An Epistemic Logic with Hypotheses

Levan Uridia[1] and Dirk Walther[2]

[1] Rey Juan Carlos University (URJC), Madrid, Spain
uridia@ia.urjc.es
[2] Technical University of Madrid (UPM), Madrid, Spain
dirk.walther@upm.es

Abstract. We introduce a variant of the standard epistemic logic S5 for reasoning about knowledge under hypotheses or background assumptions. The modal operator of necessity expressing what is known is parameterised with a hypothesis. The operator can be described as relative necessity, a notion already used by Chellas to describe conditionality. In fact, the parameterised box resembles a conditional operator and it turns out that our logic is a variant of Chellas' Conditional Logic. We present an axiomatisation of the logic and show that it bears the same expressivity and computational complexity as S5. Then we consider the extension of our logic with operators for distributed knowledge and show how it can be used to represent knowledge of agents whose epistemic capacity corresponds to any system between S4 and S5.

Keywords: epistemic logic, conditional logic, knowledge, distributed knowledge, hypotheses.

1 Introduction

We introduce an epistemic logic for reasoning about knowledge under hypotheses. The resulting logic $S5^r$ is an extension of S5 with a modal operator '$[\cdot]$' that can be parameterised with a hypothesis. The modality $[\varphi]$ represents the knowledge state under the hypothesis φ. The formula $[\varphi]\psi$ states that 'under the hypothesis φ, the agent knows ψ'. If φ happens to be true at the current world and the agent knows that φ implies ψ, then the agent knows ψ; otherwise, i.e., if φ is false, the agent knows only what it would know anyway, i.e. without any assumptions. For instance, suppose an agent is interested in whether the street is dry or wet. The agent knows that rain makes the street wet, but it does not know whether or not it is raining outside. The formula

$$[\,\textsf{it-is-raining}\,]\ \textsf{street-is-wet}$$

states that the agent knows that the street is wet when adopting the hypothesis that it is raining. We consider two situations: one, where the hypothesis is correct, i.e., it is indeed raining; and another one, where it is false, i.e., it is not raining. Clearly, in the former situation, the street is wet due to the rain and we have that the formula holds true. In case the hypothesis is in fact wrong, the formula

H. van Ditmarsch, J. Lang, and S. Ju (Eds.): LORI 2011, LNAI 6953, pp. 286–299, 2011.

is *not necessarily* true. Notice here the difference to an ordinary implication, which is true whenever the premise is false. Generally, the agent may consider a dry street possible despite adopting the hypothesis that it is raining. But this means that the agent does not know that the street is wet and, thus, the formula is false. The formula is true, however, if the agent does not consider a dry street possible. In this case, the agent already knew that the street is wet before, i.e. without assuming that it is raining. The latter is expressed by the formula:

$$[\top]\ \textsf{street-is-wet}.$$

The parameterised modal operator can be described as relative necessity, a notion already used by Chellas to describe conditionality [5]. In fact, there is a strong relation with Chellas' Condition Logic as $S5^r$ turns out to be a special case of it. We present an axiomatisation of $S5^r$ and we show it is as expressive and complex as S5.

In the second part of the paper, we extend $S5^r$ with operators for distributed knowledge but use them for combining hypotheses. Distributed knowledge is a standard notion in epistemic logic [9,16]. Generally, distributed knowledge of a group of agents equals what a single agent knows who knows everything what each member of the group knows. Suppose agent a knows p and agent b knows $p \to q$. Then the distributed knowledge between a and b is q, even though neither of them might know q individually. The notion of distributed knowledge is relevant for describing and reasoning about the combined knowledge of agents in a distributed system; see, e.g., [10]. There agents communicate with each other to combine their knowledge. Thus the notion of distributed knowledge is also central to communication protocols and relevant to reasoning about speech acts [11,20]. We may think of distributed hypotheses as the result of combining incoming information from several sources. However, the truthfulness of the incoming information is not assumed. We demonstrate another way to think about distributed hypotheses by using our logic to represent the knowledge of an agent whose knowledge capacity can be characterised by any modal system between S4 and S5.

The relative necessity operator of $S5^r$ shows up in several places, e.g., in Conditional Logic, Public Announcement Logic and Provability Logic. We now discuss the relationships to these logics in more detail and explain how far the relative necessity is paraconsistent.

Sentences in English of the form "If A, then B." are called conditional sentences. Here, A is called the *antecedent* (or *condition*) and B the *consequent.* Conditional sentences are traditionally put into different categories (according to mood or tense) such as indicative/subjunctive or factual/counterfactual. However, there is much disagreement on the logical theory of conditional sentences (in particular that of defeasible conditionals). One logical formalisation is Conditional Logic, which essentially is Propositional Logic extended with a binary operator '$\Rightarrow$' standing for conditionality. Several readings of '$\Rightarrow$' were proposed, among them counterfactual conditional, non-monotonic consequence relation, normality and belief revision. Historically several logical accounts of conditionals have been suggested, among them Stalnaker [22], Lewis [14] and Chellas [5]; for an overview on Conditional Logic, see, e.g., [17].

Our logic $S5^r$ rejects the common assumption that logics allow one to conclude anything from false premises. To be more precise we borrow the term 'explosive' from Paraconsistent Logic, but we refer to conditional operators instead of the logical consequence relation. We say that a conditional operator 'X' is *explosive* if the conditional $\varphi X \psi$ holds for all conclusions ψ whenever the antecedent φ is false. In this sense, implication of Classical Logic and even of Intuitionistic Logic is explosive, so is the conditional operator '$\Rightarrow$' of Conditional Logic [14,22]. On the other hand, the relativised necessity of our logic, which is a special case of Chellas' conditional operator [5,6], is not explosive. We have that $[\bot]\psi$ is true if, and only if, ψ is universally true. Notice that this does not mean that the consequence relation of our logic paraconsistent; it is not.

