Depth-bounded Epistemic Logic

Farid Arthaud
MIT
Cambridge, Massachusetts
farto@csail.mit.edu

Martin Rinard
MIT
Cambridge, Massachusetts
rinard@csail.mit.edu

Epistemic logics model how agents reason about their beliefs and the beliefs of other agents. Existing logics typically assume the ability of agents to reason perfectly about propositions of unbounded modal depth. We present **DBEL**, an extension of **S5** that models agents that can reason about epistemic formulas only up to a specific modal depth. To support explicit reasoning about agent depths, **DBEL** includes depth atoms E_a^d (agent a has depth exactly d) and P_a^d (agent a has depth at least d). We provide a sound and complete axiomatization of **DBEL**.

We extend **DBEL** to support public announcements for bounded depth agents and show how the resulting **DPAL** logic generalizes standard axioms from public announcement logic. We present two alternate extensions and identify two undesirable properties, amnesia and knowledge leakage, that these extensions have but **DPAL** does not. We provide axiomatizations of these logics as well as complexity results for satisfiability and model checking.

Finally, we use these logics to illustrate how agents with bounded modal depth reason in the classical muddy children problem, including upper and lower bounds on the depth knowledge necessary for agents to successfully solve the problem.

1 Introduction

Epistemic logics model how agents reason about their beliefs and the beliefs of other agents. These logics generally assume the ability of agents to perfectly reason about propositions of unbounded modal depth, which can be seen as unrealistic in some contexts [7, 19].

To model agents with the ability to reason only to certain preset modal depths, we extend the syntax of epistemic logic **S5** [8] to depth-bounded epistemic logic (DBEL). The **DBEL** semantics assigns each agent a depth in each state. For an agent to know a formula ψ in a given state of a model, the assigned depth of the agent must be at least the modal depth of ψ, i.e. $d(\psi)$. To enable agents to reason about their own and other agents' depths, **DBEL** includes **depth atoms** E_a^d (agent a has depth exactly d) and P_a^d (agent a has depth at least d). For example, the formula $K_a(P_b^5 \to K_b p)$ expresses the fact that, "agent a knows that whenever agent b is depth at least 5, agent b knows the fact p." Depth atoms enable agents to reason about agent depths and their consequences in contexts in which each agent may have complete, partial, or even no information about agent depths (including its own depth).

We provide a sound and complete axiomatization of **DBEL** (Section 2), requiring a stronger version of the LINDENBAUM lemma which ensures each agent can be assigned a depth (proven in Appendix B). Its satisfiability problem for two or more agents is immediately PSPACE-hard (because **DBEL** includes **S5** as a syntactic fragment). We provide a depth satisfaction algorithm for **DBEL** in PSPACE (Section 5), establishing that the **DBEL** satisfiability problem is PSPACE-complete for two or more agents.

Public announcement logic (PAL) [9] extends epistemic logic with public announcements. PAL includes the following public announcement and knowledge axiom (PAK), which characterizes agents' knowledge after public announcements,

$$[\varphi]K_a\psi \leftrightarrow (\varphi \to K_a[\varphi]\psi). \quad \text{(PAK)}$$

R. Verbrugge (Ed.): Theoretical Aspects of Rationality and Knowledge 2023 (TARK 2023)
EPTCS 379, 2023, pp. 46–65, doi:10.4204/EPTCS.379.7

We extend **DBEL** to include public announcements (Section 3). The resulting depth-bounded public announcement logic (DPAL) provides a semantics for public announcements in depth-bounded epistemic logic, including a characterization of how agents reason when public announcements exceed their epistemic depth. We prove the soundness of several axioms that generalize (PAK) to **DPAL**, first in a setting where each agent has exact knowledge of its own depth, then in the general setting where each agent may have partial or even no knowledge of its own depth. We provide a sound axiom set for **DPAL** as well as an upper bound on the complexity of its model checking problem [1]

We also present two alternate semantics that extend **DBEL** with public announcements (Section 3.3). The resulting logics verify simpler generalizations of (PAK) in the context of depth-bounded agents, but each has one of two undesirable properties that we call *amnesia* and *knowledge leakage*. Amnesia causes agents to completely forget about all facts they knew after announcements, whereas knowledge leakage means shallow agents can infer information from what deeper agents have learned from a public announcement. **DPAL** suffers from neither of these two undesirable properties. We provide a sound and complete axiomatization of the first of the two alternate semantics (Section 4). We also prove the PSPACE-completeness of its satisfiability problem and show that its model checking problem remains P-complete (Section 5).

Finally, we use these logics to illustrate how agents with bounded depths reason in the muddy children reasoning problem [8]. We prove a lower bound and an upper bound on the structure of knowledge of depths required for agents to solve this problem (Section 6).

Related work Logical omniscience, wherein agents are capable of deducing any fact deducible from their knowledge, is a well-known property of most epistemic logics. The ability of agents to reason about facts to unbounded modal depth is a manifestation of logical omniscience. Logical omniscience has been viewed as undesirable or unrealistic in many contexts [8] and many attempts have been made to mitigate or eliminate it [8, 15, 17]. To the best of our knowledge, only Kaneko and Suzuki [11] below have involved modal depth in the treatment of logical omniscience in epistemic logic.

Kaneko and Suzuki [11] define the logic of shallow depths GL_{EF}, which relies on a set E of chains of agents $(i_1,\ldots,i_k)$ for which chains of modal operators $K_{i_1}\cdots K_{i_m}$ can appear. A subset $F \subseteq E$ restricts chains of modal operators along which agents can perform deductions about other agents' knowledge. An effect of bounding agents' depths in **DPAL** is creating a set of allowable chains of modal operators $\cup_a\{(a,i_1,\ldots,i_{d_a}),\ (i_1,\ldots,i_{d_a}) \in \mathscr{A}^{d_a}\}$. Unlike GL_{EF}, the bound on an agent's depth is not global in **DPAL**, it can also be a function of the worlds in the Kripke possible-worlds semantics [8]. In particular, **DPAL**, unlike GL_{EF}, enables agents to reason about their own depth, the depth of other agents, and (recursively) how other agents reason about agent depths. **DPAL** also includes public announcements, which to the best of our knowledge has not been implemented in GL_{EF}.

Kline [12] uses GL_{EF} to investigate the 3-agent muddy children problem, specifically by deriving minimal epistemic structures F that solve the problem. The proof relies on a series of belief sets with atomic updates called "resolutions," with the nested length of the chains in F providing epistemic bounds on the required reasoning. **DPAL**, in contrast, includes depth atoms and public announcements as first-class features. We leverage these features to directly prove theorems expressing that for k muddy children, *(i)* (Theorem 6.2) if the problem is solvable by an agent, that agent must have depth at least $k-1$ and know that it has depth at least $k-1$ (this theorem provides a lower bound on the agent depths required to solve the problem) and *(ii)* (Theorem 6.1) if an agent has depth at least $k-1$, knows it, knows another agent is depth at least $k-2$, knows that the other agent knows of another agent of depth $k-2$, *etc.*, then it

[1] Arthaud and Rinard [3] present a lower bound for this problem, as well as additional results, proofs and content.

can solve the problem (this theorem provides an upper bound on the agent depths necessary to solve the problem). Our depth bounds match the depth bounds of Kline [12] for 3 agents (Theorems 3.1 and 3.3 in [12]), though our bounds also provide conditions on recursive knowledge of depths for the agents as described above.

Dynamic epistemic logic (DEL) [6, 18] introduces more general announcements. Private announcements are conceptually similar to public announcements in **DPAL** in that they may be perceived by only some of the agents. In DEL, model updates depend only on the relation between states in the initial model and the relations in the action model. But in **DPAL**, model updates must also take into account the agent depths in the entire connected components of each state (see Definition 3).

Resource-bounded agents in epistemic logics have been explored by Balbiani et. al [5] (limiting perceptive and inferential steps), Artemov and Kuznets [2] (limiting the computational complexity of inferences), and Alechina et. al [1] (bounding the size of the set of formulas an agent may believe at the same time and introducing communication bounds). Alechina et. al [1] also bound the modal depth of formulas agents may believe, but all agents share the same depth bound and they leave open the question of whether inferences about agent depth or memory size could be implemented, which **DPAL** does.

2 Depth-bounded epistemic logic

The modal depth $d(\varphi)$ of a formula φ, defined as the largest number of modal operators on a branch of its syntactic tree, is the determining factor of the complexity of a formula in depth-bounded epistemic logic (DBEL). Modal operators are the main contributing factor to the complexity of model checking a formula; the recursion depth when checking satisfiability of a formula is equivalent to its modal depth [14]; and bounding modal depth often greatly simplifies the complexity of the satisfiability problem in epistemic logics [16]. Humans are believed to reason within limited modal depth [7, 19].

We extend the syntax of classical epistemic logic by assigning to each agent a in a set of agents $\mathscr{A}$ a depth $d(a,s)$ in each possible world s. The language also includes **depth atoms** E_a^d and P_a^d to respectively express that agent a has depth exactly d and agent a has depth at least d.

To know a formula φ, agents are required to be at least as deep as $d(\varphi)$ and also know that the formula φ is true in the usual possible-worlds semantics sense [8]. We translate the classical modal operator K_a from multi-agent epistemic logic into the operator K_a^∞ with the same properties, therefore $K_a^\infty\varphi$ can be interpreted as "agent a would know φ if a were of infinite depth". The operator $K_a\varphi$ will now take the meaning described above, i.e. $P_a^{d(\varphi)} \wedge K_a^\infty\varphi$.

Definition 1. The language of **DPAL** is inductively defined as, for all agents $a \in \mathscr{A}$ and depths $d \in \mathbb{N}$,

$$\mathscr{L}^\infty := \varphi = p \mid E_a^d \mid P_a^d \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi \mid K_a^\infty\varphi \mid [\varphi]\varphi.$$

The K_a^∞ operator is used mainly as a tool in axiomatization proofs, we call $\mathscr{L}$ the fragment of our logic formulas without any K_a^∞ operators, which will be used in most of our theorems. We further define $\mathscr{H}^\infty$ and $\mathscr{H}$ to respectively be the syntactic fragments of $\mathscr{L}^\infty$ and $\mathscr{L}$ without public announcements $[\varphi]\psi$.

