Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Dynamic Deontic Logic for Permitted Announcements

Xu Li[1] , **Dov Gabbay**[2,1] , **Réka Markovich**[1]
[1]University of Luxembourg
[2]King's College London
{xu.li, reka.markovich}@uni.lu, dov.gabbay@kcl.ac.uk

Abstract

In this paper, we introduce and study a dynamic deontic logic for permitted announcements. In our logic framework, it is permitted to announce something if announcing it would not lead to forbidden knowledge. It is shown that the logic is not compact, and we propose a sound and weakly complete Hilbert-style axiomatisation. We also study the computational complexity of the model checking problem and the decidability of the satisfiability problem. Finally, we introduce a neighbourhood semantics with a strongly complete axiomatisation.

1 Introduction

Deontic logic is an area of logic that investigates normative concepts like obligation, permission and prohibition (Gabbay et al. 2013). It has applications in many areas of computer science, including legal knowledge representation. Recently, there has been increasing interest in studying deontic logic about knowledge or belief (Başķent, Loohuis, and Parikh 2012; Markovich and Roy 2021a). But little attention has been paid to deontic logic about epistemic actions like public announcements. Balbiani and Seban (2011) proposed a logic framework to reason about "permission to say something". In our view, however, their semantics does not reflect the intuition behind permitted announcements (see Section 6). Inspired by Aucher et al. (2011), we have developed an alternative logic of permitted announcements, called LPA, by interpreting "permitted announcements" as "announcements that would not lead to forbidden knowledge". As we will see later, it follows the well known Andersonian tradition in deontic logic. Since "knowledge" and "permission" are both well established concepts in epistemic logic and deontic logic respectively, the crucial question is how to capture all the logical principles of permitted announcements defined in this style as well as other metalogical properties of the resulting logic. Specifically, we address the following research questions in this paper:

1. What is a suitable language for talking about permitted announcements? What is a suitable semantics for characterising permitted announcements?
2. How to capture all logical principles about permitted announcements? We need to define a proof system and show that all and only the validities are provable.
3. Is the proposed axiomatisation strongly complete with respect to the formal semantics? If not, can we find an alternative semantics with which our axiomatisation is strongly complete?
4. What is the computational complexity of the model checking problem in the resulting logic framework? We propose a model checking algorithm and analyse its complexity.
5. Is the satisfiability problem decidable or not?

The resulting logic, LPA, enables us to reason about privacy policies. Aucher, Boella and van der Torre (2011) proposed a dynamic modal logic to reason about which information is permitted to be sent by a security monitor to comply with a privacy policy. A privacy policy there is just a finite and consistent set of formulas. By replacing the purely syntactical privacy policy by the general notions of permitted and forbidden knowledge characterised by relational models, we can reason about privacy policies. We explain the details of this point in Section 7.

This paper is structured as follows. We provide the definitions, explanations and semantic results of LPA in the next section. We address the model checking problem of LPA in Section 3. A sound and weakly complete axiomatisation for LPA is introduced in Section 4, and we also show the decidability there. We provide an alternative neighbourhood-like semantics in Section 5. We compare our paper with related literature in Sections 6 and 7. Section 8 concludes the paper and indicates some future work.

2 Language and Semantics

Let PROP be a countable infinite set of propositional variables and let vio be a propositional constant such that $vio \notin$ PROP.

Definition 1 (Language)**.** *The* language $\mathcal{L}$ *is defined inductively by the following BNF grammar:*

$$\varphi ::= p \mid vio \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K\varphi \mid [\varphi]\varphi$$

where $p \in$ PROP. Let $\mathcal{L}_{el}$ be the language without inductive constructs vio and $[\varphi]\varphi$. We define $\mathbb{P}\varphi$ as an abbreviation for $\neg[\varphi]vio$. Other propositional connectives are defined in the usual way.

For every formula $\varphi \in \mathcal{L}$, let $sub(\varphi)$ be the set of all the subformulas of φ. For any set of formulas $\Gamma \subseteq \mathcal{L}$, we use $\textsc{prop}(\Gamma)$ to denote the set of propositional variables occurring in at least one formula in Γ. Given some $A \subseteq \textsc{prop}, \mathcal{L}_x(A)$ is the logical language $\mathcal{L}_x$ ($\mathcal{L}$, $\mathcal{L}_{el}$) restricted to propositional variables in A.

The intuitive reading of $K\varphi$ is "φ is known" and that of vio is "a violation has occurred". The formula $[\varphi]\psi$ is read as "after the announcement of φ, it holds that ψ". The formula $\mathbb{P}\varphi$ is an abbreviation for $\neg[\varphi]vio$, which states that "it is not the case that after the announcement of φ, a violation has occurred". Thus, $\mathbb{P}\varphi$ is intended to express that "it is permitted to announce φ".

Here are some brief remarks about our language. The use of the propositional constant vio dates back to Anderson (1958). But a more relevant reference is (Meyer 1987), where deontic logic about actions is reduced to a variant of dynamic logic by setting that $F\alpha = [\alpha]V$ and $P\alpha = \neg F\alpha$, where α is an action and $[\alpha]V$ means that "after the execution of α, one gets into trouble". Thus, an action α is forbidden (that is, $F\alpha$) if, after the execution of α, one gets into trouble; an action α is permitted (that is, $P\alpha$) if α is not forbidden. Concerning our paper, since public announcements are a type of action, we define $\mathbb{P}\varphi$ as $\neg[\varphi]vio$. [1]

We do not introduce operator O for obligation in our language. This is due to two considerations. First, we want to focus on our aim of developing a logic of permitted announcements. The current language is already sufficient for talking about permitted announcements. Second, adding obligation operator O to our language would allow us to express Åqvist's paradox, i.e. $OKp \to Op$ (Åqvist 1967).

In the following definition, we define the complexity of a formula, and it will be used in several proofs in our paper, including the completeness proof.

Definition 2 (Complexity). *For every formula $\alpha \in \mathcal{L}$, the* complexity *of α, with the notation $c(\alpha)$, is a positive integer inductively defined as follows: (1) $c(p) = c(vio) = 1$; (2) $c(\neg\varphi) = c(K\varphi) = 1 + c(\varphi)$; (3) $c(\varphi \wedge \psi) = 1 + \max(c(\varphi), c(\psi))$; (4) $c([\varphi]\psi) = (4 + c(\varphi)) \cdot c(\psi)$.*

The next lemma will be used in the completeness proof.

Lemma 3. *The following hold for every $\varphi, \psi, \chi \in \mathcal{L}$ (van Ditmarsch, van der Hoek, and Kooi 2008, page 187):*

1. $c(\psi) > c(\varphi)$ *if* $\varphi \in sub(\psi) \setminus \{\psi\}$
2. $c([\varphi]p) > c(\varphi \to p)$
3. $c([\varphi]\neg\psi) > c(\varphi \to \neg[\varphi]\psi)$
4. $c([\varphi](\psi \wedge \chi)) > c([\varphi]\psi \wedge [\varphi]\chi)$
5. $c([\varphi]K\psi) > c(\varphi \to K[\varphi]\psi)$
6. $c([\varphi][\psi]\chi) > c([\varphi \wedge [\varphi]\psi]\chi)$

Now we introduce the formal models for LPA.

Definition 4 (Models). *A* model *is a tuple $M = (W_0, W, \sim_0, R_0, V_0)$ such that:*

- *W_0 is a non-empty set called* the initial domain*;*
- *$W \subseteq W_0$ is a non-empty set called* the current domain*;*
- *$\sim_0 \subseteq W_0 \times W_0$ is an equivalence relation on W_0;*
- *$R_0 \subseteq W_0 \times W_0$ is a serial relation on W_0, i.e. for any $w \in W_0$, there is a $v \in W_0$ such that wR_0v;*
- *$V_0 : \textsc{prop} \to \wp(W_0)$ is a valuation.*

Given a model $M = (W_0, W, \sim_0, R_0, V_0)$, we define the initial model *of M as $M_0 = (W_0, W_0, \sim_0, R_0, V_0)$. For every state $w \in W_0$, we define $\sim_0(w) = \{v \in W_0 \mid w \sim_0 v\}$, $R_0(w) = \{v \in W_0 \mid wR_0v\}$ and $V_0(w) = \{p \in \textsc{prop} \mid w \in V_0(p)\}$. Finally, a* pointed model *is a pair M, w where M is a model and w is a state of M.*

Intuitively, if the equivalence relation $\sim_0$ holds between two states w, u, we mean that the two states are indistinguishable (to the agent we are modelling), whereas if the serial relation R_0 holds between w, u (i.e. wR_0u), it means that u is an ideal world for w. The states in W_0 are used to model the situation before any public announcement occurred. In contrast, the states in W are intended to represent the current situation where some public announcements may have occurred. Now we are ready to introduce the formal semantics for LPA.