Epistemic logic traditionally describes the knowledge state of agents at a point in time. To be able to describe the evolution of knowledge over time, we can either combine epistemic logic with a temporal logic, or add dynamic operators for knowledge-changing actions such as communication. The latter approach is followed in the family of Dynamic Epistemic Logics (DELs) [7]. A basic DEL is Public Announcement Logic (PAL) [18] which is the extension of the basic epistemic logic S5 with an operator '$[\cdot]$' parameterised with a formula expressing the announcement. A formula of the form $[\varphi]\psi$ states that ψ holds after φ has been truthfully announced (by someone) to every agent in the system simultaneously. After the announcement, φ is incorporated in the knowledge state of every agent, i.e. φ becomes common knowledge. This is achieved by employing an update semantics, which cuts off the model all worlds in which the announcement does not hold. Here we have an important difference to our logic $S5^r$ whose semantics does not change models during the evaluation of a formula. It is well-known that PAL is as expressive as S5 [18,7], but the translation of a PAL-formula to an equivalent formula of S5 causes an exponential blow-up in size. Despite the succinctness of PAL, however, the complexity of the satisfiability checking problem coincides with that for S5, i.e. NP-complete for the single-agent PAL and PSpace-complete in the presence of multiple-agents [15].

Our relative necessity operator '$[\cdot]$' bears an interesting relationship to Provability Logic [3]. This is a modal logic, where the modality is considered to capture the metamathematical concept of 'a proposition being provable in some arithmetical theory'. An important logic in this context is the Gödel–Löb system GL as it characterises provability in Peano Arithmetic. A recent line of research is to modify GL and study the effects of these modifications wrt. provability [8]. For instance, [8] introduces three variants of the modality in GL and studies their algebraic semantics. Here is where the connection to this paper turns up as one of the modified modalities corresponds syntactically to the relative necessity operator '$[\cdot]$' considered in this paper. To be precise, the definition of the modal operator called 'Modest Enrichment (Type B)' in [8] equals Axiom (R) $[\varphi]\psi \leftrightarrow \Box\psi \vee \varphi \wedge \Box(\varphi \to \psi)$, which we introduce below. In this paper, however, we do not investigate further the relationship to Provability Logic.

The paper is organised as follows. In the next section, the logic $S5^r$ is formally introduced as extension of S5. In sections 3 and 4, we show that $S5^r$ is a special case of Chellas' Conditional Logic and we give a sound and complete axiomatisation of $S5^r$, respectively. Then, in Section 5, a polynomial reduction of $S5^r$ to S5 is presented together with a discussion on the relation between $S5^r$-modalities and the difference operator. Section 6 is used to extend $S5^r$ with operators for distributed hypotheses that are analog to distributed knowledge. We show how to knowledge of an agent can be represented as distributed hypotheses, where the agents' knowledge corresponds to any system between S4 and S5. The paper closes with conclusions in Section 7.

2 The Modal Logic

In this section, we introduce the multi-modal logic $S5^r$. Essentially, the language of $S5^r$ is the language of Propositional Logic extended with modal operators parameterised with $S5^r$-formulas. Formally, this is done as follows.

Definition 1 (Syntax of $S5^r$). *Let Π be a countably infinite set of atomic propositions. Formulas φ of the language $\mathcal{L}$ are defined inductively over Π by the following grammar:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid [\varphi]\varphi,$$

where p ranges over atomic propositions in Π. ⊣

The logical symbols '$\top$' and '$\bot$', and additional operators such as '$\wedge$', '$\rightarrow$', '$\leftrightarrow$' and the dual modality '$\langle\cdot\rangle$' are defined as usual, i.e.: $\top := p \vee \neg p$ for some atomic proposition p; $\bot := \neg\top$; $\varphi \wedge \psi := \neg(\neg\varphi \vee \neg\psi)$; $\varphi \rightarrow \psi := \neg\varphi \vee \psi$; $\varphi \leftrightarrow \psi := (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$; and $\langle\varphi\rangle\psi := \neg[\varphi]\neg\psi$.

Modal formulas are commonly evaluated in Kripke structures containing a binary relation over the domain, one for each modality in the language. In this case, however, every relation is determined by the valuation of the atomic propositions in the domain. Therefore, it is sufficient to consider Kripke structures without relations, which we call basic structures. Formally, a *basic structure* $\mathfrak{M}$ is a tuple $\mathfrak{M} = (W, V)$, where W is a non-empty set of worlds and $V : \Pi \rightarrow 2^W$ a valuation function mapping every atomic proposition p to a set of worlds $V(p)$ at which it is true. We will also refer to a basic structure simply as a *model*. The relations that are required to evaluate the modalities are defined alongside the logical consequence relation. But first we introduce an auxiliary notion, a binary operation '$\otimes$' on sets yielding a binary relation. Let X and Y be two sets. Let $X \otimes Y$ be a binary relation over $X \cup Y$ such that

$$X \otimes Y = X^2 \cup (X \times Y) \cup Y^2. \qquad (1)$$

We illustrate this notion with an example.

Example 1. Let $X = \{x_1, x_2\}$ and $Y = \{y_1, y_2, y_3\}$ be two sets. Then, according to (1), $X \otimes Y$ is a binary relation over $X \cup Y$ that is composed of the relations X^2,

$X \times Y$ and Y^2 by taking their union. We have that $X^2 = \{(x_1, x_2), (x_2, x_1)\} \cup \mathsf{id}(X)$, $X \times Y = \{(x_1, y_1), (x_1, y_2), (x_1, y_3), (x_2, y_1), (x_2, y_2), (x_2, y_3)\}$ and $Y^2 = \{(y_1, y_2), (y_2, y_1), (y_1, y_3), (y_3, y_1), (y_2, y_3), (y_3, y_2)\} \cup \mathsf{id}(Y)$. Then the relation $X \otimes Y = X^2 \cup (X \times Y) \cup Y^2$ contains two fully connected clusters X^2 and Y^2, and directed edges between every point in X to every point in Y. Figure 1 below gives a graphical representation of $X \otimes Y$ (leaving out the reflexive and symmetric edges). $\lhd$

We are now ready to introduce the semantics of $\mathcal{L}$. It differs from the semantics of Public Announcement Logic [18,7] in that the model does not change during the evaluation of formulas.