The modal depth d of a formula in $\mathscr{L}^\infty$ is inductively defined as,

$$d(p) = d\left(E_a^d\right) = d\left(P_a^d\right) = 0 \qquad d(\neg\varphi) = d(\varphi) \qquad d([\varphi]\psi) = d(\varphi) + d(\psi)$$
$$d(\varphi \wedge \psi) = \max(d(\varphi), d(\psi)) \qquad d(K_a\varphi) = 1 + d(\varphi) \qquad d(K_a^\infty\varphi) = 1 + d(\varphi).$$

We defer treatment of public announcements $[\varphi]\psi$ to Section 3. We work in the framework of **S5** [8], assuming each agent's knowledge relation to be an equivalence relation, unless otherwise specified—however, our work could be adapted to weaker epistemic logics [8] by removing the appropriate axioms.

All propositional tautologies	$p \to p$, *etc.*
Deduction	$(K_a\varphi \wedge K_a(\varphi \to \psi)) \to K_a\psi$
Truth	$K_a\varphi \to \varphi$
Positive introspection	$(K_a\varphi \wedge P_a^{d(\varphi)+1}) \to K_a(P_a^{d(\varphi)} \to K_a\varphi)$
Negative introspection	$(\neg K_a\varphi \wedge P_a^{d(\varphi)+1}) \to K_a\neg K_a\varphi$
Depth monotonicity	$P_a^d \to P_a^{d-1}$
Exact depths	$P_a^d \leftrightarrow \neg(E_a^0 \vee \cdots \vee E_a^{d-1})$
Unique depth	$\neg(E_a^{d_1} \wedge E_a^{d_2})$ for $d_1 \neq d_2$
Depth deduction	$K_a\varphi \to P_a^{d(\varphi)}$
Modus ponens	From φ and $\varphi \to \psi$, deduce ψ
Necessitation	From φ deduce $P_a^{d(\varphi)} \to K_a\varphi$

Table 1: Sound and complete axiomatization for **DBEL** over $\mathscr{H}$.

Definition 2. A model in **DBEL** is defined as a tuple $M = (\mathscr{S}, \sim, V, d)$ where $\mathscr{S}$ is a set of states, $V : \mathscr{S} \to 2^{\mathscr{P}}$ is the valuation function for atoms and $d : \mathscr{A} \times \mathscr{S} \to \mathbb{N}$ is a depth assignment function. For each agent a, $\sim_a$ is an equivalence relation on $\mathscr{S}$ modeling which states are seen as equivalent in the eyes of a. The semantics are inductively defined over $\mathscr{H}^\infty$ by,

$$(M,s) \models p \iff p \in V(s) \qquad (M,s) \models E_a^d \iff d(a,s) = d \qquad (M,s) \models P_a^d \iff d(a,s) \geq d$$

$$(M,s) \models \neg\varphi \iff (M,s) \not\models \varphi \qquad (M,s) \models \varphi \wedge \psi \iff (M,s) \models \varphi \text{ and } (M,s) \models \psi$$

$$(M,s) \models K_a^\infty\varphi \iff (\forall s',\ s \sim_a s' \implies (M,s') \models \varphi) \qquad (M,s) \models K_a\varphi \iff (M,s) \models P_a^{d(\varphi)} \wedge K_a^\infty\varphi.$$

Note that this definition does not require agents to have any (exact or approximate) knowledge of their own depth. On the other hand, it does not prohibit agents agents from having exact knowledge of their own depths, for instance we could model each agent carrying out some 'meta-reasoning' about its own depth [2] leading each agent to know its own depth exactly. These models are a subset of the class of the models we consider, which we study in more detail in Section 3.1.

As **DBEL** is an extension of **S5** up to renaming of the modal operators, one can expect for it to have a similar axiomatization: one new axiom is needed to axiomatize K_a and three others for depth atoms.

Theorem 2.1. *Axiomatization from Table 1 is sound and complete with respect to* ***DBEL*** *over* $\mathscr{H}$.

Proof. Rather than directly showing soundness and completeness, we show it is equivalent to the axiomatization of Table 3 in Appendix A on the fragment $\mathscr{H}$, which is shown to be sound and complete over $\mathscr{H}^\infty$ in Theorem A.1. We begin by proving any proposition in $\mathscr{H}$ that can be shown using Table 1 can be shown using Table 3 and then that any proof of a formula in $\mathscr{H}$ using the axioms in Table 3 can be shown using those in Table 1.

For the first direction, we prove that the axioms in Table 1 can be proven using those from Table 3. Most of them are immediate applications of bounded knowledge within the axioms of Table 3, along with tautologies when necessary. For positive and negative introspection, see equation (6) below in the proof of the opposite direction of the equivalence. We prove the least evident axiom, the deduction axiom, here as an example:

$$\text{Deduction} \quad (K_a^\infty\varphi \wedge K_a^\infty(\varphi \to \psi)) \to K_a^\infty\psi \tag{1}$$

[2] For instance deducing $P_a^{d(\varphi)}$ from the fact that it knows φ, or deducing $\neg P_a^n$ from the fact that it does not know $K_a^n\top$.

$$\begin{aligned}
\text{Bounded knowledge in (1)} \quad & (K_a^\infty \varphi \wedge K_a^\infty(\varphi \to \psi)) \to P_a^{d(\psi)} \to K_a \psi && (2)\\
\text{Tautology in (2)} \quad & P_a^{\max(d(\varphi),d(\psi))} \to K_a^\infty \varphi \to K_a^\infty(\varphi \to \psi) \to P_a^{d(\psi)} \to K_a \psi && (3)\\
\text{Repeated depth consistency} \quad & P_a^{\max(d(\varphi),d(\psi))} \to (P_a^{d(\varphi)} \wedge P_a^{d(\psi)}) && (4)\\
\text{Bounded knowledge and (3) and (4)} \quad & P_a^{\max(d(\varphi),d(\psi))} \to K_a \varphi \to K_a(\varphi \to \psi) \to K_a \psi && (5)\\
\text{Bounded knowledge in (5)} \quad & K_a \varphi \to K_a(\varphi \to \psi) \to K_a \psi.
\end{aligned}$$

In the other direction, we will show by induction over a proof of a valid formula in $\mathscr{H}$ using Table 3 that it can be transformed into a proof with the same conclusion, using only axioms from Table 1. The transformation of a proof in the first axiomatization is as follows,

- If an item of the proof is a propositional tautology, replace all $K_a^\infty \varphi$ subformulas by $P_a^{d(\varphi)} \to K_a \varphi$, clearly the tautology still holds and it is in Table 1.
- If an item is an instance of the bounded knowledge axiom, replace it with the formula $K_a \varphi \leftrightarrow (P_a^{d(\varphi)} \wedge P_a^{d(\varphi)} \to K_a \varphi)$ which is a consequence of depth deduction and a tautology (and therefore can be added to the proof with two extra steps).
- If it uses any of the other axioms, replace it with the corresponding axiom (with the same name) from Table 1.

We now have a sequence that has the same conclusion (since the conclusion is in $\mathscr{H}$) and only uses axioms from Table 1. The last thing to show for this to be a proof in this axiomatization is that all applications of *modus ponens* and necessitation are still correct within this sequence. To this end, we show by induction that each step of the sequence is the same as the original proof where every $K_a^\infty \varphi$ subformula in each step has been replaced by $P_a^{d(\varphi)} \to K_a \varphi$.

First, note that this is the case for the two first bullet points of our transformation rules above. This is also true of each axiom in the table after our transformation: a proof similar to the one in equation (1) will yield the equivalence for deduction, the only remaining non-trivial cases are positive and negative introspection. For positive introspection, performing the substitution yields,

$$(P_a^{d(\varphi)} \to K_a \varphi) \to P_a^{d(\varphi)+1} \to K_a(P_a^{d(\varphi)} \to K_a \varphi). \tag{6}$$

Through application of a tautology and the depth monotonicity axiom we find it to be equivalent to, $P_a^{d(\varphi)+1} \to K_a \varphi \to K_a(P_a^{d(\varphi)} \to K_a \varphi)$. Therefore, up to adding steps to the proof and using tautologies, we can prove the axiom from Table 1 from the axiom in Table 3 after the substitution. The same can be said of negative introspection through a similar transformation.

Finally, since *modus ponens* and necessitation also maintain the property of replacing $K_a^\infty \varphi$ subformulas in each step by $P_a^{d(\varphi)} \to K_a \varphi$, it is true that the transformed proof is indeed a proof of the same conclusion in Table 1's axiomatization. □

3 Depth-bounded public announcement logic

We next present how to incorporate depth announcements in **DBEL**, which are a key challenge in defining depth-bounded public announcement logic (DPAL). Recall the axiom (PAK) of public announcement logic, $[\varphi]K_a \psi \leftrightarrow (\varphi \to K_a[\varphi]\psi)$. For the right-hand side to be true, agent a must be of depth $d([\varphi]\psi) = d(\varphi) + d(\psi)$ according to **DBEL**. This suggests that an agent must "consume" $d(\varphi)$ of its depth every time an announcement φ is made, meaning that an agent's depth behaves like a depth budget with respect to public announcements.

Moreover, to model that some agents might be too shallow for the announcement φ, each possible world is duplicated in a *negative* version where the announcement has not taken place and a *positive* version where the announcement takes place in the same way as in PAL. Agents who are not deep enough to perceive the announcement see the negative and positive version of the world as equivalent.

Definition 3. Models in **depth-bounded public announcement logic** (DPAL) are defined the same way as in **DBEL** and the semantics is extended to $\mathscr{L}^\infty$ by $(M,s) \models [\varphi]\psi \iff ((M,s) \models \varphi \implies (M \mid \varphi, (1,s)) \models \psi)$, where we define $M \mid \varphi$ to be the model $(\mathscr{S}', \sim', V', d')$, where,

$$\begin{aligned}
&\mathscr{S}' = (\{0\} \times \mathscr{S}) \cup \{(1,s),\ s \in \mathscr{S},\ (M,s) \models \varphi\} \\
&\sim'_a \text{ is the transitive symmetric closure of } R_a \text{ such that,} \\
&(i,s)\, R_a\, (i,s') \iff s \sim_a s' \qquad \text{for } i = 0,1 \\
&(1,s)\, R_a\, (0,s) \iff (M,s) \not\models P_a^{d(\varphi)} \\
&V'((i,s)) = V(s) \qquad \text{for } i = 0,1 \\
&d'(a,(0,s)) = d(a,s) \\
&d'(a,(1,s)) = \begin{cases} d(a,s) & \text{if } d(a,s) < d(\varphi) \\ d(a,s) - d(\varphi) & \text{otherwise.} \end{cases}
\end{aligned} \tag{7}$$

Since public announcements are no longer unconditionally and universally heard by all agents, we revisit the axiom (PAK) in **DPAL**. The determining factor is **depth ambiguity**: agents that are unsure about their own depth introduce uncertainty about which agents have perceived the announcement.