Definition 5 (Semantics). *Given a model $M = (W_0, W, \sim_0, R_0, V_0)$, for every state $w \in W$ and formula $\alpha \in \mathcal{L}$, we inductively define that α is* satisfied *at w, with the notation $M, w \models \alpha$, as follows:*

$$
\begin{array}{rll}
M, w \models p & \text{iff} & w \in V_0(p) \\
M, w \models \neg\varphi & \text{iff} & \text{not } M, w \models \varphi \\
M, w \models \varphi \wedge \psi & \text{iff} & M, w \models \varphi \text{ and } M, w \models \psi \\
M, w \models K\varphi & \text{iff} & M, u \models \varphi \text{ for all } u \in W \text{ with } w \sim_0 u \\
M, w \models vio & \text{iff} & \text{there is } \varphi \in \mathcal{L}_{el} \text{ such that } M, w \models \varphi \text{ and } M_0, u \not\models \varphi \\
& & \text{for all } u \in W_0 \text{ with } wR_0u \; ^2 \\
M, w \models [\varphi]\psi & \text{iff} & M, w \models \varphi \text{ implies } M|_\varphi, w \models \psi
\end{array}
$$

where $M|_\varphi = (W_0, [\![\varphi]\!]_M, \sim_0, R_0, V_0)$ and $[\![\varphi]\!]_M = \{w \in W \mid M, w \models \varphi\}$. The notion of validity is defined as usual.

The semantics for all our operators are standard (Plaza 1989; Gerbrandy and Groeneveld 1997), except for the constant vio. The semantics for vio reflects the intuition that a violation occurred if there is a true statement which is also forbidden to be the case (i.e. false in all ideal worlds). Recall that we define $\mathbb{P}\varphi$ as an abbreviation for $\neg[\varphi]vio$. Thus, to decide whether it is permitted to announce φ, we need to check the normative status of all epistemic formulas (instead of φ itself) after the announcement. The reason is that the announcement of φ may lead the agent to infer some other fact ψ which itself is forbidden to be known by him/her. In the general case, we cannot decide which propositions are epistemically related for which agents, so a proper solution

[1] Here we define permitted announcements as weak permission (Dignum, Meyer, and Wieringa 1994a). One can also consider the version of strong permission, i.e., $\mathbb{P}^+\varphi := [\varphi]\neg vio$.

[2] One may argue that the notion of "violation" defined here is too general since the normative status of propositional formulas are also taken into account. Alternatively, one can let φ range over $\mathcal{L}_{el} \backslash \mathcal{L}_{prop}$. However, we will not consider the alternative definition in this paper. One reason is that all of our results in this paper also hold for the alternative definition.

is the introduction of existential quantification over formulas in the definition of vio.

In our semantics, we assume that the permission of statements like $p \to Kq$ does not change with public announcements, which is reflected in the fact that we always go back to the initial model M_0 to decide whether a statement is permitted or not. This assumption is plausible in many cases. For example, we are not permitted to know the medical data of others even if we have been told about them. We also recognise that there may be cases where permission to know does change with communications. We leave that for future work.

The following example illustrates our semantics.

Example 6. *Consider a company where each branch has its own website. The headquarters create a privacy policy according to which the branches are allowed to publish their employees' names* n *on their websites under the Colleagues' menu point but they are not allowed to indicate their email addresses* e *as the company doesn't want the public to know them. However, if someone knows how the company generates the email addresses, they can infer them by knowing the names. In this case, even if the public could know the names, the branch shouldn't publish them to be in compliance with the privacy policy.*

This scenario can be represented by the models M *and* $M|_n$ *in Figure 1. Model* M *characterises the public's initial knowledge and permission to know where the current state is* w. *We have* $M, w \models K(n \to e)$ *and* $M, w \models \neg vio$ *(since* wR_0w*). To characterise the situation after the announcement of* n*, we replace the current domain of* M *by* $[\![n]\!]_M = \{w, v, x\}$*, and the resulting model is* $M|_n$*. Since* $M|_n, w \models vio$ *(because* $M|_n, w \models Ke$ *and* $M_0, s \not\models Ke$ *for all* $s \in W_0$ *with* wR_0s*), we have* $M, w \models [n]vio$*, i.e.* $M, w \not\models \mathbb{P}n$*.*

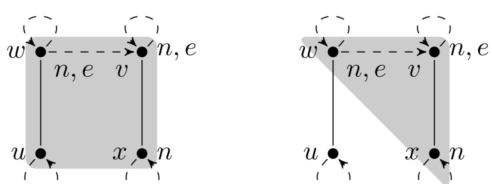

Figure 1: Two models $M = (W_0, W, \sim_0, R_0, V_0)$ (on the left) and $M|_n = (W_0, [\![n]\!]_M, R_0, V_0)$ (on the right). Worlds are nodes in the graph, and valuations are given by labelling the nodes with the true atoms. The current domains of each model consist of all the nodes in shaded areas. The epistemic relation $\sim_0$ and the deontic relation R_0 are represented by straight lines and dotted arrows respectively.

We list some semantic results in the next two propositions.

Proposition 7. *The following hold for every* $\varphi, \psi, \chi \in \mathcal{L}$*:*

1. $\models [\varphi]p \leftrightarrow (\varphi \to p)$
2. $\models [\varphi]\neg\psi \leftrightarrow (\varphi \to \neg[\varphi]\psi)$
3. $\models [\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$
4. $\models [\varphi]K\psi \leftrightarrow (\varphi \to K[\varphi]\psi)$
5. $\models [\varphi][\psi]\chi \leftrightarrow [\varphi \wedge [\varphi]\psi]\chi$
6. $\models vio \leftrightarrow [\top]vio$
7. $\models K(\varphi \leftrightarrow \psi) \to ([\varphi]vio \leftrightarrow [\psi]vio)$

Proof. We show the validity of only the last two items. Let $M = (W_0, W, \sim_0, R_0, V_0)$ be an arbitrary model and let $w \in W$.

6. Since $M, w \models \top$ and $M|_\top$ is identical to M, it is straightforward that $M, w \models vio \leftrightarrow [\top]vio$.
7. Suppose that $M, w \models K(\varphi \leftrightarrow \psi)$, we are supposed to show that $M, w \models [\varphi]vio \leftrightarrow [\psi]vio$. Without loss of generality, we only show $M, w \models [\varphi]vio \to [\psi]vio$. If $M, w \not\models \varphi$, then $M, w \not\models \psi$. By the semantics of operator $[\cdot]$, it is easy to see that $M, w \models [\varphi]vio$ and $M, w \models [\psi]vio$, thus $M, w \models [\varphi]vio \to [\psi]vio$. If $M, w \models \varphi$ (thus $M, w \models \psi$), we have $\sim_0(w) \cap [\![\varphi]\!]_M = \sim_0(w) \cap [\![\psi]\!]_M$. We can show the following claim by induction on the structure of χ:

   For every $\chi \in \mathcal{L}_{el}$ and $v \in \sim_0(w) \cap [\![\varphi]\!]_M$, $M|_\varphi, v \models \chi$ iff $M|_\psi, v \models \chi$.

   In particular, $M|_\varphi, w \models \chi$ iff $M|_\psi, w \models \chi$ for all $\chi \in \mathcal{L}_{el}$. If $M, w \models [\varphi]vio$, since $M, w \models \varphi$, then $M|_\varphi, w \models vio$. By the semantics, there must be a $\chi \in \mathcal{L}_{el}$ such that $M|_\varphi, w \models \chi$ and $M_0, v \not\models \chi$ for all $v \in R_0(w)$. Note that $M|_\psi, w \models \chi$ too, thus $M|_\psi, w \models vio$. Therefore $M, w \models [\psi]vio$ since $M, w \models \psi$. We have shown that $M, w \models [\varphi]vio \to [\psi]vio$. □

Proposition 8 (Properties of $\mathbb{P}$)**.** *The following properties hold:*

1. $\models \mathbb{P}\varphi \to \varphi$;
2. $\models [\varphi]\mathbb{P}\psi \leftrightarrow (\varphi \to \mathbb{P}(\varphi \wedge [\varphi]\psi))$;
3. *If* $\models \varphi \leftrightarrow \psi$*, then* $\models \mathbb{P}\varphi \leftrightarrow \mathbb{P}\psi$;
4. $\models \neg vio \to (K\varphi \to \mathbb{P}\varphi)$;
5. $\models vio \to \neg\mathbb{P}\top$;
6. $\models vio \to (K\varphi \to \neg\mathbb{P}\varphi)$;
7. *Even if* $\models \varphi$*, it is not necessarily the case that* $\models \mathbb{P}\varphi$;
8. $\not\models \mathbb{P}\varphi \wedge \mathbb{P}\psi \to \mathbb{P}(\varphi \wedge \psi)$;
9. $\not\models \mathbb{P}(\varphi \wedge \psi) \to \mathbb{P}\varphi$;
10. $\not\models \mathbb{P}\varphi \to \mathbb{P}\mathbb{P}\varphi$.