Definition 2 (Semantics of S5^r). *Let $\mathfrak{M} = (W, V)$ be a basic structure. The logical satisfaction relation '$\models$' is defined by induction on the structure of $\mathcal{L}$-formulas as follows: For all $p \in \Pi$ and all $\varphi, \psi \in \mathcal{L}$,*

- $\mathfrak{M}, w \models p$ *iff* $w \in V(p)$;
- $\mathfrak{M}, w \models \varphi \vee \psi$ *iff* $\mathfrak{M}, w \models \varphi$ *or* $\mathfrak{M}, w \models \psi$;
- $\mathfrak{M}, w \models [\varphi]\psi$ *iff for all* $v \in W$ *with* $(w, v) \in R_\varphi$, $\mathfrak{M}, v \models \psi$,

where $R_\varphi = (W \setminus [\![\varphi]\!]_{\mathfrak{M}}) \otimes [\![\varphi]\!]_{\mathfrak{M}}$ *as defined in (1) and* $[\![\varphi]\!]_{\mathfrak{M}} = \{x \in W \mid \mathfrak{M}, w \models \varphi\}$ *is the* extension *of* φ *in* $\mathfrak{M}$. $\dashv$

We say that a S5^r-formula φ is *satisfiable* if there is a model $\mathfrak{M}$ and a world w in $\mathfrak{M}$ such that $\mathfrak{M}, w \models \varphi$; φ is *valid in* $\mathfrak{M}$ if $\mathfrak{M}, w \models \varphi$ for all w in $\mathfrak{M}$; and φ is *valid* if φ is valid in all models. We will refer to the relation R_φ as being *determined* by φ and a model.

According to the semantics, a formula determines a binary relation in a model. The following proposition states the properties of such relations.

Proposition 1. *Let φ be an S5^r-formula and let $\mathfrak{M} = (W, V)$ be a basic structure. Then, the relation R_φ determined by φ and $\mathfrak{M}$ (cf. Definition 2) is a* one-step *total preorder, i.e. R_φ satisfies the following conditions:*

- R_φ *is transitive:* $\forall x, y, z.(xR_\varphi y) \wedge (yR_\varphi z) \to (xR_\varphi z)$;
- R_φ *is total:* $\forall x, y.(xR_\varphi y) \vee (yR_\varphi x)$;
- R_φ *is* one-step: $\forall x, y, z.(xR_\varphi y) \wedge \neg(yR_\varphi x) \wedge (xR_\varphi z) \to (zR_\varphi y)$.

Instead of 'preorder' also the term 'quasiorder' is often used in the literature. Note that totality implies reflexivity and that a symmetric total preorder is an equivalence relation. The proposition is readily checked as any relation R_φ in a model determined by φ is defined using the operation '$\otimes$', which always yields a so-called 'one-step total preorder'. As the domain of a model is non-empty, it contains at least one point and, thus, the smallest relation R_φ is the edge of a single reflexive point.

Figure 1 illustrates the relation R_φ in a model $\mathfrak{M}$. The domain of $\mathfrak{M}$ is partitioned into two clusters, the worlds in each of which are fully connected (reflexive and symmetric edges within the clusters are not shown). Between the clusters there are outgoing directed edges from worlds in the cluster on the left to worlds

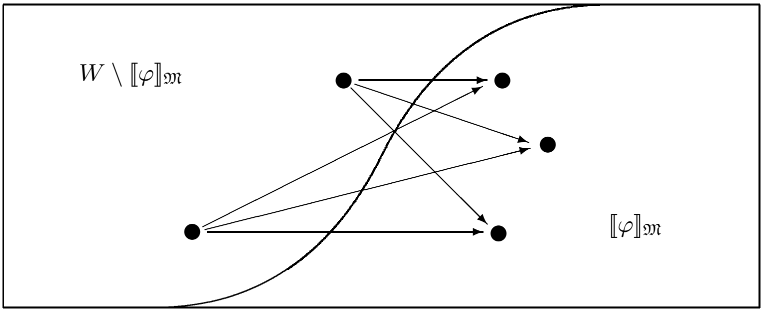

Fig. 1. Model $\mathfrak{M}$ with relation R_φ

in the cluster on the right-hand side, but not *vice versa*. Revisit Example 1 to see in detail how R_φ is computed (where $X = W \setminus [\![\varphi]\!]_{\mathfrak{M}}$ and $Y = [\![\varphi]\!]_{\mathfrak{M}}$).

Consider the following example, which demonstrates what effect hypotheses can have on an agent's knowledge.

Example 2. Let $\mathfrak{M} = (W, V)$ be a basic structure with $W = \{x, y\}$, $V(p_h) = V(p_c) = \{x\}$ and $V(p_u) = \{x, y\}$. Intuitively, the three atomic propositions p_h, p_c and p_u stand for hypothesis, conclusion and universal or already established knowledge. Then, $[p_h]q_u$ is true at x and y in $\mathfrak{M}$. In fact, we have that $\mathfrak{M}, x \models [\varphi]q_u$ for every S5^r-formula φ, because q_u holds everywhere in $\mathfrak{M}$. But $[p_h]q_c$ holds only at x and not at y, because $\mathfrak{M}, x \models p_h$ and p_h implies q_c everywhere in $\mathfrak{M}$. ◁

We conclude this section with a discussion on how S5^r could possibly be used to reason about the knowledge of multiple agents; see, e.g., [9,16] for standard references. Syntactically, S5^r is a single-agent logic. That is, it does not provide us with syntactic markers to distinguish agents such as a different modality for each agent as in S5$_n$. Consequently, there is no way to distinguish different agents other than by what they know. In S5^r we can represent the individuality of agents in the hypothesis itself. For instance, in order to represent what the agents a and b know, we can use different hypotheses p_a and p_b, which are atomic propositions labelling the states which the agents a and b, respectively, consider possible. Thus $[p_a]\varphi$ states 'a knows φ' and $[p_b]\psi$ states that 'b knows ψ'. However, this approach appears to be limited. It is unlikely to be able to encode S5$_n$ in this way as this simple complexity-theoretic argument will make clear (i.e. unlikely to the degree the following complexity classes are distinct): The satisfiability problem of S5$_n$ for $n > 1$ is PSpace-complete [12], whereas it is NP-complete for S5^r as it is shown in Section 5 below.