3.1 Unambiguous depths setting

A model verifies the **unambiguous depths** setting whenever each agent knows its own depth exactly:

$$\forall a,s,s', \quad s \sim_a s' \implies d(a,s) = d(a,s'). \tag{8}$$

The proof of the following theorem is given as Proposition C.1 in Appendix C.

Theorem 3.1. *For all $\varphi \in \mathscr{L}^\infty$, the following two properties, respectively called **knowledge preservation** and **traditional announcements**, are valid in **DPAL** in the unambiguous depths setting,*

$$\forall \psi \in \mathscr{L}_a^\infty,\ \neg P_a^{d(\varphi)} \to ([\varphi]K_a\psi \leftrightarrow (\varphi \to K_a\psi)) \tag{KP}$$

$$\forall \psi \in \mathscr{L}^\infty,\ P_a^{d(\varphi)} \to ([\varphi]K_a\psi \leftrightarrow (\varphi \to K_a[\varphi]\psi)), \tag{TA}$$

where $\mathscr{L}_a^\infty$ is the fragment of $\mathscr{L}^\infty$ without depth atoms or modal operators for agents other than a.

Discussion Knowledge preservation (KP) means that an agent who is not deep enough to perceive an announcement φ must not change its knowledge of a formula ψ. However, such a property could not be true of all formulas ψ, for instance if $\psi = K_aK_bp$ but b is deep enough to perceive φ, then the depth adjustment formula (7) could mean that b's depth is now 0, making ψ no longer hold. Even when a is certain about b's depth, its uncertainty about what the announcement entails could also mean that formulas such as $\neg K_bp$ could no longer be true if $P_b^{d(\varphi)}$ and $\varphi \to p$ in the model. This demonstrates that in depth-bounded logics public announcements must introduce uncertainty: if a is unsure what b has perceived, it can no longer hold any certainties about what b does not know. This is not the case in PAL since all agents perceive all announcements. Our treatment of the depth-ambiguous case in Section 3.2 generalizes (KP) to obtain a property (KP') that holds on all formulas in $\mathscr{L}^\infty$.

Traditional announcements (TA) ensures that announcements behave the same as in PAL when the agent is deep enough for the announcement. The caveats from the discussion of (KP) no longer apply here, as any K_b operator that appears in ψ will still appear after the same public announcement operator, meaning that depth variations or knowledge variations are accounted for.

3.2 Ambiguous depths setting

We now abandon the depth unambiguity assumption from equation (8), and explore how properties (KP) and (TA) generalize to settings without depth unambiguity. We find a condition that ensures that sufficient knowledge about other agents' depths is given to a in order to maintain its recursive knowledge about other agents. The proof to the following theorem is given as Proposition C.2 in Appendix C.

Theorem 3.2. *For any $\varphi \in \mathscr{L}^\infty$, let $\mathscr{F}_\varphi : \mathscr{L}^\infty \to \mathscr{L}^\infty$ be inductively defined as,*

$$\mathscr{F}_\varphi(p) = \mathscr{F}_\varphi(E_a^d) = \mathscr{F}_\varphi(P_a^d) = \top \qquad \mathscr{F}_\varphi(\neg\psi) = \mathscr{F}_\varphi(\psi) \qquad \mathscr{F}_\varphi(\psi \wedge \chi) = \mathscr{F}_\varphi(\psi) \wedge \mathscr{F}_\varphi(\chi)$$

$$\mathscr{F}_\varphi(K_a\psi) = \neg K_a^\infty(\varphi \to P_a^{d(\varphi)}) \wedge K_a^\infty(\varphi \to \neg P_a^{d(\varphi)} \vee P_a^{d(\varphi)+d(\psi)}) \wedge K_a^\infty \mathscr{F}_\varphi(\psi)$$

$$\mathscr{F}_\varphi(K_a^\infty\psi) = \neg K_a^\infty(\varphi \to P_a^{d(\varphi)}) \wedge K_a^\infty \mathscr{F}_\varphi(\psi) \qquad \mathscr{F}_\varphi([\psi_1]\psi_2) = \mathscr{F}_\varphi(\psi_1) \wedge \mathscr{F}_\varphi(\psi_2).$$

*For all $\varphi \in \mathscr{L}^\infty$, the following two properties are valid in **DPAL**,*

$$\forall \psi \in \mathscr{L}^\infty,\ \mathscr{F}_\varphi(K_a\psi) \to ([\varphi]K_a\psi \leftrightarrow (\varphi \to K_a\psi)) \tag{KP'}$$

$$\forall \psi \in \mathscr{L}^\infty,\ K_a^\infty(\varphi \to P_a^{d(\varphi)}) \to ([\varphi]K_a\psi \leftrightarrow (\varphi \to K_a[\varphi]\psi)). \tag{TA'}$$

3.3 Alternate treatments of model updates for public announcements

One question is whether using a definition of public announcements closer to PAL would produce a version of the above axioms closer to (PAK). Eager depth-bounded public announcement logic (EDPAL) below unconditionally decrements the depth value of all agents after public announcements.

Definition 4 (**EDPAL**)**. EDPAL** extends the **DBEL** semantics to include public announcements by defining $(M,s) \models [\varphi]\psi \iff ((M,s) \models \varphi \implies (M \mid \varphi, s) \models \psi)$, where $M \mid \varphi$ is the model $(\mathscr{S}', \sim', V, d')$ in which $\mathscr{S}' = \{s \in \mathscr{S},\ (M,s) \models \varphi\}$, $\sim'_a$ is the restriction of $\sim_a$ to $\mathscr{S}'$, $d'(a,s) = d(a,s) - d(\varphi)$, and d may take values in $\mathbb{Z}$.

EDPAL has a sound and complete axiomatization based on the axiomatization of **DBEL** (Theorem 4.1), which also allows us to prove the complexity result of Theorem 5.1.

However, another consequence of its definition is that excessive public announcements in **EDPAL** can lead an agent to a state in which it cannot reason anymore, as it has consumed its entire depth budget.

Proposition 3.3 (Amnesia)**.** *In **EDPAL**, the formula $\neg P_a^{d(\varphi)} \to [\varphi]\neg K_a\psi$ is valid for all φ and ψ.*

Proof. If $(M,s) \not\models \varphi$ then the implicand is true. If $(M,s) \models \varphi \wedge \neg P_a^{d(\varphi)}$ then the depth of a in $(M \mid \varphi, s)$ will be at most -1, meaning that $(M \mid \varphi, s) \not\models K_a\psi$ for all ψ. ☐

In particular, for $\psi = \top$ one notices that standard intuitions about knowledge fail in **EDPAL**. This property is undesirable: *(i)* one may expect agents to maintain some knowledge even after public announcements that they are not deep enough to understand and *(ii)* deeper agents should be able to continue to benefit from the state of knowledge of shallower agents even after the shallower agents have exceeded their depth.

One way to try to remedy this property is to change model updates in **EDPAL** to make agents perceive announcements only when they are deep enough to understand them. The resulting asymmetric depth-bounded public announcement logic (ADPAL) removes depth from an agent's budget only when it is deep enough for an announcement, and only updates its equivalence relation in states where it is deep enough for the announcement.

Definition 5 (**ADPAL**)**.** **ADPAL** extends the **DBEL** semantics to include public announcements by defining $(M,s) \models [\varphi]\psi \iff ((M,s) \models \varphi \implies (M \mid \varphi, s) \models \psi)$, where $M \mid \varphi$ is the model $(\mathscr{S}, \sim', V, d')$,

$$s \not\sim'_a s' \iff s \not\sim_a s' \text{ or } \begin{cases} (M,s) \models P_a^{d(\varphi)} \\ (M,s) \models \varphi \iff (M,s') \not\models \varphi, \end{cases}$$

$$d'(a,s) = \begin{cases} d(a,s) & \text{if } d(a,s) < d(\varphi) \\ d(a,s) - d(\varphi) & \text{otherwise.} \end{cases}$$

The relations $\sim_a$ are only assumed to be reflexive (as opposed to equivalence relations earlier).

Unfortunately, in **ADPAL** an agent that is too shallow for an announcement could still learn positive information that was learned by another agent who is deep enough to perceive the announcement. We call this property *knowledge leakage* as reflected in the following proposition.

Proposition 3.4 (Knowledge leakage)**.** ***ADPAL*** *does not verify the* $\rightarrow$ *direction of* (KP').

Proof. Consider three worlds, $\{0,1,2\}$ and three agents a,b,c. The relations for a and c are identity, the relation for b is the symmetric reflexive closure of, $0 \sim_b 1 \sim_b 2$. The depth of a is 1 everywhere, b's depth is $0,2,0$ in each respective state and the depth of c is 2 everywhere. The atom p_0 is true only in 0 and 1. Consider $\varphi = K_c K_c p_0$, it is true in 0 and 1 only, and consider $\psi = K_b p_0$. $K_a\psi$ is not true in state 1, however $[\varphi]K_a\psi$ is. Moreover, one can easily check that $\mathscr{F}_\varphi(K_a\psi)$ is true in that state. ☐

The proof provides a practical example of such leakage in **ADPAL** and we further demonstrate knowledge leakage in Proposition 6.4 in the muddy children reasoning problem (see Section 6).

Note how each direction of the equivalence in (KP') expresses ($\rightarrow$) that no knowledge leakage occurs and ($\leftarrow$) no amnesia occurs. As shown in Theorem 3.2, **DPAL** verifies both directions and thus has neither amnesia nor knowledge leakage. As reflected in the following proposition, although EDPAL has amnesia, it doesn't have knowledge leakage and verifies (TA).