Remark 1. Item 5 and 6 are about the behaviour of our permitted announcement operator in a violation state. Item 5 states that we are not permitted to announce tautologies in a violation state and item 6 states that we are also not permitted to announce facts that have already been known by the agent. One may argue that these are counter-intuitive since announcing them would not change the epistemic situation of the agent. Thus they are harmless and should be permitted. We answer the question by considering two cases: (1) The violation occurred because the agent must know some facts but he/she does not know it currently. In this case, we should tell them the facts, rather than hiding them by saying some tautologies or other known facts. (2) The agent knows some facts that are forbidden to be known by him/her. In this case, we can not remedy the situation by saying something more. Thus we should keep silence to prevent him/her from inferring new information from our announcements. It can be seen as that we are modelling a careful sender.

Remark 2. Questions may also arise on the relevance between announcements and the violation. For example, in item 5, announcing the tautology has nothing to do with the actual violation, so why should it be prohibited? We only mention here that the similar problem has been identified in dynamic deontic logic, i.e., whether an executed action *causes* the actual violation? So we could solve our problem using a method similar to that in (Dignum, Meyer, and Wieringa 1994a).

The next proposition shows that LPA has no compactness, so there would not be a strongly complete axiomatisation for LPA (Blackburn, de Rijke, and Venema 2001).

Proposition 9. *LPA is not compact.*

Proof. Consider a set of formulas $\Gamma = \{[\neg vio]\chi \to \chi \mid \chi \in \mathcal{L}_{el}\} \cup \{\neg vio \wedge [\neg vio]vio\}$. It is clear that Γ is unsatisfiable. Next, we show that any finite subset Γ' of Γ is satisfiable. Without loss of generality, we assume that $\neg vio \wedge [\neg vio]vio \in \Gamma'$. Let p^* be a propositional variable not in $\textsc{prop}(\Gamma')$. Let $M = (W_0, W, \sim_0, R_0, V_0)$ be a model where:

- $W_0 = \{w, u, w', u'\}$;
- $W = \{w, u\}$;
- $\sim_0 = \{(w, w), (u, u), (w', w'), (u', u'), (w, u), (u, w), (w', u'), (u', w')\}$;
- $R_0 = \{(w, w'), (u, w'), (w', w'), (u', u')\}$;
- $V(p) = \varnothing$ for every $p \neq p^*$, $V(p^*) = \{u, u'\}$.

We illustrate M in Figure 2. We are going to show that $M, w \models \varphi$ for every $\varphi \in \Gamma'$. It is clear that (1) $M, w \models \neg vio$. Since $M, u \models p^*$ and $M_0, w' \not\models p^*$, then $M, u \models vio$. So $[\![\neg vio]\!]_M = \{w\}$, thus $M|_{\neg vio} = (W_0, \{w\}, \sim_0, R_0, V_0)$, as illustrated in Figure 2. Since $M|_{\neg vio}, w \models K\neg p^*$ and $M_0, w' \not\models K\neg p^*$, then $M|_{\neg vio}, w \models vio$, thus (2) $M, w \models [\neg vio]vio$. By induction on the structure of χ, we can prove the following claim:

For every $\chi \in \mathcal{L}_{el}(\textsc{prop}(\Gamma'))$ and $x \in \{w, u\}$, it is the case that $M, x \models \chi$ iff $M|_{\neg vio}, w \models \chi$. (†)

In particular, we have $M, w \models \chi$ iff $M|_{\neg vio}, w \models \chi$ for every $\chi \in \mathcal{L}_{el}(\textsc{prop}(\Gamma'))$. It follows that (3) $M, w \models [\neg vio]\chi \to \chi$ for every $[\neg vio]\chi \to \chi \in \Gamma'$. From (1) – (3), it follows that $M, w \models \varphi$ for every $\varphi \in \Gamma'$. $\square$

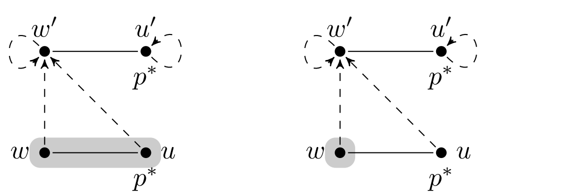

Figure 2: Two models M (on the left) and $M|_{\neg vio}$ (on the right).

3 Model Checking

Given a formula φ and a pointed model M, w, we want to check whether φ holds at M, w. This is the so-called model checking problem. It is not clear, at first sight, whether the problem is decidable in LPA, because the semantics of vio contains a quantification over all the formulas in $\mathcal{L}_{el}$. Thus, if we want to decide whether it is permitted to announce φ, it seems that we need to check all the epistemic formulas one by one. In this section, we show that if the number of propositional variables is restricted to a finite number, then we effectively resolve the "model checking" problem. More precisely, let $A \subset \textsc{prop}$ be a finite set, and we will show that the following decision problem, denoted as $\textsc{mc}(\mathcal{L}(A))$, is decidable.

Input: a finite model $M = (W_0, W, \sim_0, R_0, V_0)$ for $\mathcal{L}(A)$ (i.e. the domain of the valuation V_0 is A), a state $w \in W$, and a formula $\varphi \in \mathcal{L}(A)$,[3]
Output: determination of whether φ is satisfied at w in M.[4]

Below, we focus on the model checking of vio since the satisfaction of other operators can be checked as in standard public announcement logic (Kooi and van Benthem 2004). The key point here is that, when the number of states in the model and that of propositional variables in the language are both restricted to be finite, the constant vio is satisfied in the current state if and only if there is no ideal state that is bisimilar to the current state (as shown in Proposition 13).

Definition 10 (Largest bisimulation)**.** *Given two models $M = (W_0, W, \sim_0, R_0, V_0)$ and $M' = (W'_0, W', \sim'_0, R'_0, V'_0)$, we define the* largest bisimulation *between M and M' as a binary relation $\leftrightarroweq \subseteq W \times W'$ such that $w \leftrightarroweq v$ if:*

atoms $V_0(w) = V'_0(v)$.

forth *for all $x \in W$, if $w \sim_0 x$, then there is a $y \in W'$ such that $v \sim'_0 y$ and $V_0(x) = V'_0(y)$.*

back *for all $y \in W'$, if $v \sim'_0 y$, then there is a $x \in W$ such that $w \sim_0 x$ and $V'_0(y) = V_0(x)$.*

Definition 11 (Identity formula)**.** *Given a model $M = (W_0, W, \sim_0, R_0, V_0)$, for every $w \in W$, let $\textsc{val}(w)$ be the formula $(\bigwedge_{p \in V_0(w)} p) \wedge (\bigwedge_{p \notin V_0(w)} \neg p)$. The* identity formula *of w, denoted as $\textsc{idForm}(w)$, is defined as follows:*

$$\textsc{val}(w) \wedge K(\bigvee_{v \in W \cap \sim_0(w)} \textsc{val}(v)) \wedge \bigwedge_{v \in W \cap \sim_0(w)} \neg K \neg \textsc{val}(v)$$

Lemma 12. *Let two models $M = (W_0, W, \sim_0, R_0, V_0)$ and $M' = (W'_0, W', \sim'_0, R'_0, V'_0)$ be given and let the largest bisimulation between M and M' be $\leftrightarroweq$. For every $w \in W$ and $v \in W'$, we have*

[3]Without specifying explicitly, the "models" mentioned in this section would be finite models for $\mathcal{L}(A)$.