3 Chellas' Conditional Logic

The language $\mathcal{L}$ of S5^r coincides with the language of Chellas' Conditional Logic [5,6,21]. In this section, we make precise the relationship between the two logics. It turns out that $S5^r$ is a special case of the Conditional Logic considered by Chellas.

Chella's semantics of the Conditional Logic is given in terms of conditional structures. A *conditional structure* $\mathfrak{M}^c$ is a tuple $\mathfrak{M}^c = (W, f, V)$ which extends basic structures with a *condition function* $f : W \times 2^W \to 2^W$ that assigns worlds w and sets of worlds X to sets $f(w, X)$ of worlds. We also refer to them as *conditional models*. The set X is also understood as the extension of a proposition.

The conditional semantics is defined as follows.

Definition 3 (Conditional Semantics). *Let $\mathfrak{M}^c = (W, f, V)$ be a conditional structure. The logical consequence relation '$\models_c$' is defined as (we omit the Boolean cases): For all $\varphi, \psi \in \mathcal{L}$,*

– $\mathfrak{M}^c, w \models_c [\varphi]\psi$ *iff* $f(w, [\![\varphi]\!]^c_{\mathfrak{M}^c}) \subseteq [\![\psi]\!]^c_{\mathfrak{M}^c}$,

where $[\![\chi]\!]^c_{\mathfrak{M}^c}$ is the extension of $\mathcal{L}$-formula χ in $\mathfrak{M}^c$. ⊣

Now we can make precise the relationship between Chellas' Conditional Logic and $S5^r$. Let φ be a formula of $\mathcal{L}$ and $\mathfrak{M} = (W, V)$ a model. R_φ is the relation determined by φ and $\mathfrak{M}$. Let $\mathfrak{M}^c = (W, f, V)$ be a conditional model. We obtain

$$(w, v) \in R_\varphi \quad \text{iff} \quad v \in f(w, [\![\varphi]\!]^c_{\mathfrak{M}^c}),$$

whenever the condition function f satisfies

$$f(w, X) = \begin{cases} X & \text{if } w \in X, \\ W & \text{otherwise.} \end{cases} \tag{2}$$

That is, $S5^r$ is Chellas' Conditional Logic over the class of conditional models whose condition function f satisfies (2).

Chellas studied an axiomatisation of his logic and he showed that it is complete [5]. It is readily checked that Chellas axioms are sound for $S5^r$ as well, whereas completeness remains to be studied.

4 Axiomatisation

In this section, we present a sound and complete axiomatisation of $S5^r$. The axiom system consists of all propositional tautologies and the following axioms:

(K) $[\varphi](p \to q) \to ([\varphi]p \to [\varphi]q)$
(T) $[\top]p \to p$
(4) $[\top]p \to [\top][\top]p$
(B) $p \to [\top]\neg[\top]p$
(R) $[\varphi]\psi \leftrightarrow [\top]\psi \vee (\varphi \wedge [\top](\varphi \to \psi))$

The first four axioms are similar to the axioms known from the modal system S5 characterising any modality $[\varphi]$ in our logic $S5^r$ as epistemic operator that can be used to represent what is known under the hypothesis φ. Notice that, while

Axiom (T) establishes that what is known must in fact be true, we do not have such an axiom for hypotheses. That is, a hypothesis φ can be false and, thus, $[\varphi]\psi \to \varphi$ is not a theorem in our logic.

The axioms (T), (4), and (B) are for the modality $[\top]$ only, whereas we need additional instances of the axioms (K) and (R), namely the ones for each modal parameter φ (cf. Definition 1). The reduction axiom (R) states that every modality $[\varphi]$ is definable in terms of the basic modal operator $[\top]$, which corresponds to the S5-box or the universal modality. As it was already mentioned in the introduction, Axiom (R) corresponds to the definition of the modal operator 'Modest Enrichment (Type B)' in [8].

Theorem 1. *The system S5^r is sound and complete wrt. basic structures.* $\dashv$

We skip the proof here for reasons of space. The proof follows mainly the standard canonical model construction that can be found in textbooks on modal logic; see, e.g., [2].

5 Reduction

In what follows, we show a polynomial reduction of S5^r-formulas into S5.

We denote by $S5^r_0$ the fragment of S5^r containing formulas in which $[\top]$ is the only modality allowed to occur. It is readily checked that $S5^r_0$ is a notational variant of S5. By means of Axiom (R), we obtain a method of rewriting S5^r-formulas into formulas of $S5^r_0$. That is, we already know that S5^r is no more expressive than S5. At first sight this axiom might lead to the impression that S5^r may be exponentially more succinct[1] than S5 as happens with Public Announcement Logic [18,7]. The following proposition states this is not the case. In fact, we can translate every S5^r-formula into a formula of its fragment $S5^r_0$ that is equivalent (modulo new symbols) and without exponential blow-up in formula size. The proof Proposition 2 uses a reduction technique similar to the one used in [1], where it is shown that any hybrid language with the difference operator can polynomially be reduced to its fragment without this operator while preserving satisfiability.