Proposition 3.5. *[3]* ***EDPAL*** *verifies* (TA) *and the* $\rightarrow$ *direction in* (KP) *over* $\psi \in \mathscr{L}^\infty$*, but not the converse.*

4 Axiomatizations

Theorem 4.1. *The axiomatization in Table 2 is sound and complete with respect to* ***EDPAL*** *(Definition 4) over the fragment* $\mathscr{L}$*.*

All axioms from Table 1	
Atomic permanence	$[\varphi]p \leftrightarrow (\varphi \to p)$
Depth adjustment	$\forall d \in \mathbb{Z},\ [\varphi]E_a^d \leftrightarrow \left(\varphi \to E_a^{d(\varphi)+d}\right)$
Negation announcement	$[\varphi]\neg\psi \leftrightarrow (\varphi \to \neg[\varphi]\psi)$
Conjunction announcement	$[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$
Knowledge announcement	$[\varphi](P_a^{d(\psi)} \to K_a\psi) \leftrightarrow (\varphi \to P_a^{d(\varphi)+d(\psi)} \to K_a[\varphi]\psi)$
Announcement composition	$[\varphi][\psi]\chi \leftrightarrow ([\varphi \wedge [\varphi]\psi]\chi)$
Modus ponens	From φ and $\varphi \to \psi$, deduce ψ
Necessitation	From φ deduce $P_a^{d(\varphi)} \to K_a\varphi$

Table 2: Sound and complete axiomatization of **EDPAL** over $\mathscr{L}$.

Proof. Similarly to the proof of Proposition 2.1, rather than directly showing soundness and completeness we show it is equivalent to the axiomatization of Table 4, which is shown to be sound and complete for **EDPAL** in Theorem A.2 in Appendix A.

In the first direction, all axioms in Table 2 can be shown using those in Table 4 immediately, either from the proof of Proposition 2.1 or because they are the same. The only difficulty lies in knowledge announcement, but a proof similar to equation (1) shows it is sound.

The other direction also follows the exact same proof as in Proposition 2.1: the public announcement axioms are direct translations of the same axioms in Table 4 by replacing the $K_a^\infty\varphi$ subformulas with $P_a^{d(\varphi)} \to K_a\varphi$. The proof transformation from Proposition 2.1 therefore still yields a proof of the same formula in this axiomatization, which proves completeness. □

We now present a sound set of axioms for **DPAL**. The main missing axioms for a sound and complete axiomatization are knowledge and public announcements, which we explored in the previous section, and announcement composition. In fact, announcement composition cannot exist in **DPAL**, since making a single announcement of depth $d_1 + d_2$ can behave very differently from making an announcement of depth d_1 followed by another of depth d_2, for instance when an agent's depth is between d_1 and $d_1 + d_2$.

Theorem 4.2. *Replacing knowledge announcement by* (KP') *and* (TA') *and depth adjustment by,*

$$\forall d \in \mathbb{N}, \quad [\varphi]E_a^d \leftrightarrow \left(\varphi \to \left((P_a^{d(\varphi)} \wedge E_a^{d+d(\varphi)}) \vee (\neg P_a^{d(\varphi)} \wedge E_a^d)\right)\right)$$

*in Table 2 produces a set of sound axioms with respect to **DPAL*** [3].

Proof. Theorem 3.2 verifies the two axioms (KP') and (TA'). The proofs for most axioms follows from Theorem 4.1 and that knowledge is defined the same way in both semantics. In particular, atomic permanence and conjunction announcement axioms are proven in Theorem 3.1's induction for (KP).

We are left to show depth adjustment,

$$\begin{aligned}
(M,s) \models [\varphi]E_a^d &\iff (M,s) \models \varphi \implies (M \mid \varphi, (1,s)) \models E_a^d \\
&\iff (M,s) \models \varphi \implies \begin{cases} d(a,s) = d + d(\varphi) & \text{if } d(a,s) \geq d(\varphi) \\ d(a,s) = d & \text{if } d(a,s) < d(\varphi) \end{cases} \\
&\iff (M,s) \models \varphi \to \left((P_a^{d(\varphi)} \wedge E_a^{d+d(\varphi)}) \vee (\neg P_a^{d(\varphi)} \wedge E_a^d)\right).
\end{aligned}$$
□

[3] One could also easily add axioms for K_a^∞ modal operators, for instance using those from Table 4 in Appendix A.

5 Complexity

We first state that adding depth bounds does not change the complexity of **S5** and PAL respectively.

Theorem 5.1. *The satisfiability problems for **DBEL** with $n \geq 2$ agents and for **EDPAL** are PSPACE-complete.*

Proof. The lower bound results from PSPACE-completeness of $\mathbf{S5}_n$ for $n \geq 2$ [10] and PAL [14], respective syntactic fragments of **DBEL** and **EDPAL**.

For both logics, we begin by translating $K_a\varphi$ subformulas into $P_a^{d(\varphi)} \wedge K_a^\infty \varphi$, which only increases formula size at most linearly. Then, in the case of **EDPAL**, using the same translation as Lemma 9 of [14], we translate formulas with public announcement φ into equivalent formulas $t(\varphi)$ without public announcement such that $|t(\varphi)|$ is at most polynomial in $|\varphi|$ (this is possible because the axiomatization of K_a^∞ with relation to public announcements is the same).

We have therefore transformed our formula φ into an equivalent formula in the syntactic fragment without K_a operators or public announcements of polynomial size relative to the initial formula φ's size.

We can then use the ELE-World procedure from Figure 6 of [14] by re-defining types to accommodate for depth atoms. As a reminder, we define $\mathbf{cl}(\Gamma)$ for any set of formulas Γ to be the smallest set of formulas containing Γ and closed by single negation and sub-formulas. We then say that $\gamma \subseteq \mathbf{cl}(\Gamma)$ is a type if all of the following are true,

1. $\neg\psi \in \gamma$ if and only if $\psi \notin \gamma$ when ψ is not a negation
2. if $\psi \wedge \chi \in \mathbf{cl}(\Gamma)$ then $\psi \wedge \chi \in \gamma$ if and only if $\psi \in \gamma$ and $\chi \in \gamma$
3. if $K_a^\infty \psi \in \gamma$ then $\psi \in \gamma$
4. if $P_a^d \in \gamma$ then $\neg P_a^{d'} \notin \gamma$ and $E_a^{d'} \notin \gamma$ for all $d' < d$
5. if $E_a^d \in \gamma$ then $E_a^{d'} \notin \gamma$ for all $d' \neq d$ and $\neg P_a^{d'} \notin \gamma$ for $d' < d$
6. if $\neg P_a^d \in \gamma$ then there exists $d' < d$ such that $\neg E_a^{d'} \notin \gamma$
7. $\neg P_a^0 \notin \gamma$

Clearly, checking that a subset of $\mathbf{cl}(\Gamma)$ is not a type does not increase the space complexity of the algorithm. Lemma 18 from [14] remains true here, i.e. the procedure ELE-World returns true if and only if the formula is satisfiable. It is sufficient for this to show that any type has a consistent depth assignment for all agents, as it is clear that if any of the new rules introduced for depths are violated the formula is not satisfiable.

If the type contains E_a^d then it contains only one such depth atom per rule 5, the only $P_a^{d'}$ it contains are for $d' \leq d$ per rule 4, and it does not contain $\neg P_a^{d'}$ for $d' \leq d$ per rule 5, therefore $d(a) = d$ is a consistent setting. If it does not contain any E_a^d, it may contain a number of inequalities polynomial in $|\varphi|$, that admit a solution in $\mathbb{N}$ by rule 7. Therefore a possible algorithm is $d_0 = \max\{d',\ P_a^{d'} \in \gamma\}$ and then $d(a) = \min\{d',\ d' \geq d_0,\ \neg E_a^{d'} \notin \gamma\}$. If no P_a^d are in the type, then $d_0 = \min\{d',\ \neg P_a^{d'} \in \gamma\}$ and $d(a) = \max\{d',\ d' \leq d_0, \neg E_a^{d'} \notin \gamma\}$ are a possible choice (this choice will always be greater or equal to 0 because of rules 7 and 6 above). Finally, if there are no depth atoms in the type, the formula is clearly satisfiable for any choice of $d(a)$. □

The model checking problem remains P-complete in **DBEL**, using the same algorithm as for **S5** [8]. For **EDPAL** and **ADPAL**, the model checking problem is P-complete, as the same algorithm as PAL can be used, relying on the fact that model size can only decrease after announcements [13] (the lower bounds results from the fact that PAL is a fragment of both). This is however not the case of **DPAL**, where model size grows after announcements, potentially exponentially, in fact model checking in **DPAL** is NP-hard [3].

Theorem 5.2. *The complexity of model checking for finite models in **DPAL** is in EXPTIME. An upper bound in time complexity for checking φ in M is $O(2^{2^{|\varphi|}}\|M\|)$, where $\|M\|$ is the sum of the number of states and number of pairs in each relation of M.*

Proof. The model-checking algorithm is the same as the one for public announcement logic [13]: a tree is built from subformulas φ, with splits introduced only for subformulas of the form $[\psi]\chi$, with ψ to the left and χ to the right. Treating a node labeled ψ means labeling each state in M with either ψ or $\neg\psi$. The tree is treated from bottom-left to the top, always going up first except when a node of the type $[\psi]\chi$ is found. In that case, since the nodes in the left sub-tree have been treated, we can build $M \mid \psi$ easily in time $O(\|M\|)$ from the truth value of ψ and the depth functions of M. Moreover, the size of $M \mid \psi$ is at most $4\|M\|$.

To see this, consider an equivalence class for $\sim_a$ in M of size k, it has exactly k^2 connections within it. The number of states it creates in $M \mid \psi$ is at most $2k$, and the number of connections it creates is at most $4k^2$. Each connection being in exactly one connected component means the bound holds.

Therefore we can recurse in the right sub-tree with $M \mid \varphi$ to check χ in time $O(2^{2^{|\chi|}} \times 4\|M\|)$. Writing $O(\|M\|) \leq c\|M\|$ the time necessary to build $M \mid \varphi$, we find that checking $[\psi]\chi$ takes time at most $O((c + 2^{2^{|\psi|}} + 2^{2^{|\chi|+2}})\|M\|) = O(2^{2^{|[\psi]\chi|}}\|M\|)$. □

6 Muddy children

Consider the well-known muddy children reasoning problem, where n children convene after playing outside with mud. $k \geq 1$ of them have mud on their foreheads, but have no way of knowing it. The father, an external agent, announces that at least one child has mud on their forehead. Then, he repeatedly asks if any child would like to go wash themselves. After exactly $k-1$ repetitions of the father's question, all muddy children understand they are muddy and go wash themselves. Readers unfamiliar with the reasoning problem and its solution are directed to Van Ditmarsch et. al [18]'s treatment using PAL.