[4]Here, strictly speaking, the notion of satisfaction should be distinguished from that in Definition 5. That is, the truth definition of vio should be replaced by the following:

$M, w \models vio$ iff there is $\varphi \in \mathcal{L}_{el}(A)$ such that $M, w \models \varphi$ and $M_0, u \not\models \varphi$ for all $u \in W_0$ with wR_0u

1. *$w \leftrightarroweq v$ iff for every $\varphi \in \mathcal{L}_{el}(A)$, it is the case that $M, w \models \varphi$ iff $M', v \models \varphi$;*
2. *If $(w, v) \notin \leftrightarroweq$, then $M, w \models$ IDFORM(w) and $M', v \not\models$ IDFORM(w).*

Proof. We show only the first item. *From left to right*: Suppose that $w \leftrightarroweq v$. Note that, in S5, every formula $\varphi \in \mathcal{L}_{el}$ is equivalent to a formula of the modal degree ≤ 1 (Cresswell and Hughes 1996). Thus, we can prove the lemma by considering four cases: p, $K\alpha$ (α is a propositional formula), $\neg\varphi$ and $\varphi \wedge \psi$, all of which are straightforward. *From right to left*: Suppose that $(w, v) \notin \leftrightarroweq$. It is not hard to see that $M, w \models$ IDFORM(w) but $M', v \not\models$ IDFORM(w). □

Proposition 13. *Let $M = (W_0, W, \sim_0, R_0, V_0)$ be a model and let the largest bisimulation between M and M_0 be $\leftrightarroweq$. For every state $w \in W$, it holds that $M, w \models vio$ iff for all $u \in W_0$ with wR_0u, $(w, u) \notin \leftrightarroweq$.*

Proof. From left to right: Suppose $M, w \models vio$. Let $u \in W_0$ be an arbitrary state such that wR_0u. By the semantics of vio, there must be a formula $\varphi \in \mathcal{L}_{el}(A)$ such that $M, w \models \varphi$ and $M_0, u \not\models \varphi$. Therefore $(w, u) \notin \leftrightarroweq$ by item 1 of Lemma 12. *From right to left*: Since $(w, u) \notin \leftrightarroweq$, it follows from item 2 of Lemma 12 that $M, w \models$ IDFORM(w) and $M_0, u \not\models$ IDFORM(w). It is obvious that $M, w \models vio$. □

Now we can present a sound algorithm like that of Algorithm 1 to decide the truth of vio at some pointed model. A close look at Algorithm 1 delivers us the following theorem:

Theorem 14. MC($\mathcal{L}$ (A)) *is in $O(|W_0|^4 \times |A| \times |sub(\varphi)|)$.*

Algorithm 1 Model checking for vio

```
Input: A model M = (W_0, W, ~_0, R_0, V_0) for L(A)
Input: a state w ∈ W
Output "Yes" or "No"
  for each u ∈ R_0(w) do
      Check if ISBISIMILAR((M,w),(M_0,v)). If yes:
      return "No"
  end for
  return "Yes"

  procedure ISBISIMILAR((M,w),(M_0,v))
      if V_0(w) ≠ V_0(v) then
          return False
      else
          for each w' ∈ W ∩ ~_0(w) do
              Check if there is a v' ∈ ~_0(v) such that V_0(w') = V_0(v'). If not:
              return False
          end for
          for each v' ∈ ~_0(v) do
              Check if there is a w' ∈ W ∩ ~_0(w) such that V_0(w') = V_0(v'). If not:
              return False
          end for
      end if
      return True
  end procedure
```

4 Axiomatisation

In this section, we first define a Hilbert-style axiomatisation **LPA** for the language $\mathcal{L}$ and prove that it is sound and complete with respect to the class of all models. The general idea behind the completeness proof is to prove weak completeness by constructing a finite canonical model. This can be achieved by taking maximal consistent sets in the finite closure of one formula instead of the whole language $\mathcal{L}$. However, we do not have a modal operator to describe the relation R_0 in models, so the crucial step is how to construct the relation R_0 in the finite canonical model.

Definition 15 (Axiomatisation)**.** *The axiomatisation **LPA** for the language $\mathcal{L}$ is provided in Figure 3. A formula in $\mathcal{L}$ is a theorem of **LPA** if there is a proof for it. If φ is a theorem of **LPA**, we write $\vdash \varphi$.*

Axioms:	
(PL)	All propositional tautologies
(K)	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
(T)	$K\varphi \to \varphi$
(4)	$K\varphi \to KK\varphi$
(5)	$\neg K\varphi \to K\neg K\varphi$
(!Atom)	$[\varphi]p \leftrightarrow (\varphi \to p)$
(!Neg)	$[\varphi]\neg\psi \leftrightarrow (\varphi \to \neg[\varphi]\psi)$
(!Conj)	$[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$
(!K)	$[\varphi]K\psi \leftrightarrow (\varphi \to K[\varphi]\psi)$
(!Comp)	$[\varphi][\psi]\chi \leftrightarrow [\varphi \wedge [\varphi]\psi]\chi$
(vio)	$vio \leftrightarrow [\top]vio$
($\mathbb{P}$)	$K(\varphi \leftrightarrow \psi) \to ([\varphi]vio \leftrightarrow [\psi]vio)$
Rules:	
(MP)	from φ and $\varphi \to \psi$, infer ψ
(Nec)	from φ, infer $K\varphi$

Figure 3: The axiomatization **LPA**

Definition 16 (Closure)**.** *Given a formula $\varphi \in \mathcal{L}$, the* closure *of φ, with the notation $cl(\varphi)$, is the smallest set of formulas such that:*

- *$\varphi \in cl(\varphi)$;*
- *if $\psi \in cl(\varphi)$, then $sub(\psi) \subseteq cl(\varphi)$;*
- *if $\psi \in cl(\varphi)$ and ψ is not of the form $\neg\chi$, then $\neg\psi \in cl(\varphi)$;*
- *if $[\psi]p \in cl(\varphi)$, then $(\psi \to p) \in cl(\varphi)$;*
- *if $[\psi]\neg\chi \in cl(\varphi)$, then $(\psi \to \neg[\psi]\chi) \in cl(\varphi)$;*
- *if $[\psi](\chi \wedge \xi) \in cl(\varphi)$, then $[\psi]\chi \wedge [\psi]\xi \in cl(\varphi)$;*
- *if $[\psi]K\chi \in cl(\varphi)$, then $(\psi \to K[\psi]\chi) \in cl(\varphi)$;*
- *if $[\psi][\chi]\xi \in cl(\varphi)$, then $[\psi \wedge [\psi]\chi]\xi \in cl(\varphi)$;*
- *if $vio \in cl(\varphi)$, then $[\top]vio \in cl(\varphi)$;*
- *if $[\psi]vio, [\chi]vio \in cl(\varphi)$, then $K(\psi \leftrightarrow \chi) \in cl(\varphi)$.*

It is clear that $cl(\varphi)$ is finite for every formula φ. For any finite set of formulas Γ, let $\bigwedge\Gamma$ be the conjunction of all the formulas in Γ. The notion of "maximal consistent set" in a closure is defined as usual (van Ditmarsch, van der Hoek, and Kooi 2008). It is obvious that the set of all maximal consistent sets in $cl(\varphi)$ is finite for any formula φ.

Lemma 17. *Let φ be a formula. For any maximal consistent set of formulas Γ in $cl(\varphi)$, the following hold:*

1. *if $\vdash \bigwedge \Gamma \to \psi$ and $\psi \in cl(\varphi)$, then $\psi \in \Gamma$;*
2. *if $\psi \in cl(\varphi)$, then $\psi \in \Gamma$ iff $\neg\psi \notin \Gamma$;*
3. *if $\psi \wedge \chi \in cl(\varphi)$, then $\psi \wedge \chi \in \Gamma$ iff $\psi \in \Gamma$ and $\chi \in \Gamma$.*

In the remainder of this section, for convenience, we fix the formula $\varphi \in \mathcal{L}$. For every $\psi \in cl(\varphi)$, let $\|\psi\|$ be the set of all maximal consistent sets in $cl(\varphi)$ containing ψ. The Lindenbaum lemma can be proved in the standard way.

Lemma 18 (Lindenbaum)**.** *For every consistent set of formulas $\Sigma \subseteq cl(\varphi)$, there is a maximal consistent set Γ in $cl(\varphi)$ such that $\Sigma \subseteq \Gamma$.*

Now we define the notion of canonical pseudo-model, which is the main ingredient for constructing the canonical model in Definition 23.

Definition 19 (Canonical pseudo-model for $cl(\varphi)$)**.** *The* canonical pseudo-model *for $cl(\varphi)$ is a tuple $\mathsf{M} = (W, \sim, V)$ such that:*

- *W is the set of all maximal consistent sets in $cl(\varphi)$;*
- *$\sim \subseteq W \times W$ is such that $w \sim u$ iff $\{K\psi \mid K\psi \in w\} = \{K\psi \mid K\psi \in u\}$;*
- *$V : \text{PROP} \to \wp(W)$ is such that:*
  - *for all $p \in cl(\varphi)$, $w \in V(p)$ iff $p \in w$;*
  - *for all $w \in W$, there is a unique propositional variable $p_w \notin cl(\varphi)$ such that $V(p_w) = \{w\}$;*
  - *$V(p) = \varnothing$ for any other propositional variable.*

It is clear from the definition that $\sim$ is an equivalence relation. Note that V is well defined since W is finite. The following lemma can be shown in the standard way.