Proposition 2. *Let φ be an S5^r-formula of the form $[\psi]\theta$. Then, there is an $S5^r_0$-formula φ' of length polynomial in the length of φ such that $\varphi \equiv_\Sigma \varphi'$, where $\Sigma = \mathsf{sig}(\varphi)$.* $\dashv$

Recall that '$\equiv_\Sigma$' is logical equivalence relative to a signature Σ (also called *semantical Σ-inseparability*) defined as: $\varphi \equiv_\Sigma \varphi'$ if for every model $\mathfrak{M}$ of φ, there is a model $\mathfrak{M}'$ of φ' such that $\mathfrak{M}|_\Sigma = \mathfrak{M}'|_\Sigma$, and *vice versa*.

Proof. Let φ be as in the proposition. We describe how to construct φ_0. The idea is to replace in a bottom-up fashion all of φ's subformulas that are not

[1] A logic L_1 is *exponentially more succinct* than the logic L_2 if there is an sequence $\varphi_1, \varphi_2, \ldots$ of L_1-formulas such that for every sequence $\psi_1, \psi_2, \ldots$ of pairwise equivalent L_2-formulas (i.e. $\varphi_i \equiv \psi_i$ for all $i \geq 0$), it holds that $|\varphi_i| = \mathcal{O}(2^{|\psi_i|})$.

in $S5^r_0$ by fresh propositional variables, which are then set to be equivalent to the replaced subformula using a formula from $S5^r_0$. We proceed by inductively computing sequences $\varphi_0, \varphi_1, \ldots$ and $\chi_1, \chi_2, \ldots$ of $S5^r$-formulas as follows: Set $\varphi_0 := \varphi$. Having computed φ_i, choose a subformula of φ_i of the form $[\psi_i]\theta_i$, where ψ_i and θ_i are $S5^r_0$-formulas. If there is no such subformula, φ_i is a formula of $S5^r_0$ and we are done. Otherwise, obtain φ_{i+1} from φ_i by replacing $[\psi_i]\theta_i$ with a fresh propositional variable p_i. Let $\chi_i = [\top](p_i \leftrightarrow [\psi_i]\theta_i)$. The procedure terminates in n steps, where $n \leq |\varphi|$. Whereas φ_n is a formula of $S5^r_0$, the formulas χ_i are not. So, for every χ_i with $1 \leq i \leq n$, we construct an $S5^r_0$-formula χ'_i that is equivalent to χ_i. First observe that the extension $[\![[\psi]\theta]\!]_{\mathfrak{M}}$ of a $S5^r$-formula of the form $[\psi]\theta$ in a model $\mathfrak{M}$ can be characterised as follows:

$$[\![[\psi]\theta]\!]_{\mathfrak{M}} = \begin{cases} W & \text{if } [\![\theta]\!]_{\mathfrak{M}} = W \\ \emptyset & \text{if } [\![\theta]\!]_{\mathfrak{M}} \neq W \text{ and } ([\![\psi]\!]_{\mathfrak{M}} \neq \emptyset \text{ implies } [\![\psi]\!]_{\mathfrak{M}} \not\subseteq [\![\theta]\!]_{\mathfrak{M}}) \\ [\![\psi]\!]_{\mathfrak{M}} & \text{if } [\![\theta]\!]_{\mathfrak{M}} \neq W,\ [\![\psi]\!]_{\mathfrak{M}} \neq \emptyset \text{ and } [\![\psi]\!]_{\mathfrak{M}} \subseteq [\![\theta]\!]_{\mathfrak{M}} \end{cases} \quad (3)$$

Set $\chi'_i = \bigwedge_{j=1..4} \chi^j_i$, where:

$$\begin{aligned}
\chi^1_i &= [\top]p_i \vee [\top]\neg p_i \vee ([\top](p_i \leftrightarrow \psi_i) \wedge \langle\top\rangle p_i \wedge \langle\top\rangle\neg p_i) \\
\chi^2_i &= [\top]p_i \rightarrow [\top]\theta_i \\
\chi^3_i &= [\top]\neg p_i \rightarrow \langle\top\rangle\neg\theta_i \wedge (\langle\top\rangle\psi_i \rightarrow \langle\top\rangle(\psi_i \wedge \neg\theta_i)) \\
\chi^4_i &= [\top](p_i \leftrightarrow \psi_i) \wedge \langle\top\rangle p_i \wedge \langle\top\rangle\neg p_i \rightarrow \langle\top\rangle\neg\theta_i \wedge \langle\top\rangle\psi_i \wedge [\top](\psi_i \rightarrow \theta_i)
\end{aligned}$$

It is readily checked that $\chi_i \equiv \chi'_i$. The three disjuncts of χ^1_i reflect the three possible values of $[\psi_i]\theta_i$'s extension in a model. The disjuncts are mutually exclusive. The conjuncts χ^2_i, χ^3_i and χ^4_i express the necessary conditions for each of the three possible values of $[\psi_i]\theta_i$'s extension. Set $\varphi' = \varphi_n \wedge \bigwedge_{i=1..n} \chi'_i$. Clearly, φ' is a formula of $S5^r_0$. It is readily checked that $\varphi \equiv_\Sigma \varphi'$ with $\Sigma = \mathsf{sig}(\varphi)$, and that the reduction leads to a blow-up of at most quadratic in the size of the formula. This concludes the proof of Proposition 2. □

Recall that $S5^r_0$ is a notational variant of the (one-agent system) S5. As a corollary from Proposition 2, we obtain that several interesting or desirable properties of S5 carry over to $S5^r$. It follows that: $S5^r$ has the polynomially-bounded finite model property. The computational complexity of the satisfiability problem for $S5^r$ is NP-complete, i.e. no worse than that for propositional logic [13,4]. $S5^r$'s model checking problem is polynomial in the size of the formula and model. Moreover, every $S5^r$-formula can be translated to an equivalent formula in $S5^r_0$ without any nesting of modal operators [16].

