

Knowing Values and Public Inspection

Jan van Eijck^{1,2}, Malvin Gattinger¹, Yanjing Wang³

¹ ILLC, University of Amsterdam, Amsterdam, The Netherlands

² SEN1, CWI, Amsterdam, The Netherlands

³ Department of Philosophy, Peking University, Beijing, China

Abstract. We present a basic dynamic epistemic logic of “knowing the value”. Analogous to public announcement in standard DEL, we study “public inspection”, a new dynamic operator which updates the agents’ knowledge about the values of constants. We provide a sound and strongly complete axiomatization for the single and multi-agent case, making use of the well-known Armstrong axioms for dependencies in databases.

Keywords: Knowing what, Bisimulation, Public Announcement Logic.

1 Introduction

Standard epistemic logic studies propositional knowledge expressed by “knowing that”. However, in everyday life we talk about knowledge in many other ways, such as “knowing what the password is”, “knowing how to swim”, “knowing why he was late” and so on. Recently the epistemic logics of such expressions are drawing more and more attention (see [1] for a survey).

Merely reasoning about static knowledge is important but it is also interesting to study the changes of knowledge. Dynamic Epistemic Logic (DEL) is an important tool for this, which handles how knowledge (and belief) is updated by events or actions [2]. For example, extending standard epistemic logic, one can update the propositional knowledge of agents by making propositional announcements. They are nicely studied by public announcement logic [3] which includes reduction axioms to completely describe the interplay of “knowing that” and “announcing that”. Given this, we can also ask: What are natural dynamic counterparts the knowledge expressed by other expressions such as knowing what, knowing how etc.? How do we formalize “announcing what”?

In this paper, we study a basic dynamic operation which updates the knowledge of the values of certain constants.⁴ The action of *public inspection* is the knowing value counterpart of public announcement and we will see that it fits well with the logic of knowing value. As an example, we may use a sensor to measure the current temperature of the room. It is reasonable to say that after using the sensor you will know the temperature of the room. Note that it is not

⁴ In this paper, by *constant* we mean something which has a single value given the actual situation. The range of possible values of a constant may be infinite. This terminology is motivated by first-order modal logic as it will become more clear later.

reasonable to encode this by standard public announcement since it may result in a possibly infinite formula: $[t = 27.1^\circ\text{C}]K(t = 27.1^\circ\text{C}) \wedge [t = 27.2^\circ\text{C}]K(t = 27.2^\circ\text{C}) \wedge \dots$, and the inspection action itself may require an infinite action model in the standard DEL framework of [4] with a separate event for each possible value. Hence public inspection can be viewed as a public announcement of the actual value, but new techniques are required to express it formally. In our simple framework we define knowing and inspecting values as primitive operators, leaving the actual values out of our logical language.

The notions of knowing and inspecting values have a natural connection with dependencies in databases. This will also play a crucial role in the later technical development of the paper. In particular, our completeness proofs employ the famous set of axioms from [5]. For now, consider the following example.

Example 1. Suppose a university course was evaluated using anonymous questionnaires which besides an assessment for the teacher also asked the students for their main subject. See Table 1 for the results. Now suppose a student tells you, the teacher, that his major is Computer Science. Then clearly you know how that student assessed the course, since there is some dependency between the two columns. More precisely, in the cases of students 3 and 4, telling you the value of “Subject” effectively also tells you the value of “Assessment”. In practice, a better questionnaire would only ask for combinations of questions that do not allow the identification of students.

Student	Subject	Assessment
1	Mathematics	good
2	Mathematics	very good
3	Logic	good
4	Computer Science	bad

Table 1. Evaluation Results

Other examples abound: The author of [6] gives an account of how easily so-called ‘de-identified data’ produced from medical records could be ‘re-identified’, by matching patient names to publicly available health data.

These examples illustrate that reasoning about knowledge of values in isolation, i.e. separated from knowledge *that*, is both possible and informative. It is such knowledge and its dynamics that we will study here.

2 Existing Work

Our work relates to a collection of papers on epistemic logics with other operators than the standard “knowing that” $K\varphi$. In particular we are interested in the Kv operator expressing that an agent knows a value of a variable or constant. This operator is already mentioned in the seminal work [3] which introduced public announcement logic (PAL). However, a complete axiomatization of PAL together with Kv was only given in [7,8] using the relativized operator $Kv(\varphi, c)$ for the single and multi-agent cases. Moreover, it has been shown in [9] that by

treating the negation of Kv as a primitive diamond-like operator, the logic can be seen as a normal modal logic in disguise with binary modalities.

Inspired by a talk partly based on an earlier version of this paper, Baltag proposed the very expressive Logic of Epistemic Dependency (LED) [10], where knowing that, knowing value, announcing that, announcing value can all be encoded in a general language which also includes equalities like $c = 4$ to facilitate the axiomatization.

In this paper we go in the other direction: Instead of extending the standard PAL framework with Kv , we study it in isolation together with its dynamic counterpart $[c]$ for public inspection. In general, the motto of our work here is to see how far one can get in formalizing knowledge and inspection of values without going all the way to or even beyond PAL. In particular we do not include values in the syntax and we do not have any nested epistemic modalities.

As one would expect, our simple language is accompanied by simpler models and also the proofs are less complicated than existing methods. Still we consider our Public Inspection Logic (PIL) more than a toy logic. Our completeness proof includes a novel construction which we call “canonical dependency graph” (Definition 6). We also establish the precise connection between our axioms and the Armstrong axioms widely used in database theory [5].