Lemma 20 (Existence)**.** *For every $w \in W$ and $K\psi \in cl(\varphi)$, if $K\psi \notin w$, then there must be a $u \in W$ such that $w \sim u$ and $\psi \notin u$.*

Lemma 21. *For every $w \in W$ and formulas $[\psi]vio$,$[\xi]vio \in cl(\varphi)$, if $[\psi]vio \notin w$ and $[\xi]vio \in w$, then $\sim(w) \cap \|\psi\| \neq \sim(w) \cap \|\xi\|$.*

Proof. Suppose $[\psi]vio \notin w$ and $[\xi]vio \in w$, then $K(\psi \leftrightarrow \xi) \notin w$. Otherwise, we would have $\vdash \bigwedge w \to ([\psi]vio \leftrightarrow [\xi]vio)$ by the axiom ($\mathbb{P}$) and we would have $\vdash \bigwedge w \to (\neg[\psi]vio \wedge [\xi]vio)$ by item 2 of Lemma 17, thereby contradicting the consistency of w.

Since $K(\psi \leftrightarrow \xi) \notin w$ and $K(\psi \leftrightarrow \xi) \in cl(\varphi)$, then by Lemma 20, it follows that there must be an $u \in W$ such that $w \sim u$ and $\psi \leftrightarrow \xi \notin u$.

Note that $\psi \in cl(\varphi)$. Then, by item 2 of Lemma 17, it follows that either $\psi \in u$ or $\neg\psi \in u$ (but not both). If $\psi \in u$, then $\xi \notin u$, thus $\sim(w) \cap \|\psi\| \neq \sim(w) \cap \|\xi\|$. If $\neg\psi \in u$, then $\psi \notin u$ and $\xi \in u$ (otherwise $\neg\xi \in u$ by item 2 of Lemma 17, thus $\psi \leftrightarrow \chi \in u$. Contradiction!). It also follows that $\sim(w) \cap \|\psi\| \neq \sim(w) \cap \|\xi\|$. ☐

Let $\psi_1, \psi_2, \ldots, \psi_n$ be an enumeration of all formulas ψ such that $[\psi]vio \in cl(\varphi)$ and there is a $w \in W$ with $[\psi]vio \notin w$. Next, we define the notion of "updates to the canonical pseudo-model", which also provides the ingredients for constructing the canonical model.

Definition 22 (Updates to the canonical pseudo-model)**.** *For each $1 \leq i \leq n$, we define the* update *to the canonical pseudo-model by ψ_i as a tuple $\mathsf{M}^i = (W^i, \sim^i, V^i)$ such that:*

- *$W^i = \|\psi_i\| \times \{i\}$;*
- *$\sim^i \subseteq W^i \times W^i$ is such that $(w, i) \sim^i (u, i)$ iff $w \sim u$;*
- *$V^i(p) = (V(p) \times \{i\}) \cap W^i$.*

Note that $\sim^i$ is an equivalence relation over W^i. Now, we are ready to define the canonical model for $cl(\varphi)$.

Definition 23 (Canonical model for $cl(\varphi)$)**.** *The* canonical model *$\mathcal{M}^c = (W_0^c, W^c, \sim_0^c, R_0^c, V_0^c)$ for $cl(\varphi)$ is defined as follows:*

- *$W_0^c = W \cup \bigcup_{1 \leq i \leq n} W^i \cup \{0\}$;*
- *$W^c = W$;*
- *$\sim_0^c = \sim \cup (\bigcup_{1 \leq i \leq n} \sim^i) \cup \{(0, 0)\}$;*
- *$R_0^c(w) = \{(w, i) \mid 1 \leq i \leq n$ and $[\psi_i]vio \notin w\} \cup \{0\}$ if $w \in W$,*
  $R_0^c(w) = \{w\}$ otherwise;
- *$V_0^c(p) = V(p) \cup \bigcup_{1 \leq i \leq n} V^i(p)$.*

We can verify that $\mathcal{M}^c$ is indeed a model. In the remainder of the section, let $[\![\varphi]\!] = [\![\varphi]\!]_{\mathcal{M}^c}$.

Lemma 24. *Let $1 \leq i \leq n$ and suppose that $\|\psi_i\| = [\![\psi_i]\!]$. Then for every $w \in W^c$ with $\mathcal{M}^c, w \models \psi_i$, it is the case that $\mathcal{M}^c|_{\psi_i}, w \models \chi$ iff $\mathcal{M}_0^c, (w, i) \models \chi$ for every $\chi \in \mathcal{L}_{el}$.*

Proof. Induction on the structure of χ. In the inductive case $K\xi$, note that $\sim^i((w, i)) = (\sim(w) \cap [\![\psi_i]\!]) \times \{i\}$ since we assume that $\|\psi_i\| = [\![\psi_i]\!]$. ☐

Lemma 25. *For every $\xi \in cl(\varphi)$ and $w \in [\![\xi]\!]$, if $\sim(w) \cap [\![\xi]\!] \neq \sim(w) \cap \|\psi_i\|$ for every $1 \leq i \leq n$ with $[\psi_i]vio \notin w$, then $\mathcal{M}^c|_\xi, w \models vio$.*

Proof. Let

$$\gamma = K(\bigvee_{u \in \sim(w) \cap [\![\xi]\!]} p_u) \wedge \bigwedge_{u \in \sim(w) \cap [\![\xi]\!]} \neg K \neg p_u.$$

It is clear that $\mathcal{M}^c|_\xi, w \models \gamma$. To show that $\mathcal{M}^c|_\xi, w \models vio$, it suffices to show that $\mathcal{M}_0^c, u \not\models \gamma$ for all $u \in R_0^c(w)$. If $u = 0$, it is clear that $\mathcal{M}_0^c, 0 \not\models \gamma$. If $u \neq 0$, then by the definition of R_0^c, there must be a $1 \leq i \leq n$ such that $[\psi_i]vio \notin w$ and $u = (w, i)$. By our assumption, we have $\sim(w) \cap \|\psi_i\| \neq \sim(w) \cap [\![\xi]\!]$.

If $\sim(w) \cap \|\psi_i\| \not\subseteq \sim(w) \cap [\![\xi]\!]$, there must be a $v \in \sim(w)$ such that $v \in \|\psi_i\|$ and $v \notin [\![\xi]\!]$. Note that, by the definitions of $\sim_0^c$ and $\sim^i$, we have that $\sim_0^c((w, i)) = \sim^i((w, i)) = (\sim(w) \cap \|\psi_i\|) \times \{i\}$. Thus $(v, i) \in \sim_0^c((w, i))$. Since $v \notin [\![\xi]\!]$, we have that $\mathcal{M}_0^c, (w, i) \models \neg K \neg(\bigwedge_{u \in \sim(w) \cap [\![\xi]\!]} \neg p_u)$ by the definition of V_0^c. Thus, $\mathcal{M}_0^c, (w, i) \not\models \gamma$.

If $\sim(w) \cap \|\psi_i\| \not\supseteq \sim(w) \cap [\![\xi]\!]$, it follows that there is a $v \in \sim(w)$ such that $v \in [\![\xi]\!]$ and $v \notin \|\psi_i\|$. Note that, by the definitions of $\sim_0^c$ and $\sim^i$, we have that $\sim_0^c((w, i)) = \sim^i((w, i)) = (\sim(w) \cap \|\psi_i\|) \times \{i\}$. Thus $(v, i) \notin \sim_0^c((w, i))$. Therefore $\mathcal{M}_0^c, (w, i) \not\models \neg K \neg p_v$ by the definition of V_0^c. Thus $\mathcal{M}_0^c, (w, i) \not\models \gamma$. ☐

Lemma 26 (Truth)**.** *For every formula $\psi \in cl(\varphi)$, $\|\psi\| = [\![\psi]\!]$.*

Proof. Induction on $c(\psi)$:

Case p. For every $p \in cl(\varphi)$ and $w \in W^c$, we have that $p \in w$ iff $w \in V(p)$ iff $w \in V_0^c(p)$.

Case vio. *From left to right*: Suppose $vio \in w$. Since $[\top]vio \in cl(\varphi)$, then $[\top]vio \in w$ by the axiom (vio) and item 1 of Lemma 17. For every $1 \leq i \leq n$ such that $[\psi_i]vio \notin w$, since $[\psi_i]vio, [\top]vio \in cl(\varphi)$, $[\psi_i]vio \notin w$ and $[\top]vio \in w$, it follows from Lemma 21 that $\sim(w) \cap \|\psi_i\| \neq \sim(w) \cap \|\top\|$. Since $[\![\top]\!] = \|\top\| = W$, then $\sim(w) \cap \|\psi_i\| \neq \sim(w) \cap [\![\top]\!]$. Therefore $\mathcal{M}^c|_\top, w \models vio$ by Lemma 25. That is $\mathcal{M}^c, w \models vio$. *From right to left*: Suppose $vio \notin w$ and $\mathcal{M}^c, w \models vio$. Then $[\top]vio \notin w$ by the axiom (vio). Thus, there must be a $1 \leq i \leq n$ such that $\psi_i = \top$. By the definition of R_0^c, it must be the case that $wR_0^c(w, i)$. Note that $\|\top\| = [\![\top]\!] = W$, thus, by Lemma 24, we have

$$\mathcal{M}^c|_\top, w \models \chi \text{ iff } \mathcal{M}_0^c, (w, i) \models \chi \text{ for every } \chi \in \mathcal{L}_{el}.$$

But since we assume that $\mathcal{M}^c, w \models vio$ (thus $\mathcal{M}^c, w \models [\top]vio$ by item 6 of Proposition 7), then $\mathcal{M}^c|_\top, w \models vio$. By semantics, there must be a $\chi \in \mathcal{L}_{el}$ such that $\mathcal{M}^c|_\top, w \models \chi$ and $\mathcal{M}_0^c, (w, i) \not\models \chi$. Contradiction!