We conclude this section by pointing out an interesting similarity between the modalities of our logic $S5^r$ and the difference operator D. A discussion and axiomatisation of the difference operator can be found in, e.g., [2], its meta theory in [19] and its relation to hybrid temporal logic is investigated in [1]. Recall that the formula $\mathsf{D}\varphi$ states that φ holds at some point that is different to the current one. That is, the difference operator is rather expressive, e.g., the universal modality can be defined in terms of it as $\varphi \vee \mathsf{D}\varphi$. This means that the modalities $[\psi]$ of $\mathrm{S5}^r$ can be expressed in terms of the difference operator (cf. Axiom (R) in Section 4, where $[\top]$ is the universal modality). The other way around, however, does not hold true, i.e., the difference operator cannot be expressed in $\mathrm{S5}^r$, as a simple complexity-theoretic argument will show: While the satisfiability problem of $\mathrm{S5}^r$ is NP-complete, it is PSPACE-complete for the logic of the difference operator [13]. The difference operator can be defined in terms of the universal modality and nominals [1], but the latter are not available in $\mathrm{S5}^r$. Despite the difference modality being more expressive, an intriguing similarity between the two modalities shows up when comparing the extensions of formulas of the form $\langle\psi\rangle\theta$ and $\mathsf{D}\varphi$. In both cases, the values of their extension fall in solely three categories (cf. the $\langle\psi\rangle$-version of Equation (3) and [1]). That is, the extension is either: (i) the entire domain W of the model; (ii) the empty set; or (iii) $W \setminus [\![\psi]\!]$ and $W \setminus [\![\varphi]\!]$, respectively. Note that Case (iii) for $\mathsf{D}\varphi$ applies if, and only if, φ is a nominal (i.e., $[\![\mathsf{D}\varphi]\!] = W \setminus \{x\}$ iff $[\![\varphi]\!] = \{x\}$, for some point x in W). Even though in $\mathrm{S5}^r$ we cannot specify nominals, it seems, intuitively, that we could understand $\mathrm{S5}^r$ as the logic of the difference operator "modulo" nominals. The precise relationship between these two logics remains to be studied.

6 Distributed Knowledge

In this section, we extend $\mathrm{S5}^r$ with modalities for distributed hypotheses that are analogous to modalities for distributed knowledge. Distributed knowledge in modal logic is a well-known notion; standard references include [9,16] and for a more recent discussion, see [11,20]. We show how distributed hypotheses can be used to represent the knowledge of an agent whose epistemic capacity corresponds to any system containing S4.

Definition 4 (Syntax of $\mathbf{S5D}^r$). *Let Π be a countably infinite set of atomic propositions. Formulas φ of $\mathit{S5D}^r$ are defined inductively over Π by the following grammar:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid [\varphi]_K\,\varphi \mid [\Phi]_D\,\varphi,$$

where p ranges over atomic propositions in Π, and Φ over finite sets of $\mathit{S5D}^r$-formulas. ⊣

To improve readability we index the modalities with 'K' and 'D' to indicate that they mean knowledge and distributed knowledge, respectively.

Formulas of $\mathrm{S5D}^r$ are evaluated in basic structures as well. The operators $[\Phi]_D$ are necessities depending on the formulas in Φ. The semantics of $[\Phi]_D$ is based on the relations R_φ, where $\varphi \in \Phi$, as follows.

Definition 5 (Semantics of $\mathbf{S5D}^r$). *Let $\mathfrak{M} = (W, V)$ be a basic structure. The logical consequence relation '$\models$' and the relations R for formulas of $\mathit{S5D}^r$ are defined as for $\mathit{S5}^r$ but extended with the following clauses: For all $\mathit{S5D}^r$-formulas ψ and all finite sets Φ of $\mathit{S5D}^r$-formulas,*

– $\mathfrak{M}, w \models [\Phi]_D \psi$ *iff for all* $v \in W$ *with* $(w, v) \in R_\Phi$, $\mathfrak{M}, v \models \psi$,

where $R_\Phi = \bigcap_{\varphi \in \Phi} R_\varphi$. ⊣

The following lemma shows how a preorder (a reflexive and transitive relation) can be represented as an intersection of relations that are one-step total preorders. The latter are the type of relations that are determined by a formula in a model (cf. Proposition 1). The proof uses the binary operator '$\otimes$' introduced in Section 2. Before we state the lemma, we introduce an auxiliary notion. Let R be a binary relation over a set W and let $w \in W$. The R-image $R(w)$ at w is defined as $R(w) := \{v \in W \mid wRv\}$.

Lemma 1. *Let* $W = \{w_1, .., w_k\}$ *be a set. Let* R *be a preorder over* W*. For all* $i \in \{1, ..., k\}$*, let* $R_i = (W \setminus R(w_i)) \otimes R(w_i)$*. Then it holds that* $R = \bigcap_{i=1,..,n} R_i$. ⊣

Proof. Let the set W and the relations $R, R_1, \ldots, R_k$ be as in the lemma. We show that $R = \bigcap_{i=1..k} R_i$ holds. First consider '$\subseteq$'. Suppose $(w_i, w_j) \in R$. We need to show that $(w_i, w_j) \in R_\ell$ for all $\ell \in \{1, ..., k\}$. Suppose not, i.e. $(w_i, w_j) \notin R_\ell$ for some ℓ. We obtain $w_i \in R(w_\ell)$ by definition of R_ℓ. That is, w_i is an element of the R-image at w_ℓ. Then we have $(w_\ell, w_i) \in R$. But then, together with the assumption $(w_i, w_j) \in R$, it follows by transitivity of R that $(w_\ell, w_j) \in R$ – a contradiction.