Table 2 shows how PIL fits into the family of existing languages. Note that [10] is the most expressive language in which all operators are encoded using $K_i^{t_1, \dots, t_n} t$ which expresses that given the current values of t_1 to t_n , agent i knows the value of t . Moreover, to obtain a complete proof system for LED one also needs to include equality and rigid constants in the language. It is thus an open question to find axiomatizations for a language between PIL and LED without equality.

PAL	p	$K\varphi$			$[!\varphi]\varphi$	[3]
PAL + Kv	p	$K\varphi$	$Kv(c)$		$[!\varphi]\varphi$	[3]
PAL + Kv^r	p	$K\varphi$	$Kv(c)$	$Kv(\varphi, c)$	$[!\varphi]\varphi$	[7,8,9]
PIL			$Kv(c)$		$[c]\varphi$	this paper
PIL + K		$K\varphi$	$Kv(c)$		$[c]\varphi$	future work
LED	p	$K\varphi$	$Kv(c)$	$Kv(\varphi, c)$	$[c]\varphi$ $[!\varphi]\varphi$ $c = c$	[10]

Table 2. Comparison of Languages

All languages include the standard boolean operators \top , \neg and \wedge which we do not list in Table 2.

We also discuss other related works not in this line at the end of the paper.

3 Single-Agent PIL

We first consider a simple single-agent language to talk about knowing and inspecting values. Throughout the paper we assume a fixed set of constants \mathbb{C} .

Definition 1 (Syntax). Let c range over \mathbb{C} . The language \mathcal{L}_1 is given by:

$$\varphi ::= \top \mid \neg\varphi \mid \varphi \wedge \psi \mid Kv(c) \mid [c]\varphi$$

Besides standard interpretations of the boolean connectives, the intended meanings are as follows: $Kv(c)$ reads “the agent knows the value of c ” and the formula $[c]\varphi$ is meant to say “after revealing the actual value of c , φ is the case”. We also use the standard abbreviations $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$ and $\varphi \rightarrow \psi := \neg\varphi \vee \psi$.

Definition 2 (Models and Semantics). A model for \mathcal{L}_1 is a tuple $\mathcal{M} = \langle S, \mathcal{D}, V \rangle$ where S is a non-empty set of worlds (also called states), \mathcal{D} is a non-empty domain and V is a valuation $V : (S \times \mathbb{C}) \rightarrow \mathcal{D}$. To denote $V(s, c) = V(t, c)$, i.e. that c has the same value at s and t according to V , we write $s =_c t$. If this holds for all $c \in C \subseteq \mathbb{C}$ we write $s =_C t$. The semantics are as follows:

$\mathcal{M}, s \models \top$	always
$\mathcal{M}, s \models \neg\varphi$	$\Leftrightarrow \mathcal{M}, s \not\models \varphi$
$\mathcal{M}, s \models \varphi \wedge \psi$	$\Leftrightarrow \mathcal{M}, s \models \varphi$ and $\mathcal{M}, s \models \psi$
$\mathcal{M}, s \models Kv(c)$	\Leftrightarrow for all $t \in S : s =_c t$
$\mathcal{M}, s \models [c]\varphi$	$\Leftrightarrow \mathcal{M} _c^s, s \models \varphi$

where $\mathcal{M}|_c^s$ is $\langle S', \mathcal{D}, V|_{S' \times \mathbb{C}} \rangle$ with $S' = \{t \in S \mid s =_c t\}$. If for a set of formulas Γ and a formula φ we have that whenever a model \mathcal{M} and a state s satisfy $\mathcal{M}, s \models \Gamma$ then they also satisfy $\mathcal{M}, s \models \varphi$, then we say that φ follows semantically from Γ and write $\Gamma \models \varphi$. If this hold for $\Gamma = \emptyset$ we say that φ is semantically valid and write $\models \varphi$.

Note that the actual state s plays an important role in the last clause of our semantics: Public inspection of c at s reveals the *local actual* value of c to the agent. The model is restricted to those worlds which agree on c with s . This is different from PAL and other DEL variants based on action models, where updates are usually defined on models directly and not on pointed models.

We employ the usual abbreviation $\langle c \rangle \varphi$ as $\neg[c]\neg\varphi$. Note however, that public inspection of c can always take place and is deterministic. Hence the determinacy axiom $\langle c \rangle \varphi \leftrightarrow [c]\varphi$ is semantically valid and we include it in the following system.

Definition 3. The proof system SPIL_1 for PIL in the language \mathcal{L}_1 consists of the following axiom schemata and rules. If a formula φ is provable from a set of premises Γ we write $\Gamma \vdash \varphi$. If this holds for $\Gamma = \emptyset$ we also write $\vdash \varphi$.

Axiom Schemata

<i>TAUT</i>	<i>all instances of propositional tautologies</i>
<i>DIST</i>	$[c](\varphi \rightarrow \psi) \rightarrow ([c]\varphi \rightarrow [c]\psi)$
<i>LEARN</i>	$[c]Kv(c)$
<i>NF</i>	$Kv(c) \rightarrow [d]Kv(c)$
<i>DET</i>	$\langle c \rangle \varphi \leftrightarrow [c]\varphi$
<i>COMM</i>	$[c][d]\varphi \leftrightarrow [d][c]\varphi$
<i>IR</i>	$Kv(c) \rightarrow ([c]\varphi \rightarrow \varphi)$

Rules

<i>MP</i>	$\frac{\varphi, \varphi \rightarrow \psi}{\psi}$
<i>NEC</i>	$\frac{\varphi}{[c]\varphi}$

Intuitively, LEARN captures the effect of the inspection; NF says that the agent does not forget; DET says that inspection is deterministic; COMM says that inspections commute; finally, IR expresses that inspection does not bring any new information if the value is known already. Note that DET says that $[c]$ is a function. It also implies seriality which we list in the following Lemma.

