

Indicative Conditionals and Dynamic Epistemic Logic

Wesley H. Holliday

University of California, Berkeley
wesholliday@berkeley.edu

Thomas F. Icard, III

Stanford University
icard@stanford.edu

Recent ideas about epistemic modals and indicative conditionals in formal semantics have significant overlap with ideas in modal logic and dynamic epistemic logic. The purpose of this paper is to show how greater interaction between formal semantics and dynamic epistemic logic in this area can be of mutual benefit. In one direction, we show how concepts and tools from modal logic and dynamic epistemic logic can be used to give a simple, complete axiomatization of Yalcin's [16] semantic consequence relation for a language with epistemic modals and indicative conditionals. In the other direction, the formal semantics for indicative conditionals due to Kolodny and MacFarlane [9] gives rise to a new dynamic operator that is very natural from the point of view of dynamic epistemic logic, allowing succinct expression of *dependence* (as in dependence logic) or *supervenience* statements. We prove decidability for the logic with epistemic modals and Kolodny and MacFarlane's indicative conditional via a full and faithful computable translation from their logic to the modal logic K45.

Logic and the formal semantics of natural language are related by blood and yet somewhat estranged. Today it is rare that formal semanticists consider questions of axiomatizability or decidability of the consequence relations defined by model-theoretic accounts of natural language fragments. Meanwhile logicians focus more on logics motivated by mathematical or philosophical concerns than on logics arising from semantic theories in linguistics.

The cost of estrangement is that insights from one field that would be useful for the other may go unnoticed or efforts may be unnecessarily duplicated. The aim of this paper is to help encourage a family reunion between logic and formal semantics of natural language, by way of concrete examples. The topic of modals and conditionals is a prime example of overlap between formal semantics and logic. In this paper, we consider the case of epistemic modals and indicative conditionals.

In §1, we show how concepts and tools from modal logic and dynamic epistemic logic can be used to give a simple, complete axiomatization of Yalcin's [16] semantic consequence relation for a language with epistemic modals and indicative conditionals. Then in §2, we show that the formal semantics for indicative conditionals due to Kolodny and MacFarlane [9] gives rise to a new dynamic operator that is very natural from the point of view of dynamic epistemic logic, allowing succinct expression of *dependence* (as in dependence logic) or *supervenience* statements. We prove decidability for the logic with epistemic modals and Kolodny and MacFarlane's indicative conditional via a full and faithful computable translation from their logic to the modal logic K45.

There are other examples of clear overlap between formal semantics and dynamic epistemic logic, such as the connection between the *dynamic logical consequence* of [14] and the *dynamic consequence* of [1], between the notions of *epistemic contradictions* in [15] and of *Moorean sentences* in [6], and more. Thus, the examples to follow by no means exhaust the connections to be made between formal semantics and dynamic epistemic logic.

1 Applying DEL to Formal Semantics

Throughout we work with the language $\mathcal{L}(\Rightarrow)$ defined by:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid \Box\varphi \mid (\varphi \Rightarrow \psi),$$

where p comes from a fixed set of propositional variables. The connectives \vee , \rightarrow , \leftrightarrow , \perp , and \diamond are defined as usual, so $\diamond\varphi := \neg\Box\neg\varphi$. In the intended interpretation, \Box and \diamond stand for “must” and “might”, and \Rightarrow stands for the indicative conditional “if... then”. \mathcal{L} is the set of formulas that do not contain \Rightarrow . *Nonmodal* formulas are formulas of \mathcal{L} that do not contain \Box (and hence \diamond).

We begin by reviewing Yalcin’s [16] semantics for $\mathcal{L}(\Rightarrow)$. Models are tuples $\mathcal{M} = \langle W, V \rangle$ where W is a nonempty set and V is a function assigning to each propositional variable p a proposition $V(p) \subseteq W$. Formulas are evaluated in a model \mathcal{M} at a world $w \in W$ relative to an *information state* $X \subseteq W$ as follows:

- $\mathcal{M}, w, X \models p$ iff $w \in V(p)$;
- $\mathcal{M}, w, X \models \neg\varphi$ iff $\mathcal{M}, w, X \not\models \varphi$;
- $\mathcal{M}, w, X \models \varphi \wedge \psi$ iff $\mathcal{M}, w, X \models \varphi$ and $\mathcal{M}, w, X \models \psi$;
- $\mathcal{M}, w, X \models \Box\varphi$ iff for all $v \in X$, $\mathcal{M}, v, X \models \varphi$;
- $\mathcal{M}, w, X \models \diamond\varphi$ iff for some $v \in X$, $\mathcal{M}, v, X \models \varphi$;
- $\mathcal{M}, w, X \models \varphi \Rightarrow \psi$ iff $\mathcal{M}, w, [\varphi]^{\mathcal{M}, X} \models \Box\psi$,

where $[\varphi]^{\mathcal{M}, X} = \{v \in X \mid \mathcal{M}, v, X \models \varphi\}$.¹

A formula φ is *valid* iff it is true at every world relative to every information state in every model. To define *consequence*, first let $\mathcal{M}, X \models \varphi$ (X “accepts” φ) iff for all $w \in X$, we have $\mathcal{M}, w, X \models \varphi$. Then for a set Σ of formulas, Yalcin defines φ to be an *informational consequence* of Σ iff for every model $\mathcal{M} = \langle W, V \rangle$ and information state $X \subseteq W$, if $\mathcal{M}, X \models \sigma$ for all $\sigma \in \Sigma$, then $\mathcal{M}, X \models \varphi$. (If φ is valid, then φ is an informational consequence of \emptyset , but the converse fails for, e.g., $p \vee \diamond\neg p$.)

The above “domain semantics” for \Box and \diamond is presented by Yalcin [15] as an alternative to standard Hintikka-style relational semantics for epistemic logic. Conceptually, the two semantics are different. Mathematically, the domain semantics is equivalent to a special case of the relational semantics, using what we might call *uniform* relational models $\langle W, R, V \rangle$ in which every two worlds have the same set of R -successors. In fact, since the definition of “informational consequence” only involves pairs w, X such that $w \in X$, it is as if we were working with relational models in which R is the universal relation. In this setting, informational consequence is equivalent to the notion of *global consequence* from modal logic (see [3, §1.5]). Thus, from the point of view of modal logic, the above semantics for \Box and \diamond can be summarized as follows: the proposed consequence relation for epistemic modals is the *global consequence* relation over *universal* relational models according to the standard semantics (cf. [12]).²

As for the indicative conditional \Rightarrow , this is where dynamic epistemic logic [5, 2] enters the story. Consider a system of dynamic epistemic logic that extends the language of propositional modal logic with formulas of the form $[\varphi]\psi$, intuitively interpreted as “after *information update* with φ , ψ is the case.” Starting with relational models $\langle W, R, V \rangle$, the semantics for $[\varphi]\psi$ as in [5, §4.9] is:

¹Equivalently, we could define $[\varphi]^{\mathcal{M}, X} = \{v \in W \mid \mathcal{M}, v, X \models \varphi\}$ and $\mathcal{M}, w, X \models \varphi \Rightarrow \psi$ iff $\mathcal{M}, w, X \cap [\varphi]^{\mathcal{M}, X} \models \Box\psi$, but we prefer less notation. Also note that if one wants to require that information states be nonempty, then the clause for \Rightarrow must be changed, e.g., to $\mathcal{M}, w, X \models \varphi \Rightarrow \psi$ iff $\mathcal{M}, w, X \not\models \diamond\varphi$ or $\mathcal{M}, w, [\varphi]^{\mathcal{M}, X} \models \Box\psi$.

²Yalcin [15] also proposes domain semantics for a language with both epistemic modals and attitude verbs such as ‘believe’ and ‘suppose’, in which case the comparison of domain semantics and relational semantics is not as straightforward.

- $\langle W, R, V \rangle, w \models [\varphi]\psi$ iff $\langle W, R^\varphi, V \rangle, w \models \psi$,

where R^φ is defined by: $wR^\varphi v$ iff wRv and $\langle W, R, V \rangle, v \models \varphi$. A superficial difference between this semantics and that of \Rightarrow above is that here we are “changing the model,” whereas above we “shifted the information state.”³ The semantics for $\varphi \Rightarrow \psi$ is in fact equivalent to the semantics for $[\varphi]\Box\psi$. The form $[\varphi]\Box\psi$ has been studied extensively in dynamic epistemic logic, where the main interest is in reasoning about what is known or believed after information update, so \Rightarrow turns out to be a familiar modality.