Consider the other direction '$\supseteq$'. Suppose $(w_i, w_j) \in \bigcap_{\ell=1..k} R_\ell$. That is, $(w_i, w_j) \in R_\ell$ for all $\ell \in \{1, ..., k\}$. In particular, we have that $(w_i, w_j) \in R_i$. We obtain $w_i \in R(w_i)$ by reflexivity of R. But then the construction of R_i yields that $w_j \in R(w_i)$. That is, w_j belongs to the R-image at w_i. Thus $(w_i, w_j) \in R$. □

To give an intuition for understanding the lemma, observe that a preorder R induces a partial order (i.e. an antisymmetic preorder) on the set of R-clusters, which are sets of points fully connected by R. In other words, R gives rise to a collection of directed graphs whose nodes are R-clusters. Note that the graph is loopless (and thus antisymmetric). Now, if R is total, all points are connected which gives rise to just one such graph. If additionally R is 'one-step', the graph consists of merely two nodes. Intersecting one-step total preorders has the effect of erasing some directed edges from the universal relation. It is not hard to check that the intersection of preorders is again a preorder. Lemma 1 shows that by intersecting a certain selection of one-step total preorders, we can "carve out" the desired preorder. The following example illustrates the scenario.

Example 3. Let $W = \{x, y, z\}$ be a set and $R = \{(x, y), (x, z)\} \cup \mathsf{id}(W)$. It is readily checked that R is a reflexive and transitive relation. Now let $R_w = (W \setminus R(w)) \otimes R(w)$ for all $w \in W$. That is, according to Equation (1), we have $R_x = W \times W$, $R_y = \{(x, y), (z, y), (x, z), (z, x)\} \cup \mathsf{id}(W)$ and $R_z = \{(y, z), (x, z), (x, y), (y, x)\} \cup \mathsf{id}(W)$. Intersecting these relations we obtain $R_x \cap R_y \cap R_z = \{(x, y), (x, z)\} \cup \mathsf{id}(W)$, which is equivalent to R. ◁

The intersection in Lemma 1 reminds us on the relations R_Φ determined by a finite set Ψ of S5D^r-formulas in a model. In fact, this is the connection we seek to establish in order to represent the knowledge of an agent as distributed knowledge. In the following, we state how this is done.

Take an arbitrary uni-modal logic L between S4 and S5 (whose satisfaction relation is denoted by $\models_{\mathrm{L}}$). The necessity operator '$\Box$' of L is thought of as representing the knowledge of the agent. Note that the system L contains the axioms (T) and (4), each of which represent important epistemic properties, namely, veridicality and positive introspection, respectively. Of course, L may contain other axioms, in fact, any axiom that can be derived in system S5. For instance, prominent axioms that are considered relevant for epistemics are:

(.2) $\neg\Box\neg\Box p \rightarrow \Box\neg\Box\neg p$,
(.3) $\Box(\Box p \rightarrow \Box q) \vee \Box(\Box q \rightarrow \Box p)$,
(.4) $p \rightarrow (\neg\Box\neg\Box p \rightarrow \Box p)$.

We assume that L is determined by a class $\mathcal{C}$ of Kripke structures (i.e., the theorems of L are exactly the formulas that are valid on all structures in $\mathcal{C}$). The class $\mathcal{C}$ is not required to be first-order definable or definable in any other formalism. In fact, $\mathcal{C}$ may be given by manual selection. Clearly, the structures in $\mathcal{C}$ are reflexive and transitive. What we require as a precondition is that L has the finite-model property wrt. $\mathcal{C}$. This means that, if a formula φ is not a theorem of L then there is a finite Kripke structure $\mathfrak{M}^k$ in $\mathcal{C}$ that falsifies φ, i.e. $\mathfrak{M}^k, w \not\models \varphi$ for some world w in $\mathfrak{M}^k$.

Before we can state the theorem, we need one more auxiliary notion. Let $\mathfrak{M}^k = (W, R, V)$ be a finite Kripke structure such that the relation R is a preorder. We say that the valuation function V *covers* R if for every world $w \in W$, there is an atomic proposition p_w such that $V(p_w) = R(w)$, i.e. the R-image at w.

Theorem 2. *Let $\mathcal{C}$ be a class of Kripke structures whose relations are preorders. Let $\mathfrak{M}^k = (W, R, V)$ be a finite structure from $\mathcal{C}$ such that V covers R. Let $\mathfrak{M} = (W, V)$ be a basic structure and let $w \in W$ be a world. Let φ be a Boolean formula over Π. Then, there is a finite set Ψ of atomic propositions such that the following are equivalent:*

(i) $\mathfrak{M}^k, w \models_L \Box\varphi$;
(ii) $\mathfrak{M}, w \models_{S5D^r} [\Psi]_D\,\varphi$. ⊣

Proof. For every $w \in W$, select an atomic proposition p_w such that $V(p_w) = R(w)$. Note that such p_w exists since V covers R. Set $\Psi = \{p_w \mid w \in W\}$. Using Lemma 1 the equivalence of (i) and (ii) can be shown by induction on the structure of φ. □

We remark that the theorem can be generalised since the condition of using finite models is a bit too strict. Recall the metaphor that views a preorder R as a collection of loopless graphs whose nodes are R-clusters. What is actually required is that the collection of graphs and the graphs themselves are finite. So, we can still find a finite intersection of relations as desired.

The following example illustrates Theorem 2 and discusses the presented notions.

Example 4. Consider the Kripke model $\mathfrak{M}^k = (W, R, V)$, where W and R are as in Example 3, and $V(p) = \{x, z\}$ and $V(q) = \{z\}$. Clearly, $\mathfrak{M}^k$ is not an S5-model as R is not symmetric. Let $\varphi_{.2}$, $\varphi_{.3}$, $\varphi_{.4}$ be the instances of the axioms (.2), (.3) and (.4) as shown above. It turns out that only $\varphi_{.3}$ holds at x, but not $\varphi_{.2}$ nor $\varphi_{.4}$. In fact, $\varphi_{.3}$ holds at all worlds in $\mathfrak{M}^k$. Let us assume that the box (i.e. the epistemic capacity of the agent) is characterised by the system S4.3.