Lemma 1. *The following schemes are provable in SPILL_1 :*

- $\langle c \rangle \top$ (*seriality*)
- $Kv(c) \rightarrow (\varphi \rightarrow [c]\varphi)$ (*IR'*)
- $[c](\varphi \wedge \psi) \leftrightarrow [c]\varphi \wedge [c]\psi$ (*DIST'*)
- $[c_1] \dots [c_n](\varphi \rightarrow \psi) \rightarrow ([c_1] \dots [c_n]\varphi \rightarrow [c_1] \dots [c_n]\psi)$ (*multi-DIST*)
- $[c_1] \dots [c_n](\varphi \wedge \psi) \leftrightarrow [c_1] \dots [c_n]\varphi \wedge [c_1] \dots [c_n]\psi$ (*multi-DIST'*)
- $[c_1] \dots [c_n](Kv(c_1) \wedge \dots \wedge Kv(c_n))$ (*multi-LEARN*)
- $(Kv(c_1) \wedge \dots \wedge Kv(c_n)) \rightarrow [d_1] \dots [d_n](Kv(c_1) \wedge \dots \wedge Kv(c_n))$ (*multi-NF*)
- $(Kv(c_1) \wedge \dots \wedge Kv(c_n)) \rightarrow ([c_1] \dots [c_n]\varphi \rightarrow \varphi)$ (*multi-IR*)

Moreover, the multi-NEC rule is admissible: If $\vdash \varphi$, then $\vdash [c_1] \dots [c_n]\varphi$.

Proof. For reasons of space we only prove three of the items and leave the others as an exercise for the reader. For IR', we use DET and TAUT:

$$\frac{\frac{\frac{}{Kv(c) \rightarrow ([c]\neg\varphi \rightarrow \neg\varphi)}{(IR)}}{Kv(c) \rightarrow (\neg[c]\varphi \rightarrow \neg\varphi)}{(DET)}}{Kv(c) \rightarrow (\varphi \rightarrow [c]\varphi)}{(TAUT)}$$

To show multi-NEC, we use DIST, NEC and TAUT. For simplicity, consider the case where $C = \{c_1, c_2\}$.

$$\frac{\frac{\frac{\frac{}{[c_2](\varphi \rightarrow \psi) \rightarrow ([c_2]\varphi \rightarrow [c_2]\psi)}{(DIST)}}{[c_1]([c_2](\varphi \rightarrow \psi) \rightarrow ([c_2]\varphi \rightarrow [c_2]\psi))}{(NEC)}}{[c_1][c_2](\varphi \rightarrow \psi) \rightarrow [c_1]([c_2]\varphi \rightarrow [c_2]\psi)}{(DIST, TAUT)}}{[c_1][c_2](\varphi \rightarrow \psi) \rightarrow ([c_1][c_2]\varphi \rightarrow [c_1][c_2]\psi)}{(DIST, TAUT)}$$

For multi-LEARN, we use LEARN, NEC, COMM, DIST' and TAUT:

$$\begin{array}{c}
\frac{}{[c_1]Kv(c_1)} \text{ (LEARN)} \\
\frac{[c_2][c_1]Kv(c_1)}{[c_1][c_2]Kv(c_1)} \text{ (NEC)} \quad \frac{[c_2]Kv(c_2)}{[c_1][c_2]Kv(c_2)} \text{ (LEARN)} \\
\frac{[c_1][c_2]Kv(c_1)}{[c_1][c_2]Kv(c_1)} \text{ (COMM)} \quad \frac{[c_1][c_2]Kv(c_2)}{[c_1][c_2]Kv(c_2)} \text{ (NEC)} \\
\frac{[c_1]([c_2]Kv(c_1) \wedge [c_2]Kv(c_2))}{[c_1][c_2](Kv(c_1) \wedge Kv(c_2))} \text{ (DIST', TAUT)} \\
\frac{[c_1]([c_2]Kv(c_1) \wedge [c_2]Kv(c_2))}{[c_1][c_2](Kv(c_1) \wedge Kv(c_2))} \text{ (DIST', TAUT)}
\end{array}$$

Definition 4. We use the following abbreviations for any two finite sets of constants $C = \{c_1, \dots, c_m\}$ and $D = \{d_1, \dots, d_n\}$.

- $Kv(C) := Kv(c_1) \wedge \dots \wedge Kv(c_m)$
- $[C]\varphi := [c_1] \dots [c_m]\varphi$
- $Kv(C, D) := [C]Kv(D)$.

Note that by multi-DIST' and COMM the exact enumeration of C and D in Definition 4 do not matter modulo logical equivalence.

In particular, these abbreviations allow us to shorten the “multi” items from Lemma 1 to $Kv(C, C)$, $Kv(C) \rightarrow Kv(D, C)$ and $Kv(C) \rightarrow ([C]\varphi \rightarrow \varphi)$. The abbreviation $Kv(C, D)$ allows us to define dependencies and it will be crucial in our completeness proof. We have that:

$$\overline{\mathcal{M}, s \models Kv(C, D) \Leftrightarrow \text{for all } t \in S : \text{if } s =_C t \text{ then } s =_D t}$$

Definition 5. Let \mathcal{L}_2 be the language given by $\varphi ::= \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid Kv(C, C)$.