The cases for $\neg$, $\wedge$ and K are all straightforward and the cases for $[\xi]p, [\xi]\neg\xi', [\xi](\xi_1 \wedge \xi_2), [\xi]K\xi'$, and $[\xi][\xi_1]\xi_2$ follow directly from the induction hypothesis, Proposition 7, and the axioms (!Atom) – (!Comp). The case for $[\xi]vio$ remains to be considered.

From left to right: Suppose $[\xi]vio \in w$. If $\xi \notin w$, by the induction hypothesis, $\mathcal{M}^c, w \not\models \xi$. Thus $\mathcal{M}^c, w \models [\xi]vio$ by semantics. Therefore, we assume that $\xi \in w$. By the induction hypothesis, this implies that $\mathcal{M}^c, w \models \xi$. To prove that $\mathcal{M}^c, w \models [\xi]vio$, it suffices to show that $\mathcal{M}^c|_\xi, w \models vio$. For every $1 \leq i \leq n$ with $[\psi_i]vio \notin w$, since $[\psi_i]vio, [\xi]vio \in cl(\varphi)$, $[\psi_i]vio \notin w$ and $[\xi]vio \in w$, it follows from Lemma 21 that $\sim(w) \cap \|\psi_i\| \neq \sim(w) \cap \|\xi\|$. By the induction hypothesis that $\|\xi\| = [\![\xi]\!]$, we have that $\sim(w) \cap \|\psi_i\| \neq \sim(w) \cap [\![\xi]\!]$. Now we can apply Lemma 25 and obtain $\mathcal{M}^c|_\xi, w \models vio$. Therefore, $\mathcal{M}^c, w \models [\xi]vio$. *From right to left*: Suppose $[\xi]vio \notin w$ (thus $\xi \in w$) and $\mathcal{M}^c, w \models [\xi]vio$. It follows that there must be a $1 \leq i \leq n$ such that $\psi_i = \xi$. By the definition of R_0^c, it must be the case that $wR_0^c(w, i)$. Note that, by the induction hypothesis, we have $\|\xi\| = [\![\xi]\!]$ and $\mathcal{M}^c, w \models \xi$. Thus, by Lemma 24, we have

$$\mathcal{M}^c|_\xi, w \models \chi \text{ iff } \mathcal{M}_0^c, (w, i) \models \chi \text{ for every } \chi \in \mathcal{L}_{el}.$$

But since we assume $\mathcal{M}^c, w \models [\xi]vio$, then $\mathcal{M}^c|_\xi, w \models vio$. By semantics, there must be a $\chi \in \mathcal{L}_{el}$ such that $\mathcal{M}^c|_\xi, w \models \chi$ and $\mathcal{M}_0^c, (w, i) \not\models \chi$. Contradiction! □

Theorem 27 (Soundness and completeness)**.** *For every $\varphi \in \mathcal{L}$, $\models \varphi$ iff $\vdash \varphi$.*

Proof. It is not hard to see that all the axioms and rules of standard epistemic logic S5 are valid. This, together with Proposition 7, implies that **LPA** is sound. For the completeness, suppose $\nvdash \varphi$. Then $\neg\varphi$ is consistent. Thus, by Lemma 18, there must be a maximal consistent set Γ in $cl(\varphi)$ such that $\neg\varphi \in \Gamma$. By Lemma 26, it follows that $\mathcal{M}^c, \Gamma \models \neg\varphi$. Thus, $\not\models \varphi$. □

Since the canonical model we constructed is finite, the next lemma follows immediately.

Lemma 28 (Finite model)**.** *For any formula $\varphi \in \mathcal{L}$, if φ is satisfiable, then φ must be satisfied in a finite model.*

Theorem 29 (Decidability)**.** *The satisfiability problem for LPA is decidable.*

5 A Neighbourhood-like Semantics

We have proven that **LPA** is weakly complete with respect to the semantics in Definition 5. However, by Proposition 9, **LPA** is not strongly complete, and we can never find such an axiomatisation. From a logical point of view, it would be interesting to find an alternative semantics with which **LPA** is strongly complete. In this section, we propose such a semantics for LPA in which the permitted announcements $\mathbb{P}\varphi$ undergoes a neighbourhood characterisation.

Definition 30 (Neighbourhood epistemic models)**.** *A* neighbourhood epistemic model*, or simply* n-model*, is a tuple $\mathsf{M} = (W, \sim, N, V)$ such that:*

- *W is a non-empty set;*
- *$\sim \subseteq W \times W$ is an equivalence relation on W;*
- *$N : W \to \wp(\wp(W))$ is a neighbourhood function;*[5]
- *$V : \text{PROP} \to \wp(W)$ is a valuation.*

Definition 31 (n-model semantics)**.** *Given an n-model $\mathsf{M} = (W, \sim, N, V)$, for every state $w \in W$ and formula $\alpha \in \mathcal{L}$, we inductively define that $\mathsf{M}, w \models^n \alpha$ as follows:*

$$\begin{array}{rcl}
\mathsf{M}, w \models^n p & \textit{iff} & w \in V(p) \\
\mathsf{M}, w \models^n vio & \textit{iff} & \{v \in W \mid w \sim v\} \notin N(w) \\
\mathsf{M}, w \models^n \neg\varphi & \textit{iff} & \textit{not } \mathsf{M}, w \models^n \varphi \\
\mathsf{M}, w \models^n (\varphi \wedge \psi) & \textit{iff} & \mathsf{M}, w \models^n \varphi \textit{ and } \mathsf{M}, w \models^n \psi \\
\mathsf{M}, w \models^n K\varphi & \textit{iff} & \textit{for all } u \in W \textit{ with } w \sim u, \\
 & & \mathsf{M}, u \models^n \varphi \\
\mathsf{M}, w \models^n [\varphi]\psi & \textit{iff} & \mathsf{M}, w \models^n \varphi \textit{ implies} \\
 & & \mathsf{M}|_\varphi, w \models^n \psi
\end{array}$$

where $\mathsf{M}|_\varphi = (W_\varphi, \sim_\varphi, N_\varphi, V_\varphi)$ is defined as follows:

- *$W_\varphi = [\![\varphi]\!]_\mathsf{M}$;*
- *$\sim_\varphi = \sim \cap ([\![\varphi]\!]_\mathsf{M} \times [\![\varphi]\!]_\mathsf{M})$;*
- *$N_\varphi(w) = \{U \mid U \in N(w) \,\&\, U \subseteq [\![\varphi]\!]_\mathsf{M}\}$;*
- *$V_\varphi(p) = V(p) \cap [\![\varphi]\!]_\mathsf{M}$.*

Thus, in the n-model semantics, we have $\mathsf{M}, w \models^n \mathbb{P}\varphi$ iff

$$\mathsf{M}, w \models^n \varphi \text{ and } [\![\varphi]\!]_\mathsf{M} \cap \{v \in W \mid w \sim v\} \in N(w)$$

However, the relationship between the n-model semantics and the semantics in Definition 5 is not so straightforward.

Proposition 32. *For any model $M = (W_0, W, \sim_0, R_0, V_0)$ and $w \in W$, there is an n-model M and state w' in M such that $M, w \models \varphi$ iff $\mathsf{M}, w' \models \varphi$ for all formulas $\varphi \in \mathcal{L}$.*

[5]For every $w \in W$, each set in the neighbourhood $N(w)$ induces an equivalence relation over W. Thus, intuitively, the function N can be understood as specifying which epistemic situations are "good" at each state.

Sketch of Proof. We can construct the model $\mathsf{M} = (W', \sim', N', V')$ such that:

- $W' = W$, $\sim' = \sim_0 \cap (W \times W)$, $V'(p) = V_0(p) \cap W$;
- $N'(w) = \{[\![\psi]\!]_M \cap \sim_0(w) \mid \psi \in \mathcal{L} \;\&\; M, w \models \mathbb{P}\psi\}$;
- $w' = w$.