Now label the worlds with fresh atomic propositions p_x, p_y, p_z, i.e., we set $V'(p_w) = \{w\}$ for all $w \in W$. Notice that V' covers R. Let R_{p_x}, R_{p_y}, R_{p_z} be the relations determined by the basic structure $\mathfrak{M} = (W, V')$ and the fresh propositions (cf. Definition 2). Notice that R_{p_w} equals R_w from Example 3, for every $w \in W$. Thus $R_{p_x} \cap R_{p_y} \cap R_{p_z} = R$. Now it is immediate that $\mathfrak{M}, w \models [\{p_x, p_y, p_z\}]_D \varphi$ iff $\mathfrak{M}^k, w \models \Box\varphi$, for all $w \in W$ and all propositional formulas φ without occurrence of any of p_x, p_y and p_z. In other words, $[\{p_x, p_y, p_z\}]_D$ simulates the S4.3-box. We can see p_x, p_y, p_z as hypotheses that another agent has to adopt in order to know what the S4.3-agent knows.

In some cases we have an alternative to introducing fresh propositions even though V does not cover R. This means that V covering R is a sufficient but not necessary condition for Theorem 2. Here $\neg p$ and q are hypotheses so that $[\{\neg p, q\}]_D$ simulates S4.3-box as well. That is, hypotheses do not need to be atomic propositions. Moreover, (parts of) hypotheses may occur in the conclusion as in $(W, V), x \models [\{\neg p, q\}]_D \varphi_{.3}$. A candidate for a refined notion of V covering R is the condition that, for every $w \in W$, there is a S5D^r-formula φ_w such that $[\![\varphi_w]\!] = R(w)$. In this paper, however, we do not explore this notion further. $\lhd$

7 Conclusion

In this paper, we introduced the logic S5^r, which is both an extension of S5 and a special case of Chella's Conditional Logic. We presented an axiomatisation for this logic and showed that it is as expressive and complex as S5. However, it turns out that unlike Public Announcement Logic S5^r is not exponentially more succinct than S5. Nevertheless we argue that S5^r is a more intuitive formalism for describing and reasoning about knowledge under hypotheses. In the second part of the paper, we extend the logic with modalities for distributed hypotheses that are analog to modalities for distributed knowledge. We showed how distributed hypotheses can be employed to represent the knowledge of agents whose epistemic capacity corresponds to any system containing S4. Possible directions for future work are to investigate S5D^r in more detail and axiomatise it, generalise Theorem 2 and to investigate an abductive reasoning service, e.g.: Given an agent a with certain epistemic capacity and that a knows φ, compute "appropriate" hypotheses such that the combined or distributed hypotheses imply φ in S5D^r.

Acknowledgements. We would like to thank the anonymous reviewers for their useful comments. This work was supported in part by the Spanish Ministry of Science and Innovation through the Juan de la Cierva programme, the project "Agreement Technologies" (Grant CONSOLIDER CSD2007-0022, INGENIO 2010), and the MICINN projects TIN2006-15455 and TIN2009-14562-CO5.

References

1. Areces, C., Blackburn, P., Marx, M.: The computational complexity of hybrid temporal logics. Logic Journal of the IGPL 8, 653–679 (2000)
2. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. Cambridge University Press, Cambridge (2001)
3. Boolos, G.S.: The Logic of Provability. Cambridge University Press, Cambridge (1993)
4. Chagrov, A., Zakharyaschev, M.: Modal Logic, Oxford Logic Guides, vol. 35. Clarendon Press, Oxford (1997)
5. Chellas, B.F.: Basic conditional logic. Journal of Philosophical Logic 4, 133–153 (1975)
6. Chellas, B.F.: Modal Logic: An Introduction. Cambridge University Press, Cambridge (1980)
7. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Synthese Library: Studies in Epistemology, Logic, Methodology, and Philosophy of Science, vol. 337. Springer, Dordrecht (2007)
8. Esakia, L.: Around provability logic. Annals of Pure and Applied Logic 161(2), 174–184 (2009); Festschrift on the occasion of Franco Montagna's 60th birthday
9. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning about Knowledge. The MIT Press, Cambridge (1995)
10. Fagin, R., Halpern, J.Y., Vardi, M.Y.: What can machines know?: On the properties of knowledge in distributed systems. Journal of the ACM 39, 328–376 (1992)
11. Gerbrandy, J.: Distributed knowledge. In: Twendial 1998: Formal Semantics and Pragmatics of Dialogue. TWLT, vol. 13, pp. 111–124. Universiteit Twente, Enschede (1998)
12. Halpern, J.Y., Moses, Y.: A guide to completeness and complexity for modal logics of knowledge and belief. Artificial Intelligence 54, 319–379 (1992)
13. Ladner, R.E.: The computational complexity of provability in systems of modal propositional logic. SIAM Journal of Computing 6, 467–480 (1977)
14. Lewis, D.: Counterfactuals. Harvard University Press, Cambridge (1973) (revised printing Blackwell)
15. Lutz, C.: Complexity and succinctness of public announcement logic. In: Proceedings of AAMAS 2006, pp. 137–143. ACM, New York (2006)
16. Meyer, J.J.C., van der Hoek, W.: Epistemic Logic for AI and Computer Science. Cambridge Tracts in Theoretical Computer Science, vol. 41. Cambridge University Press, Cambridge (1995)
17. Nute, D.: Conditional logic. In: Handbook of Philosophical Logic, vol. 2, pp. 387–439. D. Reidel Publishing Company, Dordrecht (1984)
18. Plaza, J.A.: Logics of public communications. In: Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems, pp. 201–216. Oak Ridge National Laboratory, ORNL/DSRD-24 (1989)
19. de Rijke, M.: The modal logic of inequality. The Journal of Symbolic Logic 57(2), 566–584 (1992)
20. Roelofsen, F.: Distributed knowledge. Journal of Applied Non-Classical Logics 17(2), 255–273 (2007)
21. Segerberg, K.: Notes on conditional logic. Studia Logica 48, 157–168 (1989)
22. Stalnaker, R.: A theory of conditionals. Studies in Logical Theory, 98–112 (1968)