Note that this language is essentially a fragment of \mathcal{L}_1 due to the above abbreviation, where (possibly multiple) $[c]$ operators only occur in front of Kv operators (or conjunctions thereof). Moreover, the next Lemma might count as a small surprise.

Lemma 2. \mathcal{L}_1 and \mathcal{L}_2 are equally expressive.

Proof. As $Kv(\cdot, \cdot)$ was just defined as an abbreviation, we already know that \mathcal{L}_1 is at least as expressive as \mathcal{L}_2 : we have $\mathcal{L}_2 \subseteq \mathcal{L}_1$. We can also translate in the other direction by pushing all sensing operators through negations and conjunctions. Formally, let $t : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ be defined by

$$\begin{array}{lll}
Kv(d) \mapsto Kv(\emptyset, \{d\}) & [c]\neg\varphi & \mapsto \neg t([c]\varphi) \\
\neg\varphi \mapsto \neg t(\varphi) & [c](\varphi \wedge \psi) & \mapsto t([c]\varphi) \wedge t([c]\psi) \\
\varphi \wedge \psi \mapsto t(\varphi) \wedge t(\psi) & [c]\top & \mapsto \top \\
& [c_1] \dots [c_n]Kv(d) & \mapsto Kv(\{c_1, \dots, c_n\}, \{d\})
\end{array}$$

Note that this translation preserves and reflects truth because determinacy and distribution are valid (determinacy allows us to push $[c]$ through negations, distribution to push $[c]$ through conjunctions). At this stage we have not yet established completeness, but determinacy is also an axiom. Hence we can note separately that $\varphi \leftrightarrow t(\varphi)$ is provable and that t preserves and reflects provability and consistency.

Example 2. Note that the translation of $[c]\varphi$ formulas also depends on the top connective within φ . For example we have

$$\begin{aligned} t([c](\neg K_V(d) \wedge [e]K_V(f))) &= t([c]\neg K_V(d)) \wedge t([c][e]K_V(f)) \\ &= \neg K_V(\{c\}, \{d\}) \wedge K_V(\{c, e\}, \{f\}) \end{aligned}$$

The language \mathcal{L}_2 allows us to connect PIL to the maybe most famous axioms about database theory and dependence logic from [5].

Lemma 3. *Armstrong's axioms are semantically valid and derivable in SPIL_1 :*

- $K_V(C, D)$ for any $D \subseteq C$ (projectivity)
- $K_V(C, D) \wedge K_V(D, E) \rightarrow K_V(C, E)$ (transitivity)
- $K_V(C, D) \wedge K_V(C, E) \rightarrow K_V(C, D \cup E)$ (additivity)

Proof. The semantic validity is easy to check, hence we focus on the derivations.

For projectivity, take any two finite sets $D \subseteq C$. If $D = C$, then we only need a derivation like the following which basically generalizes learning to finite sets.

$$\frac{\frac{\frac{[c_1]K_V(c_1)}{[c_2][c_1]K_V(c_1)} \text{ (NEC)}}{[c_1][c_2]K_V(c_1)} \text{ (COMM)}}{[c_1]([c_2]K_V(c_1) \wedge [c_2]K_V(c_2))} \text{ (DIST)}}{[c_1][c_2](K_V(c_1) \wedge K_V(c_2))} \text{ (DIST)}$$

If $D \subsetneq C$, then continue by applying NEC for all elements of $C \setminus D$ to get $K_V(C, D)$.

Transitivity follows from IR and NF as follows. For simplicity, first we only consider the case where C, D and E are singletons.

$$\frac{\frac{\frac{K_V(e) \rightarrow [c]K_V(e)}{[d](K_V(e) \rightarrow [c]K_V(e))} \text{ (NEC)}}{[d]K_V(e) \rightarrow [d][c]K_V(e)} \text{ (COMM)}}{[d]K_V(e) \rightarrow [c][d]K_V(e)} \text{ (TAUT)}}{\frac{\frac{\frac{K_V(d) \rightarrow ([d]K_V(e) \rightarrow K_V(e))}{[c](K_V(d) \rightarrow ([d]K_V(e) \rightarrow K_V(e)))} \text{ (NEC)}}{[c]K_V(d) \rightarrow [c]([d]K_V(e) \rightarrow K_V(e))} \text{ (DIST)}}{[c]K_V(d) \rightarrow ([c][d]K_V(e) \rightarrow [c]K_V(e))} \text{ (DIST)}}{[c]K_V(d) \rightarrow ([d]K_V(e) \rightarrow [c]K_V(e))} \text{ (TAUT)}$$

Now consider any three finite sets of constants $C = \{c_1, \dots, c_i\}$. Using the abbreviations from Definition 4 and the ‘‘multi’’ rules given in Lemma 1 it is easy to generalize the proof. In fact, the proof is exactly the same with capital letters.

Similarly, additivity follows immediately from multi-DIST’.

We can now use Armstrong’s axioms to prove completeness of our logic. The crucial idea is a new definition of a canonical dependency graph.

Theorem 1 (Strong Completeness). *For all sets of formulas $\Delta \subseteq \mathcal{L}_1$ and all formulas $\varphi \in \mathcal{L}_1$, if $\Delta \models \varphi$, then also $\Delta \vdash \varphi$.*

Proof. By contraposition using a canonical model. Suppose $\Delta \not\vdash \varphi$. Then $\Delta \cup \{\neg\varphi\}$ is consistent and there is a maximally consistent set $\Gamma \subseteq \mathcal{L}_1$ such that $\Gamma \supseteq \Delta \cup \{\neg\varphi\}$. We will now build a model \mathcal{M}_Γ such that for the world \mathbb{C} in that model we have $\mathcal{M}_\Gamma, \mathbb{C} \models \Gamma$ which implies $\Delta \not\vdash \varphi$.