Consider the set of states $\{0,1\}^n$, where each tuple contains n entries indicating for each child if they are muddy (1) or not (0). For the sake of simplicity and since it is of depth 0, we assume the father's announcement has taken place and therefore define the Kripke structure M_n with states $\{0,1\}^n \setminus \{0\}^n$ with the usual definition of the agents' knowledge relations [8]. We define the **DPAL** class of muddy children models to be models $\hat{M}_n$ extending M_n with any depth function. We name m_i the atom expressing that child i is muddy.

We number the agents in $[|0; n-1|]$, where the first k are muddy, and focus on the reasoning of one agent (without loss of generality agent 0) to understand that it is muddy. Recall the definition of the dual of public announcements, $\langle\varphi\rangle\psi := \neg[\varphi]\neg\psi$ and define the following series of formulas for $i \leq k$,

$$\varphi_i = \langle\neg K_{i-1}m_{i-1}\rangle\langle\neg K_{i-2}m_{i-2}\rangle\cdots\langle\neg K_1 m_1\rangle K_0 m_0.$$

Here φ_k states that if each of the children from $k-1$ to 1 announce one after the other they don't know they are muddy, then child 0 knows that they (child 0) are muddy [4] It is well known this formula is true for unbounded agents in M_n in PAL (it is also a consequence of Theorem 6.1 below). The following two

[4]These announcements are a sufficient subset of the full announcements $\wedge_{j=1,\ldots,n}\neg(K_j m_j \vee K_j \neg m_j)$ in the usual formulation.

theorems define a sufficient structure of knowledge of depths for the formula to be true and a necessary condition on the structure of knowledge of depths for it to be true.

Theorem 6.1 (Upper bound)**.** *For all three semantics,* $K_0\left(P_0^{k-1} \wedge K_1(P_1^{k-2} \wedge \cdots K_{k-1}(P_{k-1}^0)\cdots)\right) \rightarrow \varphi_k$ *is true in all muddy children models* $\hat{M}_n$ *in the initial state.*

Note that this formula directly provides an upper bound on the structure of depths and knowledge about depths: it shows a sufficient condition on the knowledge of depths for the problem to be solvable by agent 0. Moreover, the upper bound for one child readily generalizes to a sufficient condition for all children to understand they are muddy: each muddy child must know they are of depth at least $k-1$, know at least some other muddy child knows they are of depth at least $k-2$, and know that that other child knows some other muddy child knows they are of depth at least $k-3$, *etc*.

Proof. For the sake of simplicity and since it does not change the treatment of the problem, we assume $n = k$. We show the result for **DPAL**, as the treatments for **EDPAL** and **ADPAL** are similar.

We will show the result by induction over k. Denote $s_k = (1,\ldots,1)$ the true state of the world where all the children are muddy.

For $k = 2$, we assume $K_0P_0^1$ and want to show $\neg K_1m_1 \wedge [\neg K_1m_1]K_0m_0$. First notice that $(\hat{M}_2, s_2) \models \neg K_1m_1$, simply because it considers the state $(1,0)$ to also be possible. In the state $(0,1)$, child 1 knows it is muddy. Therefore, the set of states for the successful part of the model update will be $(1,(1,1))$ and $(1,(1,0))$. Moreover, since $K_0P_0^1$, it is deep enough in s_2 to not have any links to the unsuccessful part of the model update, therefore it knows m_0.

Consider some $k > 2$, we denote S_i the set of states that are "active" when considering φ_i. More precisely, we set $S_i = \{0,1\}^i \times \{1\}^{k-i} \setminus \{0\}^k$. We will show that after $k-i$ announcements, the remainder of the problem is equivalent to checking φ_i on the subgraph induced by the states S_i. This is evident for $i = k$ by definition, we now show by descending induction that it is equivalent to checking φ_2 on S_2, which we have just verified to be true.

Firstly, it is true that $(\hat{M}_n, s_k) \models \neg K_{k-1}m_{k-1}$ since child $k-1$ considers possible the state $(1,\ldots,1,0)$. The set of states in which $K_{k-1}m_{k-1}$ holds is exactly $(0,\ldots,0,1)$. Therefore, the model update will create a copy of all other states. We then notice that the set of states whose last component is 0 can be ignored in the rest of the problem: they are not reachable from s_k by any sequence of $\sim_i$ that does not contain $\sim_{k-1}$ and the rest of the formula φ_{k-1} to be checked does not use any modal operators for agent $k-1$ any more. These states will never be reached and can therefore be removed without altering the result of the rest of the execution.

We are therefore restricting ourselves, after the model update, to the set of states S_{k-1} in the positive part of the model. Note however there are still possibly links between the negative part of the model and S_{k-1} in the positive part of the model. We will show that these links have no effect on the checking of the rest of the formula, by showing that links for child i find themselves in $S_{k-1} \setminus S_i$: therefore, by the time we query modal operator i, the set of ignored states will contain all states with a link for child i.

For child $i < k-1$, the information we have about its depth is $K_0K_1\cdots K_iP_i^{k-1-i}$ before the model update. Therefore, we in particular know it is deep enough for the announcement (which is of depth $1 \leq k-1-i$) in the set of states in which the i first components might have changed compared to s_k but the last $k-1-i$ are all fixed to 1: this is exactly S_i.

We have shown that the recursive check in $M \mid \neg K_{k-1}m_{k-1}$ will take place on a set of states for which the execution is equivalent to S_{k-1} and on which we will have to check the formula φ_{k-1}. Finally, since the depths of each agent other than $k-1$ was at least 1 on S_{k-2}, they are reduced by 1 and the induction hypothesis on depths for $k-2$ is also verified. $\square$

Theorem 6.2 (Lower bound)**.** *For **DPAL**, the formula $\varphi_k \to K_0 P_0^{k-1}$ is true in all models $\hat{M}_n$.*

Proof. We use the notations from the proof of Theorem 6.1 above. Notice first that all of the announcements remain true when they are performed, because $\neg K_{k-1}^{\infty} m_{k-1} \to \neg K_{k-1} m_{k-1}$ and the implicant is true by the usual lower bound for muddy children (it takes k announcements for any child to know they are muddy).

Assume by contraposition that $d(0, s_k) = i < k-1$ or $d(0, \tilde{s}_k) = i < k-1$ initially, where $\tilde{s}_k$ is the state $(0, 1, \ldots, 1)$ of $\hat{M}_n$. After i public announcements, it will be true that $\neg K_0 m_0$ still, as well as $\neg K_0 \neg E_0^0$ since each public announcement is of depth 1. The former is a consequence of the usual lower bound for muddy children, and can be derived from the proof in Theorem 6.1 using symmetry between 0 and $k-1-i$ after the i announcements and monotonicity of knowledge of atoms: if the depths are lower than they were in the previous proof, there are more states and more links in the updated model and therefore $\neg K_{k-1-i} m_{k-1-i}$ remains true.

Therefore in this model after i announcements, either s_k or $\tilde{s}_k$ sees agent 0 of depth 0 and both states are still connected by $\sim_0$. This means that for the next announcement, since $\neg K_0 m_0$ after each announcement except potentially the last using the same argument as above, we will have the chain of connections $(1, s_k) \sim_0' (0, s_k) \sim_0' (0, s_k')$ or $(1, s_k) \sim_0' (1, \tilde{s}_k) \sim_0' (0, \tilde{s}_k)$. This means that by an immediate induction, after the $k-i$ announcements it is still true that $\neg K_0 m_0$: this is a contradiction with φ_k. □

A stronger lower bound for each child is available [3], with recursive conditions on the depth of all agents similarly to Theorem 6.1. This formula provides a lower bound on the knowledge of depths of the agent 0 to be able to solve the problem: it must be depth at least $k-1$ and know so. By symmetry, this generalizes to any child or any set of children solving the problem.

Finally, we present propositions that illustrate how *amnesia* in **EDPAL** (Proposition 3.3) and *knowledge leakage* in **ADPAL** (Proposition 3.4) manifest in the muddy children problem. These propositions are easily verified by computing explicitly the models after updates.

Proposition 6.3 (Amnesia in **EDPAL**)**.** *Consider the instance of muddy children M_3, where child i is unambiguously of depth $2-i$, i.e. $d(i, \cdot) = 2-i$. The formula $\langle \neg K_2 m_2 \rangle \langle \neg K_1 m_1 \rangle \neg K_2 \top$ is true in **EDPAL** but not in **DPAL** or **ADPAL**. This means that in **EDPAL**, after the first two announcements, agent 2 does not know anything anymore.*

Proposition 6.4 (Knowledge leakage in **ADPAL**)**.** *The formula $\langle K_1 \neg K_2 m_2 \rangle K_1 K_0 m_0$ is true in **ADPAL** but not in **DPAL** or **EDPAL**. In **ADPAL**, agent 1 has deduced the conclusion of agent 0's reasoning, despite not being deep enough to perceive the announcement. Moreover, if agent 0 were of depth 1 it would not be true that $\langle K_1 \neg K_2 m_2 \rangle K_0 m_0$: agent 0 would not be able to deduce what agent 1 has deduced.*

Library Alongside this paper, we publish code for a library for multi-agent epistemic logic model checking and visualization in Python. It implements depth-unbounded PAL models as well as **DPAL**, **EDPAL** and **ADPAL**. The code is available in an online repository [4]. The code can also be used to generate illustrations of model updates in the muddy children reasoning problem [3] under the assumptions of Theorem 6.1 above.

Conclusion We have shown how **S5** and public announcement logic (PAL) can be extended to incorporate bounded-depth agents. We have shown completeness results for several of the resulting logics and explored the relationship between public announcements and knowledge in **DPAL**, as well as complexity bounds for these logics. We finally illustrated the behavior of depth-bounded agents in the muddy children reasoning problem, where we showed upper and lower bounds on depths (and recursive knowledge of depths) necessary and sufficient to solve the problem. These results extend epistemic logics to support formal reasoning about agents with limited modal depth.