It is not hard to prove the proposition by induction on the structure of φ.

But the converse of the above proposition does not hold. That is, there is an n-model for which we can not find an equivalent model. Since it is not hard to find an n-model satisfying the set Γ introduced in the proof of Proposition 9, but Γ is unsatisfiable in our original semantics. The exact relationship between the n-model semantics and the original semantics is given by the following proposition.

Proposition 33. *For every $\varphi \in \mathcal{L}$, φ is satisfiable (under the semantics in Definition 5) iff φ is satisfied by an n-model.*

We can also show the following completeness theorem by the canonical model method.

Theorem 34. *The axiomatisation **LPA** is sound and strongly complete with respect to the n-model semantics.*

6 The Principles of Permitted Announcements

The work of Balbiani and Seban (2011) is the first attempt to study the notion of "permission to say something". Since our work is devoted to the same topic, we compare our paper with their work in a separate section.

In Balbiani and Seban's (2011) article, the operators $P(\psi, \varphi)$ and $O(\psi, \varphi)$ are introduced to express the notions that "after announcing ψ, it is permitted/obligatory to announce φ". It is clear that our operator $\mathbb{P}\varphi$ can be expressed in their framework as $P(\top, \varphi)$. Conversely, one can also interpret their operator $P(\psi, \varphi)$ as $[\psi]\mathbb{P}\varphi$ in LPA. We do not have an operator $\mathbb{O}\varphi$ to express that "it is obligatory to announce φ", which is clearly a direction for future work. However, the semantics for their logic is quite different from ours. In their framework, a ternary relation $\mathcal{P} \subseteq S \times \wp(S) \times \wp(S)$ is used to provide the semantics for $P(\psi, \varphi)$ in such a way that $\mathcal{M}, s \models P(\psi, \varphi)$ iff for some $(s, [\![\psi]\!], S'') \in \mathcal{P}, S'' \subseteq [\![\varphi]\!]$, where S is the domain of the model and $[\![\varphi]\!]$ is the truth set of φ. We observe that those permitted announcements are directly encoded in the ternary relation $\mathcal{P}$. This encoding is purely technical and serves only to distinguish between what is permitted and what is not permitted. It does not reflect the semantic intuition behind the permitted announcements. In contrast, our work gives an intuitively appealing interpretation of "permitted announcements". We also find different logical principles pertaining to permitted announcements.

The prominent principle of permitted announcements that is absent from the work of Balbiani and Seban (2011) but present in ours is the axiom ($\mathbb{P}$): $K(\varphi \leftrightarrow \psi) \to (\mathbb{P}\varphi \leftrightarrow \mathbb{P}\psi)$. It states that if two statements are known by the agent to be equivalent, then it is impossible that we are permitted to announce one statement to the agent while being forbidden to announce the other statement, because the actions of announcing φ and announcing ψ would bring the same set of knowledge in this case. From ($\mathbb{P}$), we can infer that $\mathbb{P}\top \to (K\varphi \to \mathbb{P}\varphi)$. It states that we are permitted to announce things that have already been known if the current situation is good. This is a trivial principle, but it is not valid in the work of Balbiani and Seban (2011).

Another missing principle is $[\varphi]\mathbb{P}\psi \leftrightarrow (\varphi \to \mathbb{P}(\varphi \wedge [\varphi]\psi))$, which is equivalent to $\langle\varphi\rangle\mathbb{P}\psi \leftrightarrow \mathbb{P}\langle\varphi\rangle\psi$, where $\langle\varphi\rangle\psi = \neg[\varphi]\neg\psi$. It can be understood as saying that "it is permitted to make an announcement that 'φ is true and after the announcement of φ, ψ is true' if and only if φ is true and it is permitted to announce ψ after φ has been announced". One may notice that it is a variant of a well-known property pertaining to sequel actions in dynamic deontic logic: $P(\alpha_1; \alpha_2) \leftrightarrow \langle\alpha_1\rangle P\alpha_2$ (Meyer 1987).

LPA also invalidates some principles pertaining to permitted announcements in the work of Balbiani and Seban (2011). One is the weakening of a permitted announcement: $\mathbb{P}(\varphi \wedge \psi) \to \mathbb{P}\varphi$.[6] This is because the announcement of φ may provide less information than the announcement of $\varphi \wedge \psi$. However, sometimes ignorance about some facts may be forbidden or, in other words, knowledge of these facts is obligatory. For example, in many countries, it is not permitted for a pharmacist to only advertise over-the-counter medicines to customers while not telling them the possible side-effects.

7 Related Work

Deontic logic about knowledge and belief. The notions of permission and prohibition to know have been studied by Cuppens and Demolombe (1996) in the context of security of databases. In their framework, the modalities $PK\varphi$ and $FK\varphi$ are introduced to express that "some users are permitted/forbidden to know that the database believes φ".[7] There are two binary relations in their models, one for knowledge and one for obligation. Thus $PK\varphi$ is true at some state w if there is some ideal world of w in which $K\varphi$ holds. LPA implicitly employs the same characterisation of permission to know in its truth definition for vio.

An important paradox of deontic logic about knowledge is Åqvist's paradox (Åqvist 1967). Our work is not aimed at solving that paradox. Our concern is how the change of knowledge brought by public announcements affects the normative status of the agent concerned. Various work are devoted to the study of how knowledge affects obligation, e.g. (Horty 2001; Pacuit, Parikh, and Cogan 2006; Grossi et al. 2021). In our work, we assume that knowledge does not affect permission to know. It would be interesting to extend our work such that permission to know interacts with knowledge. Recently, there have also been some work studying other deontic concepts of knowledge, like "the right to know" (Markovich and Roy 2021a; Markovich and Roy 2021b). Our work concerns the permission to say something. We plan to extend our work to

[6] We point out that this principle does not hold in the earlier work of Balbiani et al. (2009).

[7] We reformulate their notations for a clear presentation.

studying the notion of "having the right to say" and its role in the formalisation of freedom of speech.

Deontic logic about epistemic actions. We are heavily influenced by the work of Aucher et al. (2011) which proposes a dynamic modal logic for reasoning about which information is permitted to be sent by a security monitor in order to comply with a privacy policy. We improve their logic in the following respects: 1. A privacy policy in their logic is just a finite and consistent set of formulas (like OKq or $Kp \to PKq$) specifying the permitted and obligatory knowledge of the receiver. By encoding the privacy policy into relational models, we can reason about the privacy policy, rather than just check the permitted and forbidden knowledge according to a purely syntactic list. 2. We simplify the semantics. In their logic, they use the machinery of action models in dynamic epistemic logic (Baltag and Moss 2004). Though dynamic epistemic logic is powerful enough, we show that we do not need it if we are only dealing with permitted announcements. Our semantics is much simpler and the number of states will not grow exponentially with announcements. The latter is important for model checking. The idea of distinguishing between the initial and current domains in relational models can be found in the work of Baltag et al. (2018). 3. We also enhance their formal language so that we can talk about high-order expressions about permitted announcements like $\mathbb{P}\mathbb{P}\varphi$. But a problem follows from our approach. Aucher et al. (2011) can express the notion of compliance as a single compound formula. Thus, they easily get a complete axiomatisation for their logic. However, the constant vio in LPA can not be reduced. Thus, it is not so easy to propose a complete axiomatisation. Meanwhile, we also think that our axiomatisation is more revealing.

Van Benthem et al. (2009) propose a logic for protocols in dynamic epistmemic logic. We can interpret it as a logic for permitted epistemic actions. Due to the space limitation, we only mention here that the permitted epistemic actions (represented by the event models of dynamic epistemic logic) are listed in the protocols, and there is no the notion of "permitted knowledge" in their framework. Besides, our work can easily be extended to other epistemic actions by adding the vio constant to the language of dynamic epistemic logic (Baltag and Moss 2004). Van Ditmarsch and Seban (2012) extend their prior work (Balbiani and Seban 2011) to the multi-agent case. We can also study how to extend our work to the multi-agent case.

Dynamic deontic logic. Our work has a close relationship to dynamic deontic logic (Meyer 1987). In dynamic deontic logic, an action is permitted if performing it would not lead to getting into trouble. Thus, permissions about actions are reduced to the normative status after executing the actions. Our work can be seen as instantiating actions by public announcements. Compared to (Meyer 1987), our work has three differences: 1. The truth of vio depends on the structure of the model rather than the propositional truth values of a state. 2. Our language allows high-order expressions like $\mathbb{P}\mathbb{P}\varphi$. 3. There are no compositions of public announcements like the negation of a public announcement. In the later developments of dynamic deontic logic, different definitions of permission have been proposed (Dignum, Meyer, and Wieringa 1994a). For example, the strong permission $P_{Strong}(\alpha)$ is defined as $[\alpha]\neg V$. Similarly, we can also study different variants of permitted announcements.