Definition 6 (Canonical Graph and Model). Let the graph $G_\Gamma := (\mathcal{P}(\mathbb{C}), \rightarrow)$ be given by $A \rightarrow B$ iff $Kv(A, B) \in \Gamma$. By Lemma 3 this graph has properties corresponding to the Armstrong axioms: projectivity, transitivity and additivity. We call a set of variables $s \subseteq \mathbb{C}$ closed under G_Γ iff whenever $A \subseteq s$ and $A \rightarrow B$ in G_Γ , then also $B \subseteq s$. Then let the canonical model be $\mathcal{M}_\Gamma := (S, \mathcal{D}, V)$ where

$$S := \{s \subseteq \mathbb{C} \mid s \text{ is closed under } G_\Gamma\}, \mathcal{D} := \{0, 1\} \text{ and } V(s, c) = \begin{cases} 0 & \text{if } c \in s \\ 1 & \text{otherwise} \end{cases}$$

Note that our domain is just $\{0, 1\}$. This is possible because we do not have to find a model where the dependencies hold globally. Instead, $Kv(C, d)$ only says that given the C -values at the actual world, also the d values are the same at the other worlds. The dependency does not need to hold between two non-actual worlds. This distinguishes our models from relationships as discussed in [5] where no actual world or state is used, see Example 4 below.

Given the definition of a canonical model we can now show:

Lemma 4 (Truth Lemma). $\mathcal{M}_\Gamma, \mathbb{C} \models \varphi$ iff $\varphi \in \Gamma$.

Before going into the proof, let us emphasize two peculiarities of our truth lemma: First, the states in our canonical model are not maximally consistent sets of formulas but sets of constants. Second, we only claim the truth Lemma at one specific state, namely \mathbb{C} where all constants have value 0. As our language does not include nested epistemic modalities, we actually never evaluate formulas at other states of our canonical model.


Proof (Truth Lemma). Note that it suffices to show this for all φ in \mathcal{L}_2 : Given some $\varphi \in \mathcal{L}_1$, by Lemma 2 we have that $\mathcal{M}_\Gamma, \mathbb{C} \models \varphi \iff \mathcal{M}_\Gamma, \mathbb{C} \models t(\varphi)$ because the translation preserves and reflects truth. Moreover, we have $\varphi \in \Gamma \iff t(\varphi) \in \Gamma$, because $\varphi \leftrightarrow t(\varphi)$ is provable in SPILL_1 . Hence it suffices to show that $\mathcal{M}_\Gamma, \mathbb{C} \models t(\varphi)$ iff $t(\varphi) \in \Gamma$, i.e. to show the Truth Lemma for \mathcal{L}_2 . Again, negation and conjunction are standard, the crucial case are dependencies.

Suppose $Kv(C, D) \in \Gamma$. By definition $C \rightarrow D$ in G_Γ . To show $\mathcal{M}_\Gamma, \mathbb{C} \models Kv(C, D)$, take any t such that $\mathbb{C} =_C t$ in \mathcal{M}_Γ . Then by definition of V we have $C \subseteq t$. As t is closed under G_Γ , this implies $D \subseteq t$. Now by definition of V we have $\mathbb{C} =_D t$.

For the converse, suppose $Kv(C, D) \notin \Gamma$. Then by definition $C \not\rightarrow D$ in G_Γ . Now, let $t := \{c' \in \mathbb{C} \mid C \rightarrow \{c'\} \text{ in } G_\Gamma\}$. This gives us $C \subseteq t$. But we also have $D \not\subseteq t$ because otherwise additivity would imply $C \rightarrow D$ in G_Γ . Moreover, because G_Γ is transitive it is enough to “go one step” in G_Γ to get a set that is closed under G_Γ . This means that t is closed under G_Γ and therefore a state in our model, i.e. we have $t \in S$. Now by definition of V and projectivity, we have $\mathbb{C} =_C t$ but $\mathbb{C} \neq_D t$. Thus t is a witness for $\mathcal{M}_\Gamma, \mathbb{C} \not\models Kv(C, D)$.

This also finishes the completeness proof. Note that we used all three properties corresponding to the Armstrong axioms.

Example 3. To illustrate the idea of the canonical dependency graph, let us study a concrete example of what the graph and model look like. Consider the maximally consistent set $\Gamma = \{\neg K\nu(c), \neg K\nu(d), K\nu(e), K\nu(c, d), \dots\}$. The interesting part of the canonical graph G_Γ then looks as follows, where the nodes are subsets of $\{c, d, e\}$. For clarity we only draw $\rightarrow \cap \not\subseteq$, i.e. we omit edges given by inclusions. For example all nodes will also have an edge going to the \emptyset node.



To get a model out of this graph, note that there are exactly three subsets of \mathbb{C} closed under following the edges. Namely, let $S = \{s : \{e\}, t : \{d, e\}, u : \{c, d, e\}\}$ and use the binary valuation which says that a constant has value 0 iff it is an element of the state. It is then easy to check that $\mathcal{M}, u \models \Gamma$.

s	t	u
c	1	1
d	1	0
e	0	0

It is also straightforward to define an appropriate notion of bisimulation.