References

[1] Natasha Alechina, Brian Logan, Nguyen Hoang Nga & Abdur Rakib (2008): *Reasoning about Other Agents' Beliefs under Bounded Resources*. In John-Jules Ch. Meyer & Jan M. Broersen, editors: *Knowledge Representation for Agents and Multi-Agent Systems, KRAMAS 2008, Sydney, Australia, September 17, 2008*, *Lecture Notes in Computer Science* 5605, Springer, pp. 1–15, doi:10.1007/978-3-642-05301-6_1.

[2] Sergei N. Artëmov & Roman Kuznets (2009): *Logical omniscience as a computational complexity problem*. In Aviad Heifetz, editor: *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2009), Stanford, CA, USA, July 6-8, 2009*, pp. 14–23, doi:10.1145/1562814.1562821.

[3] Farid Arthaud & Martin Rinard (2023): *Depth-bounded epistemic logic*. *CoRR* abs/2305.08607, doi:10.48550/arXiv.2305.08607.

[4] Farid Arthaud & Martin Rinard (2023): *Library for depth-bounded epistemic logic*. https://gitlab.com/farid-fari/depth-bounded-epistemic-logic.

[5] Philippe Balbiani, David Fernández-Duque & Emiliano Lorini (2016): *A Logical Theory of Belief Dynamics for Resource-Bounded Agents*. In: *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, ACM, pp. 644–652, doi:10.5555/2936924.2937020.

[6] Alexandru Baltag & Lawrence S. Moss (2004): *Logics for Epistemic Programs*. *Synth.* 139(2), pp. 165–224, doi:10.1023/B:SYNT.0000024912.56773.5e.

[7] Thomas Bolander, Robin Engelhardt & Thomas S. Nicolet (2020): *The Curse of Shared Knowledge: Recursive Belief Reasoning in a Coordination Game with Imperfect Information*. *CoRR* abs/2008.08849, doi:10.48550/arXiv.2008.08849.

[8] Ronald Fagin, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi (1995): *Reasoning About Knowledge*. MIT Press, doi:10.7551/mitpress/5803.001.0001.

[9] Jelle Gerbrandy & Willem Groeneveld (1997): *Reasoning about Information Change*. *J. Log. Lang. Inf.* 6(2), pp. 147–169, doi:10.1023/A:1008222603071.

[10] Joseph Y. Halpern & Yoram Moses (1992): *A Guide to Completeness and Complexity for Modal Logics of Knowledge and Belief*. *Artif. Intell.* 54(2), pp. 319–379, doi:10.1016/0004-3702(92)90049-4.

[11] Mamoru Kaneko & Nobu-Yuki Suzuki (2000): *Epistemic Logic of Shallow Depths and Game Theoretical Applications*. In Frank Wolter, Heinrich Wansing, Maarten de Rijke & Michael Zakharyaschev, editors: *Advances in Modal Logic 3, papers from the third conference on "Advances in Modal logic," held in Leipzig, Germany, 4-7 October 2000*, World Scientific, pp. 279–298, doi:10.1142/9789812776471_0015.

[12] J Jude Kline (2013): *Evaluations of epistemic components for resolving the muddy children puzzle*. *Economic Theory* 53(1), pp. 61–83, doi:10.1007/s00199-012-0735-x.

[13] Barteld Kooi & Johan Benthem (2004): *Reduction axioms for epistemic actions*. In: *Advances in Modal Logic 5, papers from the fifth conference on "Advances in Modal logic," held in Manchester, UK, 9-11 September 2004*, King's College Publications, pp. 197–211. Available at https://core.ac.uk/download/pdf/148195794.pdf.

[14] Carsten Lutz (2006): *Complexity and succinctness of public announcement logic*. In Hideyuki Nakashima, Michael P. Wellman, Gerhard Weiss & Peter Stone, editors: *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8-12, 2006*, ACM, pp. 137–143, doi:10.1145/1160633.1160657.

[15] John-Jules Ch Meyer (2003): *Modal epistemic and doxastic logic*. *Handbook of philosophical logic*, pp. 1–38, doi:10.1007/978-94-017-4524-6_1.

[16] Linh Anh Nguyen (2004): *On the Complexity of Fragments of Modal Logics*. In Renate A. Schmidt, Ian Pratt-Hartmann, Mark Reynolds & Heinrich Wansing, editors: *Advances in Modal Logic 5, papers from the fifth conference on "Advances in Modal logic," held in Manchester, UK, 9-11 September 2004*, King's College Publications, pp. 249–268. Available at http://www.aiml.net/volumes/volume5/Nguyen.ps.

[17] Kwang Mong Sim (1997): *Epistemic logic and logical omniscience: A survey*. Int. J. Intell. Syst. 12(1), pp. 57–81, doi:10.1002/(SICI)1098-111X(199701)12:1<57::AID-INT3>3.0.CO;2-X.

[18] Hans Van Ditmarsch, Wiebe van Der Hoek & Barteld Kooi (2007): *Dynamic epistemic logic*. 337, Springer Science & Business Media, doi:10.1007/978-1-4020-5839-4.

[19] Rineke Verbrugge & Lisette Mol (2008): *Learning to Apply Theory of Mind*. J. Log. Lang. Inf. 17(4), pp. 489–511, doi:10.1007/s10849-008-9067-4.

A Axiomatization proofs

All propositional tautologies	$p \to p$, *etc.*
Deduction	$(K_a^\infty \varphi \wedge K_a^\infty(\varphi \to \psi)) \to K_a^\infty \psi$
Truth	$K_a^\infty \varphi \to \varphi$
Positive introspection	$K_a^\infty \varphi \to K_a^\infty K_a^\infty \varphi$
Negative introspection	$\neg K_a^\infty \varphi \to K_a^\infty \neg K_a^\infty \varphi$
Depth monotonicity	$P_a^d \to P_a^{d-1}$
Exact depths	$P_a^d \leftrightarrow \neg(E_a^0 \vee \cdots \vee E_a^{d-1})$
Unique depth	$\neg(E_a^{d_1} \wedge E_a^{d_2})$ for $d_1 \neq d_2$
Bounded knowledge	$K_a \varphi \leftrightarrow P_a^{d(\varphi)} \wedge K_a^\infty \varphi$
Modus ponens	From φ and $\varphi \to \psi$, deduce ψ
Necessitation	From φ deduce $K_a^\infty \varphi$

Table 3: Sound and complete axiomatization of **DBEL** over $\mathscr{H}^\infty$.

Theorem A.1. *Axiomatization from Table 3 is sound and complete with respect to* ***DBEL*** *over* $\mathscr{H}^\infty$.

Proof. Soundness of all of these axioms is immediate: the definition of K_a^∞ follows that of **S5** and so do the axioms, those concerning depth atoms are consequences of linear arithmetic, and the bounded knowledge axiom follows immediately from the definition of K_a in the semantics.

For completeness, first note we can translate any formula φ in $\mathscr{H}^\infty$ into an equivalent formula $t(\varphi)$ that does not contain any P_a^d atoms or K_a modal operators using the exact depths and bounded knowledge axioms (which we know to be sound). Call **S5D** this fragment of **DBEL**.

We will use a proof through the LINDENBAUM lemma and the truth lemma, to this end we need to complete the definition for the canonical model to add a depth function. As a reminder, the proof is as follows: if φ cannot be shown within the axiomatization in Table 1, i.e. $\nvdash \varphi$, then we show that $\not\models \varphi$ by showing there is a state in the canonical model in which it does not hold.

The canonical model M^c is the model whose states are maximally consistent sets Γ of formulas for our axiomatization and whose states are related by $\sim_a$ if the set of formulas a knows is the same in both states. Its valuation function for atoms $V(\Gamma)$ is simply the set of axioms in Γ, i.e. $\Gamma \cap \mathscr{P}$.

We restrict M^c to sets Γ that contain at least some E_a^d for each agent $a \in \mathscr{A}$ and by the unique depth axiom we define $d(a,\Gamma) = \max\{d,\ E_a^d \in \Gamma\}$, since Γ contains exactly one depth to be consistent. This completes M^c into a **DBEL** model.

All propositional tautologies	$p \to p$, *etc.*
Deduction	$(K_a^\infty \varphi \wedge K_a^\infty(\varphi \to \psi)) \to K_a^\infty \psi$
Truth	$K_a^\infty \varphi \to \varphi$
Positive introspection	$K_a^\infty \varphi \to K_a^\infty K_a^\infty \varphi$
Negative introspection	$\neg K_a^\infty \varphi \to K_a^\infty \neg K_a^\infty \varphi$
Atomic permanence	$[\varphi]p \leftrightarrow \varphi \to p$
Depth adjustment	$\forall d \in \mathbb{Z},\ [\varphi]E_a^d \leftrightarrow \left(\varphi \to E_a^{d(\varphi)+d}\right)$
Negation announcement	$[\varphi]\neg\psi \leftrightarrow (\varphi \to \neg[\varphi]\psi)$
Conjunction announcement	$[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$
Knowledge announcement	$[\varphi]K_a^\infty \psi \leftrightarrow (\varphi \to K_a^\infty[\varphi]\psi)$
Announcement composition	$[\varphi][\psi]\chi \leftrightarrow ([\varphi \wedge [\varphi]\psi]\chi)$
Depth monotonicity	$P_a^d \to P_a^{d-1}$
Exact depths	$P_a^d \leftrightarrow \neg(E_a^0 \vee \cdots \vee E_a^{d-1})$
Unique depth	$\neg(E_a^{d_1} \wedge E_a^{d_2})$ for $d_1 \neq d_2$
Bounded knowledge	$K_a \varphi \leftrightarrow P_a^{d(\varphi)} \wedge K_a^\infty \varphi$
Modus ponens	From φ and $\varphi \to \psi$, deduce ψ
Necessitation	From φ deduce $K_a^\infty \varphi$

Table 4: Sound and complete axiomatization of **EDPAL**.

Since $\not\vdash \varphi$, the set $\{\neg\varphi\}$ is consistent for our axiomatization. We must now show we can extend this set into a maximal consistent set of formulas that contains a depth atom E_a^d for each agent a.

