg\varphi$. As T is maximal, we have $K[s]\Box\neg\varphi \in T$, thus $[s]\Box\neg\varphi \in \{\theta : K\theta \in T\}$.

Proof (Lemma 8 and Corolary 1). From left-to-right follows directly from the definition of X^c and \sim_K . For the right-to-left direction, we prove the contrapositive: Let $\varphi \in \mathcal{L}$ such that $K\varphi \notin T$. Then, by Lemma 7 and Lemma 4, we obtain that $\{\psi : K\psi \in T\} \cup \{\neg\varphi\}$ is an O-witnessed theory. We can then apply Lindenbaum's Lemma (Lemma 5) and extend it to a maximal O-witnessed theory S such that $\varphi \notin S$.

Proof (Lemma 9). The proof is straightforward. 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 $\vec{u} \in \text{Prop}_O^*$, $[\vec{u}]\varphi \in T$ 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 $v \in O$ such that $\langle v \rangle\neg\varphi$ is consistent with T . Since T is also maximally consistent, we obtain $\langle v \rangle\neg\varphi \in T$, contradicting our initial assumption.

8 Proofs of Section 4: Expressivity

Proof (Proposition 8). Suppose, towards a contradiction, that φ is a formula which is *not* free of dynamic modalities, and moreover that φ is not provably equivalent to any formula of lower complexity (in the sense of $<$) that is free of dynamic modalities. We construct an infinite descending sequence

$$\varphi_0 > \varphi_1 > \dots > \varphi_n > \dots$$

of provably equivalent formulas, none of which is free of dynamic modalities. The construction goes as follows: we first put $\varphi_0 := \varphi$. Then, at step n , assuming given φ_n not free of dynamic modalities and provably equivalent to all the previous formulas, we chose the first dynamic modality occurring in φ_n and apply once to it the relevant Reduction Axiom (from left to right), obtaining a provably equivalent formula φ_{n+1} , which by the previous Lemma has the property that $\varphi_{n+1} < \varphi_n$. By transitivity of provable equivalence, φ_{n+1} is provably equivalent to $\varphi_0 = \varphi$, and (by transitivity of $<$) it is of lower complexity than $\varphi_0 = \varphi$; so, by our assumption above, φ_{n+1} is still not free of dynamic modalities. But the existence of this infinite descending sequence contradicts the well-foundedness of $<$.