Definition 7. *Two pointed models $((S, \mathcal{D}, V), s)$ and $((S', \mathcal{D}', V'), s')$, are bisimilar iff (i) For all finite $C \subseteq \mathbb{C}$ and all $d \in \mathbb{C}$: If there is a $t \in S$ such that $s =_C t$ and $s \neq_d t$, then there is a $t' \in S'$ such that $s' =_C t'$ and $s' \neq_d t'$; and (ii) Vice versa.*

Note that we do not need the bisimulation to also link non-actual worlds. This is because all formulas are evaluated at the same world. In fact it would be too strong for the following characterization.

Theorem 2. *Two pointed models satisfy the same formulas iff they are bisimilar.*

Proof. By Lemma 2 we only have to consider formulas of \mathcal{L}_2 . Moreover, it suffices to consider formulas $K\nu(C, d)$ with a singleton in the second set because $K\nu(C, D)$ is equivalent to $\bigwedge_{d \in D} K\nu(C, d)$. Then it is straightforward to show that if \mathcal{M}, s and \mathcal{M}', s' are bisimilar then $\mathcal{M}, s \models \neg K\nu(C, d) \iff \mathcal{M}', s' \models \neg K\nu(C, d)$ by definition of our bisimulation. The other way around is also obvious since the two conditions for bisimulation are based on the semantics of $\neg K\nu(C, d)$.

Note that a bisimulation characterization for a language without the dynamic operator can be obtained by restricting Definition 7 to $C = \emptyset$. We leave it as an exercise for the reader to use this and Theorem 2 to show that $[c]$ is not reducible, which distinguishes it from the public announcement $[\varphi]$ in PAL.

Example 4 (Pointed Models Make a Difference). It seems that the following theorem of our logic does not translate to Armstrong’s system from [5].

$$[c](Kv(d) \vee Kv(e)) \leftrightarrow ([c]Kv(d) \vee [c]Kv(e))$$

First, to see that this is provable, note that it follows from determinacy and seriality. Second, it is valid because we consider pointed models which convey more information than a simple list of possible values. Consider the following table which represents 4 possible worlds.

c	d	e
1	1	3
1	1	2
2	2	1
2	3	1

Here we would say that “After learning c we know d or we know e ”, i.e. the antecedent of above formula holds. However, the consequent only holds if we evaluate formulas while pointing at a specific world/row: It is globally true that given c we will learn d or that given c we will learn e . But none of the two disjuncts holds globally which would be needed for a dependency in Armstrong’s sense. Note that this is more a matter of expressiveness than of logical strength. In Armstrong’s system there is just no way to express $[c](Kv(d) \vee Kv(e))$.

4 Multi-Agent PIL

We now generalize the Public Inspection Logic to multiple agents. In the language we use Kv_i to say that agent i knows the value of c and in the models an accessibility relation for each agent is added to describe their knowledge. To obtain a complete proof system we can leave most axioms as above but have to restrict the irrelevance axiom. Again the completeness +proof uses a canonical model construction and a truth lemma for a +restricted but equally expressive syntax. The only change is that we now define a dependency graph for each agent in order to define accessibility relations instead of restricted sets of worlds.

Definition 8 (Multi-Agent PIL). We fix a non-empty set of agents I . The language \mathcal{L}_1^I of multi-agent Public Inspection Logic is given by

$$\varphi ::= \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid Kv_i c \mid [c]\varphi$$

where $i \in I$. We interpret it on models $\langle S, \mathcal{D}, V, R \rangle$ where S , \mathcal{D} and V are as before and R assigns to each agent i an equivalence relation \sim_i over S . The semantics are standard for the booleans and as follows:

$$\begin{array}{l} \overline{\mathcal{M}, s \models Kv_i c \iff \forall t \in S : s \sim_i t \Rightarrow s =_c t} \\ \overline{\mathcal{M}, s \models [c]\varphi \iff \mathcal{M}|_c^s, s \models \varphi} \end{array}$$

where $\mathcal{M}|_c^s$ is $\langle S', \mathcal{D}, V|_{S' \times \mathbb{C}}, R|_{S' \times S'} \rangle$ with $S' = \{t \in S \mid s =_c t\}$.

Analogous to Definition 4 we define the following abbreviation to express dependencies known by agent i and note its semantics:

$$Kv_i(C, D) := [c_1] \dots [c_n](Kv_i(d_1) \wedge \dots \wedge Kv_i(d_m))$$

$$\overline{\mathcal{M}, s \models Kv_i(C, D) \Leftrightarrow \text{for all } t \in S : \text{if } s \sim_i t \text{ and } s =_C t \text{ then } s =_D t}$$

The proof system SPiLL for PiL in the language \mathcal{L}_1^I is obtained by replacing each Kv in the axioms of SPiLL_1 by Kv_i , and replacing IR by the following restricted version:

$$\text{RIR} \quad Kv_i c \rightarrow ([c]\varphi \rightarrow \varphi) \text{ where } \varphi \text{ does not mention any agent besides } i$$

Before summarizing the completeness proof for the multi-agent setting, let us highlight some details of this definition.

As before the actual state s plays an important role in the semantics of $[c]$. However, we could also use an alternative but equivalent definition: Instead of deleting states, only delete the \sim_i links between states that disagree on the value of c . Then the update no longer depends on the actual state.

For traditional reasons we define \sim_i to be an equivalence relation. This is not strictly necessary, because our language can not tell whether the relation is reflexive, transitive or symmetric. Removing this constraint and extending the class of models would thus not make any difference in terms of validities.

For the proof system, note that the original irrelevance axiom IR is *not* valid in the multi-agent setting because φ might talk about other agents for which the inspection of c does matter.