However, this stronger requirement is not satisfied by the usual LINDENBAUM lemma, since a consistent set of formulas could be $\{P_a^d,\ d \in \mathbb{N}\}$ (which is not consistent with any E_a^d). Note however we only need it to hold for a finite set of formulas (namely $\{\neg\varphi\}$): Lemma B.1 below proves this version of the LINDENBAUM lemma, by showing there must exist some E_a^d that is consistent with any finite set for each a, and then a maximally consistent set can be derived using the traditional LINDENBAUM lemma.

Finally, the truth lemma shows that $\varphi \in \Gamma \iff (M^c, \Gamma) \models \varphi$ by induction on φ and is enough to conclude (since the maximal consistent set containing $\neg\varphi$ will not verify φ). Most induction cases are the same as for **S5**, the only new symbols left in our formula φ are the E_a^d atoms, and the truth lemma is immediately true for them by definition of the depth function of M^c.

Finally, if $\models \varphi$, then $\models t(\varphi)$ by the soundness of the axiomatization and definition of the transformation, then **S5D** $\vdash t(\varphi)$ since we have just shown the completeness of this fragment. Finally, this must mean **DBEL** $\vdash t(\varphi)$ and then $\vdash \varphi$ since the transformations of t can be performed using equivalences in our axiomatization: we have shown completeness. □

Theorem A.2. *The axiomatization in Table 4 is sound and complete with respect to **EDPAL**.*

Proof. Soundness of the axioms of **DBEL** is proven in Theorem A.1. Soundness of all axioms for public announcement is also a consequence of their definition in PAL with which they share their definition, except for depth adjustment for which the proof is relatively immediate.

For completeness, we translate any formula φ into $t(\varphi)$ by removing public announcements, K_a modal operators and P_a^d atoms by using the sound axioms from Table 4. The formula $t(\varphi)$ is in the syntactic fragment **S5D**, thus we can use completeness shown in Theorem 2.1 to show $\vdash t(\varphi)$, which implies $\vdash \varphi$ within the axiomatization of Table 4 by using the same axioms in the opposite direction. □

B LINDENBAUM lemma with depth assignments

Lemma B.1. *For every agent a and finite consistent set of formulas* Γ *without public announcement,* P_b^d *literals or* K_b *operators for all b, there exists some* $d \in \mathbb{N}$ *such that* $\Gamma \cup \{E_a^d\}$ *is a consistent set.*

Proof. Fix agent a. As Γ is a finite set of finite formulas, the set of exact depth atoms for a that appear in its formulas is included in a finite set $F = \{E_a^0, \ldots, E_a^D\}$ for some $D \in \mathbb{N}$.

We can add to Γ instances of the unique depth axiom for each pair of integers in $[|0;D|]$ while maintaining consistency. The set Γ can then be seen as a consistent set of formulas for **S5** over the set of atoms $F \cup \mathscr{P}$, i.e. consistent in the axiomatization of Table 3 without depth axioms or bounded knowledge (or tautologies involving symbols not in the language of **S5**). Therefore there is an **S5** model (M,s) that satisfies it by the usual LINDENBAUM lemma and the truth lemma (the canonical model here).

In (M,s), if any of the E_a^d are valued to $\top$, then at most one of them is satisfied (since we added the unique depth axiom for all pair of depths). If all of the E_a^d are valued to $\bot$, then we can introduce a new atom E_a^{D+1} and set its value to $\top$ in all states of the model. All of the unique depths axioms for $D+1$ and $d \leq D$ can be added to Γ without making it inconsistent.

In both cases, let d_0 be the value of the unique $E_a^{d_0}$ valued to $\top$ in this final model. We claim that $\{\varphi, E_a^{d_0}\}$ must be a consistent set. Indeed, a proof of its inconsistency with the axioms from Table 3 must only involve axioms from **S5** and unique depths axioms for the set F, since none of the symbols P_a^d or K_a are necessary in a proof (they can be replaced by their equivalents with E_a^d and K_a^∞ without changing the conclusion) and any occurrence of E_a^d for $d > D+1$ can be replaced by $\bot$ while maintaining the truthfulness and conclusion of the proof.

Therefore, such an inconsistency proof would also hold within **S5**, which is a contradiction with soundness since these formulas are verified in a consistent set (the set of true formulas in (M,s)). □

C Proofs for Section 3

Proposition C.1. *Formulas* (KP) *and* (TA) *are valid for* ***DPAL*** *in the unambiguous depths setting.*

Proof. To prove (KP), suppose without loss of generality that $(M,s) \models \neg P_a^{d(\varphi)} \wedge \varphi$. In particular, this means that in $M \mid \varphi$, we have $(0,s) \sim_a' (1,s)$ and therefore the equivalence class of $(1,s)$ in $M \mid \varphi$ contains all $(0,s')$ whenever $s' \sim_a s$. Then,

$$
\begin{aligned}
(M,s) \models [\varphi]K_a\psi &\iff (M,s) \models \varphi \implies (M \mid \varphi,(1,s)) \models K_a\psi \\
&\iff (M \mid \varphi,(1,s)) \models P_a^{d(\psi)} \text{ and } \forall s', j, (j,s') \sim_a' (1,s) \implies (M \mid \varphi,(j,s')) \models \psi \\
&\iff (M \mid \varphi,(1,s)) \models P_a^{d(\psi)} \text{ and } \forall s' \sim_a s, \begin{cases} (M \mid \varphi,(0,s')) \models \psi \\ (M,s') \models \varphi \implies (M \mid \varphi,(1,s')) \models \psi. \end{cases}
\end{aligned} \tag{9}
$$

On the other hand,

$$
(M,s) \models K_a\psi \iff (M,s) \models P_a^{d(\psi)} \text{ and } \forall s', s' \sim_a s \implies (M,s') \models \psi. \tag{10}
$$

We prove by structural induction over $\psi \in \mathscr{H}_a$ the stronger equivalence,

$$
\forall s' \sim_a s, \quad \begin{cases} (M,s') \models \psi \iff (M \mid \varphi,(0,s')) \models \psi \\ (M,s') \models \varphi \implies ((M,s') \models \psi \iff (M \mid \varphi,(1,s')) \models \psi). \end{cases} \tag{11}
$$

Given that $(M,s) \not\models P_a^{d(\varphi)}$, we have $(M \mid \varphi,(1,s)) \models P_a^{d(\psi)} \iff (M,s) \models P_a^{d(\psi)}$. Therefore the depth conditions in equations (9) and (10) are the same and since both sides are true if $(M,s) \not\models \varphi$, equation (11) is enough to prove (KP).

For $\psi \in \mathscr{P}$, it is true because $V'((j,s')) = V(s')$ for all j and s' (note that $\mathscr{P}$ does not include depth atoms). For depth atoms about a, it is a consequence of $(M,s) \models K_a \neg P_a^{d(\varphi)}$ by the depth unambiguity condition (8), which means the depth of a is unchanged in all $s' \sim_a s$ after the model update.

The cases where $\psi = \psi_1 \wedge \psi_2$ and $\psi = \neg\chi$ are immediate, by the way these operators coincide with the usual propositional logic definition on both sides of the equivalences.

If $\psi = K_a\chi$ and $s' \sim_a s$, recall that by the depth unambiguity condition (8) we have $(M,s') \models \neg P_a^{d(\varphi)}$. Therefore, if $(M,s') \models \varphi$,

$$\begin{aligned}
(M \mid \varphi,(1,s')) \models \psi &\iff d(a,s') \geq d(\chi) \text{ and } \forall (j,s'') \sim'_a (1,s'),\ (M \mid \varphi,(j,s'')) \models \chi \\
&\iff d(a,s') \geq d(\chi) \text{ and } \forall s'' \sim_a s,\ \begin{cases} (M \mid \varphi,(0,s'')) \models \chi \\ (M,s'') \models \varphi \implies (M \mid \varphi,(1,s'')) \models \chi \end{cases} \\
&\iff d(a,s') \geq d(\chi) \text{ and } \forall s'' \sim_a s',\ (M,s'') \models \chi \\
&\iff (M,s') \models \psi,
\end{aligned}$$

where we have used the induction hypothesis (11) for χ once in each direction. The first equivalence in equation (11) is even easier to verify, by the same technique. The case for $\psi = K_a^\infty\chi$ is directly implied by this proof, as there are no depth conditions to verify.

To prove public announcements, we will need a stronger induction hypothesis than (11). Write for any s, $1_0(s) = s$ and $1_n(s) = (1, 1_{n-1}(s)) = (1,\ldots,(1,s))$. We posit,

$$\begin{aligned}
&\forall n \in \mathbb{N},\ \forall \psi_1,\ldots,\psi_n,\ \forall s' \sim_a s,\ (M,s') \models P_a^{d(\psi_1)+\cdots d(\psi_n)+d(\psi)} \text{ and } (M,s') \models \neg P_a^{d(\varphi)} \implies \\
&(M,s') \models \psi_1 \text{ and } (M \mid \psi_1,(1,s')) \models \psi_2 \text{ and } \ldots \text{ and } (M \mid \psi_1 \mid \cdots \mid \psi_{n-1}, 1_{n-1}(s')) \models \psi_n \implies \\
&\begin{cases} (M \mid \psi_1 \mid \cdots \mid \psi_n, 1_n(s')) \models \psi \iff (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n, 1_n((0,s'))) \models \psi \\ (M,s') \models \varphi \implies ((M \mid \psi_1 \mid \cdots \mid \psi_n, 1_n(s')) \models \psi \iff (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n, 1_n((1,s'))) \models \psi). \end{cases}
\end{aligned} \tag{12}$$

Note we slightly abuse notation here and some of these states might not exist, the convention is that the equivalences need only hold when the states exist in the models on both sides. The implicant implies that the left-hand term always exists.

Taking this for $n = 0$ is sufficient to conclude on (KP), since both equations (9) and (10) will be false whenever $(M,s) \not\models P_a^{d(\psi)}$.

The cases for atoms, negations and conjunction are clear for the same reasons as they were in equation (11). The case for depth atoms for a is direct, since the assumption $(M,s') \models P_a^{d(\psi_1)+\cdots d(\psi_n)}$ implies that the depth of a after the $\psi_1,\ldots,\psi_n$ announcements is its initial depth minus the sum of the depths of all the announcements, and the assumption that it is not deep enough for φ means its depth does not change with the announcement of φ.