Van der Meyden et al. (1996) criticise Meyer's dynamic deontic logic because it contains a problematic validity: $\langle\alpha\rangle P\beta \to P(\alpha;\beta)$. The idea is that the action α may be a forbidden action. Thus, as per their example, even if we are permitted to remain silent after shooting the president, we are not permitted to shoot the president and remain silent. In LPA, we have an analogy: $\langle\varphi\rangle\mathbb{P}\psi \to \mathbb{P}\langle\varphi\rangle\psi$. However, what falls into the scope of $\mathbb{P}$ in the consequent of the implication is not a *sequel announcement* but a *single announcement* of formula $\langle\varphi\rangle\psi$. The free choice paradox (Ross 1941; Dignum, Meyer, and Wieringa 1994b) is another paradox concerning dynamic deontic logic. But, again, there is no notion of choice between two announcements in LPA.

The notion of "the negation of an action" in dynamic deontic logic has also been doubted by several authors (Broersen 2004; Wansing 2004; Sun and Huimin 2014; Ju and van Eijck 2016). Our work suggests the importance of a proper definition for the negation of public announcements.

8 Conclusions and Future Work

In this paper, we introduced a logic of permitted announcements by interpreting "permitted announcements" as "announcements that would not lead to forbidden knowledge". The language of LPA is obtained by augmenting public announcement logic with a constant vio, and the semantics for permitted announcements contains a quantification over all epistemic formulas. We captured all the logical principles of permitted announcements by the axiomatisation **LPA** and provided an alternative neighbourhood-like semantics with which **LPA** is strongly complete. It should be pointed out that the axiom ($\mathbb{P}$) would not be valid if assume a weaker logic for the knowledge operator K, e.g., the modal logic $\mathbf{K}$. We conjecture that our technique can be adapted to show the completeness results when the base epistemic logic is stronger than $\mathbf{K4}$. We also studied the computational complexity of the model checking problem and showed the decidability of the satisfiability problem. The computational complexity of the satisfiability problem for LPA is still unknown.

Our work gives new insights into the notion of "permitted announcements" and enables us to reason about privacy policies and other forms of regulated communication. Many issues are left for future research. The vital issue is how to incorporate the notion of "obligatory announcements". When we provided our semantics in this work, we assumed that permission to know does not change with public announcements, but it is natural to consider how to enable our framework to reason about the dynamics of permission to know. We can also consider, for instance, how to extend our framework to the multi-agent case, and how to incorporate epistemic actions other than public announcements.

Acknowledgments

We thank Leendert van der Torre for valuable discussions and comments on earlier drafts of this paper. We also thank three anonymous reviewers of KR 2022 for insightful remarks and Llio Humphreys for proofreading the paper. This work was supported by the Fonds National de la Recherche Luxembourg through the project Deontic Logic for Epistemic Rights (OPEN O20/14776480).

References

Anderson, A. R. 1958. A reduction of deontic logic to alethic modal logic. *Mind* 67(265):100–103.

Aucher, G.; Boella, G.; and van der Torre, L. 2011. A dynamic logic for privacy compliance. *Artificial Intelligence and Law* 19(2):187.

Balbiani, P., and Seban, P. 2011. Reasoning about permitted announcements. *Journal of Philosophical Logic* 40(4):445–472.

Balbiani, P.; van Ditmarsch, H.; and Seban, P. 2009. Reasoning about permitted announcements. In Herzig, A., and Lorini, E., eds., *(Electronic) proceedings of: ESSLLI 2009 workshop logical methods for social concepts, Bordeaux, France*.

Baltag, A., and Moss, L. S. 2004. Logics for epistemic programs. *Synthese* 139(2):165–224.

Baltag, A.; Özgün, A.; and Vargas Sandoval, A. L. 2018. APAL with memory is better. In Moss, L. S.; de Queiroz, R.; and Martinez, M., eds., *Logic, Language, Information, and Computation*, 106–129. Berlin, Heidelberg: Springer Berlin Heidelberg.

Başķent, C.; Loohuis, L. O.; and Parikh, R. 2012. On knowledge and obligation. *Episteme* 9(2):171–188.

Blackburn, P.; de Rijke, M.; and Venema, Y. 2001. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press.

Broersen, J. 2004. Action negation and alternative reductions for dynamic deontic logics. *Journal of Applied Logic* 2(1):153–168. The Sixth International Workshop on Deontic Logic in Computer Science.

Cresswell, M. J., and Hughes, G. E. 1996. *A new introduction to modal logic*. Routledge.

Cuppens, F., and Demolombe, R. 1996. A deontic logic for reasoning about confidentiality. In Brown, M. A., and Carmo, J., eds., *Deontic Logic, Agency and Normative Systems*, 66–79. London: Springer London.

Dignum, F.; Meyer, J.-J. C.; and Wieringa, R. 1994a. A dynamic logic for reasoning about sub-ideal states. In *ECAI'94 Workshop "Artificial Normative Reasoning"*.

Dignum, F.; Meyer, J.-J. Ch..; and Wieringa, R. J. 1994b. Contextual permission: A solution to the free choice paradox. In *Proceedings of the 2nd International Workshop on Deontic Logic in Computer Science (DEON 1994), Oslo, Norway, January 6-8, 1994*, 107–135.

Gabbay, D.; Horty, J.; Parent, X.; van der Meyden, R.; and van der Torre, L., eds. 2013. *Handbook of Deontic Logic and Normative Systems*. College Publications.

Gerbrandy, J., and Groeneveld, W. 1997. Reasoning about information change. *Journal of Logic, Language, and Information* 6(2):147–169.

Grossi, D.; Kooi, B.; Su, X.; and Verbrugge, R. 2021. How knowledge triggers obligation. In Ghosh, S., and Icard, T., eds., *Logic, Rationality, and Interaction*, 201–215. Cham: Springer International Publishing.

Horty, J. F. 2001. *Agency and Deontic Logic*. Oxford University Press.

Ju, F., and van Eijck, J. 2016. To do something else. In *Deontic Logic and Normative Systems. 13th International Conference, DEON 2016, Bayreuth, Germany, July 18-21, 2018*, 109–122.

Kooi, B., and van Benthem, J. 2004. Reduction axioms for epistemic actions. In Schmidt, R. A.; Pratt-Hartmann, I.; Reynolds, M.; and Wansing, H., eds., *AiML-2004: Advances in Modal Logic (Preliminary Proceedings)*, 197–211.

Markovich, R., and Roy, O. 2021a. Formalizing the right to know: Epistemic rights as normative positions. In *The First International Workshop on Logics for New-Generation Artificial Intelligence (LNGAI 2021)*.

Markovich, R., and Roy, O. 2021b. A logical analysis of freedom of thought. In *Deontic Logic and Normative Systems. 14th International Conference, DEON 2018, The Netherlands, 3-6 July, 2018*, 323–338.

Meyer, J.-J. C. 1987. A different approach to deontic logic: deontic logic viewed as a variant of dynamic logic. *Notre Dame Journal of Formal Logic* 29(1):109 – 136.

Pacuit, E.; Parikh, R.; and Cogan, E. 2006. The logic of knowledge based obligation. *Synthese* 149(2):311–341.

Plaza, J. 1989. Logics of public communications. In Emrich, M.; Pfeifer, M.; Hadzikadic, M.; and Ras, Z., eds., *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, 201–216.

Ross, A. 1941. Imperatives and logic. *Theoria* 7:53–71.

Sun, X., and Huimin, D. 2014. Stratified action negation for dynamic logic. In *26th BENELUX Conference on Artificial Intelligence (BNAIC)*.

van Benthem, J.; Gerbrandy, J.; Hoshi, T.; and Pacuit, E. 2009. Merging frameworks for interaction. *Journal of Philosophical Logic* 38(5):491–526.

van der Meyden, R. 1996. The Dynamic Logic of Permission. *Journal of Logic and Computation* 6(3):465–479.

van Ditmarsch, H., and Seban, P. 2012. A logical framework for individual permissions on public announcements. https://personal.us.es/hvd/newerpubs/hanspablo070212.pdf. Accessed: 2022-05-15.

van Ditmarsch, H.; van der Hoek, W.; and Kooi, B. 2008. *Dynamic Epistemic Logic*. Springer, Dordrecht.

Wansing, H. 2004. *On the Negation of Action Types: Constructive Concurrent PDL.* Dresden Preprints in theoretical philosophy and philosophical logic. Techn. Univ.

Åqvist, L. 1967. Good samaritans, contrary-to-duty imperatives, and epistemic obligations. *Noûs* 1(4):361–379.