Theorem 3 (Strong Completeness for SPiLL). *For all sets of formulas $\Delta \subseteq \mathcal{L}_1^I$ and all formulas $\varphi \in \mathcal{L}_1^I$, if $\Delta \models \varphi$, then also $\Delta \vdash \varphi$.*

Proof. By the same methods as for Theorem 1. Given a maximally consistent set $\Gamma \subseteq \mathcal{L}_1^I$ we want to build a model \mathcal{M}_Γ such that for the world \mathbb{C} in that model we have $\mathcal{M}_\Gamma, \mathbb{C} \models \Gamma$.

First, for each agent $i \in I$, let G_Γ^i be the graph given by $A \rightarrow_i B \quad : \Leftrightarrow \Gamma \vdash Kv_i(A, B)$. Given that the proof system SPiLL was obtained by indexing the axioms of SPiLL_1 , it is easy to check that indexed versions of the Armstrong axioms are provable and therefore all the graphs G_Γ^i for $i \in I$ will have the corresponding properties. In particular RIR suffices for this.

Second, define the canonical model $\mathcal{M}_\Gamma := (S, \mathcal{D}, V, R)$ where $S := \mathcal{P}(\mathbb{C})$, $\mathcal{D} := \{0, 1\}$, $V(s, c) := 0$ if $c \in s$ and $V(s, c) := 1$ otherwise, and $s \sim_i t$ iff s and t are both closed or both not closed under G_Γ^i .

Lemma 5 (Multi-Agent Truth Lemma). $\mathcal{M}_\Gamma, \mathbb{C} \models \varphi$ iff $\varphi \in \Gamma$.

Proof. Again it suffices to show the Truth Lemma for a restricted language and we only consider the state \mathbb{C} . We proceed by induction on φ . The crucial case is when φ is of form $Kv_i(C, D)$.

Suppose $Kv_i(C, D) \in \Gamma$. Then by definition $C \rightarrow D$ in G_Γ^i . To show $\mathcal{M}_\Gamma, \mathbb{C} \models Kv_i(C, D)$, take any t such that $\mathbb{C} \sim_i t$ and $\mathbb{C} =_C t$ in \mathcal{M}_Γ . Then by definition

of V we have $C \subseteq t$. Moreover, \mathbb{C} is closed under G_T^i . Hence by definition of \sim_i also t must be closed under G_T^i , which implies $D \subseteq t$. Now by definition of V we have $\mathbb{C} =_D t$.

For the converse, suppose $Kv_i(C, D) \notin \Gamma$. Then by definition $C \not\rightarrow D$ in G_T^i . Now, let $t := \{c' \in \mathbb{C} \mid C \rightarrow \{c'\} \text{ in } G_T^i\}$. This gives us $C \subseteq t$. But we also have $D \not\subseteq t$ because otherwise additivity would imply $C \rightarrow D$ in G_T^i . Moreover, because G_T^i is transitive it is enough to “go one step” in G_T^i to get a set that is closed under G_T^i . This means that t is closed under G_T^i and therefore by definition of \sim_i we have $\mathbb{C} \sim_i t$. Now by definition of V and projectivity, we have $\mathbb{C} =_C t$ but $\mathbb{C} \neq_D t$. Thus t is a witness for $\mathcal{M}_\Gamma, \mathbb{C} \not\models Kv_i(C, D)$.

Again the Truth Lemma also finishes the completeness proof.




Fig. 1. Two canonical dependency graphs and the resulting canonical model.

Example 5. Analogous to Example 3, the following illustrates the multi-agent version of our canonical construction. Consider the maximally consistent set $\Gamma = \{\neg Kv_1(d), Kv_1(c, d), \neg Kv_1(d, c), \neg Kv_2(c), \neg Kv_2(c, d), Kv_1(d, c), \dots\}$. Note that agents 1 and 2 do not differ in which values they know right now but there is a difference in what they will learn from inspections of c and d . The two canonical dependency graphs generated from Γ are shown in Figure 1. Again for clarity we only draw the non-inclusion arrows. The subsets of $\mathbb{C} = \{c, d\}$ closed under the graphs are thus $\{\{c, d\}, \{d\}, \emptyset\}$ and $\{\{c, d\}, \{c\}, \emptyset\}$ for agent 1 and 2 respectively, inducing the equivalence relations as shown in Figure 1.

It is also not hard to find the right notion of bisimulation for $\mathbb{S}PIL$.

Definition 9. Given two models (S, \mathcal{D}, V, R) and $(S', \mathcal{D}', V', R')$, a relation $Z \subseteq S \times S'$ is a multi-agent bisimulation iff for all sZs' we have (i) For all finite $C \subseteq \mathbb{C}$, all $d \in \mathbb{C}$ and all agents i : If there is a $t \in S$ such that $s \sim_i t$ and $s =_C t$ and $s \neq_d t$, then there is a $t' \in S'$ such that tZt' and $s' \sim_i t'$ and $s =_C t$ and $s' \neq_d t'$; and (ii) Vice versa.

Theorem 4. Two pointed models satisfy the same formulas of the multi-agent language \mathcal{L}_1^I iff there is a multi-agent bisimulation linking them.

As it is very similar to the one of Theorem 2, we omit the proof here.

5 Future Work

Between our specific approach and the general language of [10], a lot can still be explored. An advantage of having a weaker language with explicit operators, instead of encoding them in a more general language, is that we can clearly see the properties of those operators showing up as intuitive axioms.