The case for modal operators K_a relies on the fact that depth atoms are preserved (by the induction hypothesis for depth atoms) and the relations verify in $M \mid \psi_1 \mid \cdots \mid \psi_n$ when these states exist,

$$(1,(1,\ldots(1,s_1))) \sim_a^{(n)} (1,(1,\ldots(1,s_2))) \iff s_1 \sim_a s_2,$$

by denoting $\sim_a^{(k)}$ the relation for a in a model after k announcements. And similarly in $M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n$,

$$(1,(1,\dots(j,s_1))) \sim_a^{(n+1)} (1,(1,\dots(k,s_2))) \iff (j,s_1) \sim_a' (k,s_2) \iff s_1 \sim_a s_2.$$

This also implies the case for K_a^∞, since the verification is the same without the depth condition.

Finally, if $\psi = [\psi']\chi$, we verify that for $s' \sim_a s$ such that $(M,s') \models \varphi$,

$$\begin{aligned}
&(M \mid \psi_1 \mid \cdots \mid \psi_n, 1_n(s')) \models \psi \\
&\iff (M \mid \psi_1 \mid \cdots \mid \psi_n, 1_n(s')) \models \psi' \implies (M \mid \psi_1 \mid \cdots \mid \psi_n \mid \psi', 1_{n+1}(s')) \models \chi \\
&\iff (M \mid \psi_1 \mid \cdots \mid \psi_n, 1_n(s')) \models \psi' \implies (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n \mid \psi', 1_{n+1}((1,s'))) \models \chi \\
&\iff (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n, 1_n((1,s'))) \models \psi' \implies (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n \mid \psi', 1_{n+1}((1,s'))) \models \chi \\
&\iff (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n, 1_n((1,s'))) \models \psi
\end{aligned} \tag{13}$$

when the latter state exists. Our first use of the induction hypothesis on χ is justified because the left-hand side of the implication is the $n+1$ term in the assumptions for the induction hypothesis in equation (12) (and $d(\psi) = d(\psi') + d(\chi)$). The second use of the induction hypothesis on ψ' is justified for the same depth reason and the other assumptions remain the same. Once more the case for $(0,s')$ is very similar.

For (TA), we assume without loss of generality that $(M,s) \models K_a(P_a^{d(\varphi)}) \wedge \varphi$ (using the depth unambiguity condition (8)), this means in particular the equivalence class of $(1,s)$ in $M \mid \varphi$ is $\{(1,s'),\ s' \sim_a s,\ (M,s') \models \varphi\}$ since no state equivalent to s by $\sim_a$ has a not deep enough for φ. Using the same reasoning as in equation (9), we have,

$$(M,s) \models [\varphi]K_a\psi \iff (M \mid \varphi, (1,s)) \models P_a^{d(\psi)} \text{ and } \forall s' \sim_a s,\ (M,s') \models \varphi \implies (M \mid \varphi, (1,s')) \models \psi.$$

Moreover, we have,

$$(M,s) \models K_a[\varphi]\psi \iff (M,s) \models P_a^{d(\varphi)+d(\psi)} \text{ and } \forall s' \sim_a s,\ (M,s') \models \varphi \implies (M \mid \varphi, (1,s')) \models \psi. \tag{14}$$

Since $(M,s) \models P_a^{d(\varphi)}$, the depth of a in $(M \mid \varphi, (1,s))$ is its depth in (M,s) minus $d(\varphi)$. This means that

$$(M \mid \varphi, (1,s)) \models P_a^{d(\psi)} \iff (M,s) \models P_a^{d(\varphi)+d(\psi)}.$$ □

Proposition C.2. ***DPAL** verifies* (KP') *and* (TA').

Proof. For (KP'), in light of equations (9) and (10), we use the following induction hypothesis,

$$\begin{aligned}
&\forall s,a, \quad (M,s) \models K_a^\infty \mathscr{F}_\varphi(\psi) \implies \\
&\forall s' \sim_a s, \quad \begin{cases} (M \mid \varphi, (0,s')) \models \psi \iff (M,s') \models \psi \\ (M,s') \models \varphi \implies ((M \mid \varphi, (1,s')) \models \psi \iff (M,s') \models \psi). \end{cases}
\end{aligned} \tag{15}$$

Assume that $(M,s) \models \varphi \wedge \mathscr{F}_\varphi(K_a\psi)$. In particular, $(M,s) \models \neg K_a^\infty(\varphi \to P_a^{d(\varphi)})$. First notice that this condition allows us to write, $(0,s') \sim_a' (1,s) \iff s' \sim_a s$. Indeed, since there exists some $s'' \sim_a s$ where a is of depth strictly less than $d(\varphi)$ and φ holds, we deduce the chain of connections, $(1,s) \sim_a' (1,s'') \sim_a' (0,s'') \sim_a' (0,s')$ for any $s' \sim_a s$ (and the direct implication is immediate).

Moreover, we have assumed $(M,s) \models \varphi \wedge (\varphi \to \neg P_a^{d(\varphi)} \vee P_a^{d(\varphi)+d(\psi)})$. In either case of the disjunction, the depth conditions of equations (9) and (10) become equivalent as they did in the proof of (KP). Therefore, proving the induction hypothesis (15) is sufficient to conclude (KP') here.

The cases for atoms, negations and conjunctions is the same as in the proof of Proposition C.1, as the induction hypothesis holds because $\mathscr{F}_\varphi(\neg\psi) = \mathscr{F}_\varphi(\psi)$, $\mathscr{F}_\varphi(\psi_1 \wedge \psi_2) = \mathscr{F}_\varphi(\psi_1) \wedge \mathscr{F}_\varphi(\psi_2)$, and by commutativity of K_a^∞ with conjunction.

If $\psi = K_b\chi$ for some agent $b \in \mathscr{A}$, for some fixed $s' \sim_a s$, we know that $(M,s') \models \neg K_b^\infty(\varphi \to P_b^{d(\varphi)})$ as well as $(M,s') \models K_b^\infty \mathscr{F}_\varphi(\chi)$. Moreover, the condition $(M,s') \models K_b^\infty(\varphi \to \neg P_b^{d(\varphi)} \vee P_b^{d(\varphi)+d(\chi)})$ implies that the depth of b will be greater or equal to $d(\chi)$ in $(M \mid \varphi, (1,s'))$ if and only if it was in (M,s'). If $(M,s') \models \varphi$, by once more using the induction hypothesis (15) for b in s', we obtain that,

$$
\begin{aligned}
(M \mid \varphi,(1,s')) \models \psi &\iff d(b,s') \geq d(\chi) \text{ and } \forall (j,s'') \sim_b' (1,s'),\ (M \mid \varphi,(j,s'')) \models \chi \\
&\iff d(b,s') \geq d(\chi) \text{ and } \forall s'' \sim_b s',\ \begin{cases} (M \mid \varphi,(0,s'')) \models \chi \\ (M,s'') \models \varphi \implies (M \mid \varphi,(1,s'')) \models \chi \end{cases} \\
&\iff d(b,s') \geq d(\chi) \text{ and } \forall s'' \sim_b s',\ (M,s'') \models \chi \\
&\iff (M,s') \models \psi.
\end{aligned}
$$

The case for $(0,s')$ is the same, since its equivalence class in $M \mid \varphi$ is the same and the depth condition is the same. The case for $\psi = K_b^\infty\chi$ is implied by this proof, as there are no depth conditions to verify.

Finally, checking public announcements involves performing the same induction hypothesis strengthening as in the proof of (KP) in its equation (12). The new induction hypothesis becomes,

$$
\begin{aligned}
&\forall s,a, \quad (M,s) \models K_a^\infty \mathscr{F}_\varphi(\psi) \implies \\
&\forall n \in \mathbb{N},\ \forall \psi_1,\ldots,\psi_n,\ \forall s' \sim_a s,\ (M,s') \models P_a^{d(\psi_1)+\cdots d(\psi_n)+d(\psi)} \text{ and } (M,s') \models \neg P_a^{d(\varphi)} \implies \\
&(M,s') \models \psi_1 \text{ and } (M \mid \psi_1,(1,s')) \models \psi_2 \text{ and } \ldots \text{ and } (M \mid \psi_1 \mid \cdots \mid \psi_{n-1}, 1_{n-1}(s')) \models \psi_n \implies \\
&\begin{cases} (M \mid \psi_1 \mid \cdots \mid \psi_n, 1_n(s')) \models \psi \iff (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n, 1_n((0,s'))) \models \psi \\ (M,s') \models \varphi \implies ((M \mid \psi_1 \mid \cdots \mid \psi_n, 1_n(s')) \models \psi \iff (M \mid \varphi \mid \psi_1 \mid \cdots \mid \psi_n, 1_n((1,s'))) \models \psi). \end{cases}
\end{aligned}
$$

Note we slightly abuse notation here and some of these states might not exist, the convention is that the equivalences need only hold when the states exist in the models on both sides. The implicant implies that the left-hand term always exists.

Checking atoms, depth atoms, negation and conjunction is the same as in the proof of (KP) once more. Checking modal operators K_a and K_a^∞ is similar to the proof of (KP) using induction hypothesis (12), but using the same reasoning as above for induction hypothesis (15): the induction hypothesis contained in $\mathscr{F}_\varphi$ tells us that the announcement is not perceived by the agent at each modal operator.

Finally, public announcements follow the exact same proof as they did in (KP) in equation (13), with the extra information that $\mathscr{F}_\varphi([\psi']\chi) = \mathscr{F}_\varphi(\psi') \wedge \mathscr{F}_\varphi(\chi)$, allowing us to obtain the assumption of the inductive hypothesis in both inductive hypothesis applications (one for ψ' and one for χ).

For (TA'), we assume without loss of generality that $(M,s) \models K_a^\infty(\varphi \to P_a^{d(\varphi)}) \wedge \varphi$, this means in particular the equivalence class of $(1,s)$ in $M \mid \varphi$ is $\{(1,s'),\ s' \sim_a s,\ (M,s') \models \varphi\}$ since no state equivalent to s by $\sim_a$ has a not deep enough for φ. Using once more the same re-writings as in equation (14), it is sufficient to prove that the depth conditions are the same. This is the case because $(M,s) \models \varphi$, therefore by the truth axiom, $(M,s) \models P_a^{d(\varphi)}$. □