The framework can be extended in different directions. We could for example add equalities $c = d$ to the language, together with knowledge $K(c = d)$ and announcement $[c = d]$. No changes to the models are needed, but axiomatizing these operators seems not straightforward. Alternatively, just like Plaza added Kv to PAL, we can also add K to PIL. Another next language to be studied is thus $PIL + K$ from Table 2 above and given by

$$\varphi ::= \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid K v_i c \mid K_i \varphi \mid [c]\varphi.$$

Note that in this language, we can also express *knowledge of dependency* in contrast to *de facto* dependency. For example, $K_i[c]K v_i d$ expresses that agent i knows that d functionally depends on c , while $[c]K v_i d$ express that the value of d (given the information state of i) is determined by the *actual value* of c *de facto*. In particular the latter does not imply that i knows this. The agent can still consider other values of c possible that would not determine the value of d . To see the difference technically, we can spell out the truth condition for $K_i[c]K v_i(d)$ under standard Kripke semantics for K_i on S5 models:

$$\mathcal{M}, s \models K_i[c]K v_i(d) \Leftrightarrow \text{for all } t_1 \sim_i s, t_2 \sim_i s : t_1 =_c t_2 \implies t_1 =_d t_2$$

Now consider Example 4: $[c]K v(d)$ holds in the first row, but $K[c]K v(d)$ does not hold since the semantics of K require $[c]K v(d)$ to hold at *all* worlds considered possible by the agent. This also shows that $[c]K v(d)$ is not positively introspective (i.e. the formula $[c]K v(d) \rightarrow K_i[c]K v(d)$ is not valid), and it is essentially not a subjective epistemic formula.

In this way, $K[c]K v(d)$ can also be viewed as the atomic formula $=(c, d)$ in *dependence logic* (DL) from [11]. A *team model* of DL can be viewed as the set of epistemically accessible worlds, i.e., a single-agent model in our case. The connection with dependence logic also brings PIL closer to the first-order variant of *epistemic inquisitive logic* by [12], where knowledge of entailment of interrogatives can also be viewed as the knowledge of dependency. For a detailed comparison with our approach, see [13, Sec. 6.7.4].

Another approach is to make the dependency more explicit and include functions in the syntax. In [14] a functional dependency operator $\mathcal{K}f_i$ is added to the epistemic language with $K v_i$ operators: $\mathcal{K}f_i(c, d) := \exists f K_i(d = f(c))$ where f ranges over a pool of functions.

Finally, there is an independent but related line of work on (in)dependency of variables using predicates, see for example [15,16,17,18]. In particular, [17] also uses a notion of dependency as an epistemic implication “Knowing c implies knowing d .”, similar to our formula $K v(c, d)$. In [18] also a “dependency graph” is used to describe how different variables, in this case payoff functions in strategic

games, may depend on each other. Note however, that these graphs are not the same as our canonical dependency graphs from Definition 6. Our graphs are directed and describe determination between sets of variables. In contrast, the graphs in [18] are undirected and consist of singleton nodes for each player in a game. We leave a more detailed comparison for a future occasion.

Acknowledgements. We thank the following people for useful comments on this work: Alexandru Baltag, Peter van Emde Boas, Hans van Ditmarsch, Jie Fan, Kai Li and our anonymous reviewers.

References

1. Wang, Y.: Beyond knowing that: a new generation of epistemic logics. In van Ditmarsch, H., Sandu, G., eds.: Jaakko Hintikka on knowledge and game theoretical semantics. Springer (2016) forthcoming.
2. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic epistemic logic. Volume 1. Springer Heidelberg (2007)
3. Plaza, J.: Logics of public communications. *Synthese* **158**(2) (2007) 165–179
4. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In Bilboa, I., ed.: TARK'98. (1998) 43–56
5. Armstrong, W.W.: Dependency structures of data base relationships. In: IFIP congress. Volume 74., Geneva, Switzerland (1974) 580–583
6. Sweeney, L.: Only you, your doctor, and many others may know. *Technology Science* (2015) published online: <http://techscience.org/a/2015092903/>.
7. Wang, Y., Fan, J.: Knowing that, knowing what, and public communication: Public announcement logic with kv operators. In: IJCAI '13. (2013) 1147–1154
8. Wang, Y., Fan, J.: Conditionally knowing what. *Advances in Modal Logic* **10** (2014) 569–587
9. Gu, T., Wang, Y.: "knowing value" logic as a normal modal logic. *Advances in Modal Logic* **11** (2016) 362–381
10. Baltag, A.: To know is to know the value of a variable. *Advances in Modal Logic* **11** (2016) 135–155
11. Väänänen, J.: *Dependence Logic: A New Approach to Independence Friendly Logic*. Cambridge University Press (2007)
12. Ciardelli, I., Roelofsen, F.: Inquisitive dynamic epistemic logic. *Synthese* **192**(6) (2015) 1643–1687
13. Ciardelli, I.: *Questions in Logic*. PhD thesis, University of Amsterdam (2016)
14. Ding, Y.: *Epistemic logic with functional dependency operator*. Bachelor's thesis (in Chinese), Peking University (2015)
15. More, S.M., Naumov, P.: An independence relation for sets of secrets. *Studia Logica* **94**(1) (2010) 73–85
16. Naumov, P.: Independence in information spaces. *Studia Logica* **100**(5) (2012) 953–973
17. Naumov, P., Nicholls, B.: Rationally functional dependence. *J. Philosophical Logic* **43**(2-3) (2014) 603–616
18. Harjes, K., Naumov, P.: Functional dependence in strategic games. *Notre Dame Journal Formal Logic* **57**(3) (2016) 341–353