Elsewhere [7, 8] we have argued that it is valuable not only in logic but also in natural language semantics to accompany a formal semantic proposal with a complete *axiomatization* (when possible), capturing basic entailment predictions of the semantics from which all of its other entailment predictions may be derived. In this spirit, and taking advantage of the connections with modal and dynamic epistemic logic above, we establish as our first main result a complete axiomatization of the logic of Yalcin’s [16] semantics for epistemic modals and indicative conditionals.⁴

Theorem 1.1. The set of $\mathcal{L}(\Rightarrow)$ formulas that are valid according to Yalcin’s semantics is the Yalcin logic: the smallest set of formulas that is closed under *replacement of equivalents*,⁵ *modus ponens* for the material conditional \rightarrow , and *necessitation* for \Box , and contains all substitution instances of propositional tautologies and all instances of the axioms in Figure 1. Moreover, φ is an informational consequence of $\{\sigma_1, \dots, \sigma_n\}$ iff $(\Box\sigma_1 \wedge \dots \wedge \Box\sigma_n) \rightarrow \Box\varphi$ is a theorem of the Yalcin logic.

$$\begin{array}{ll}
\mathbf{K} & \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi) \\
\mathbf{4} & \Diamond\Diamond\varphi \rightarrow \Diamond\varphi \quad \mathbf{5} \quad \Diamond\Box\varphi \rightarrow \Box\varphi \\
\mathbf{I1} & (\varphi \Rightarrow \pi) \leftrightarrow \Box(\varphi \rightarrow \pi) \text{ for } \pi \text{ nonmodal} \\
\mathbf{I2} & (\varphi \Rightarrow (\alpha \wedge \beta)) \leftrightarrow ((\varphi \Rightarrow \alpha) \wedge (\varphi \Rightarrow \beta)) \\
\mathbf{I3} & (\varphi \Rightarrow \alpha) \rightarrow (\varphi \Rightarrow (\alpha \vee \beta)) \\
\mathbf{I4} & (\varphi \Rightarrow \alpha) \rightarrow (\varphi \Rightarrow \Box\alpha) \\
\mathbf{I5} & ((\varphi \Rightarrow (\alpha \vee \Box\beta)) \wedge \neg(\varphi \Rightarrow \beta)) \rightarrow (\varphi \Rightarrow \alpha) \\
\mathbf{I6} & ((\varphi \Rightarrow (\alpha \vee \Diamond\beta)) \wedge (\varphi \Rightarrow \neg\beta)) \rightarrow (\varphi \Rightarrow \alpha) \\
\mathbf{I7} & \neg(\varphi \Rightarrow \beta) \rightarrow (\varphi \Rightarrow \Diamond\neg\beta)
\end{array}$$

Figure 1: Axioms of the Yalcin logic.

Before proving Theorem 1.1, let us observe how the axioms in Figure 1 can be seen as corresponding to natural language inferences. The following examples are taken verbatim from [8].

³The semantics of dynamic epistemic logic can be equivalently repackaged by pulling the relation R out of the model, so we would define ‘ $\langle W, V \rangle, w, R \models \varphi$ ’, and then in the case where R is uniform, by replacing R with the image $X = R[W] = \{v \in W \mid \exists w \in W : wRv\}$, so we would define ‘ $\langle W, V \rangle, w, X \models \varphi$ ’ as in domain semantics.

⁴Bledin [4] also proposes a logic for Yalcin’s modal-indicative semantics, but with a Fitch-style natural deduction system for an extension of the language that includes symbols for information states and information acceptance relations.

⁵Closure under replacement of equivalents means that if $\alpha \leftrightarrow \beta \in L$, and φ' is obtained from φ by replacing some occurrence of α in φ by β , then $\varphi \leftrightarrow \varphi' \in L$.

Example 1.2. Axiom **I4** corresponds to a key prediction of the semantics that a sentence like

- (1) If Miss Scarlet didn't do it, then Colonel Mustard did it.

entails

- (2) If Miss Scarlet didn't do it, then it must be that Colonel Mustard did it.

Axiom **I5** suggests the prediction that the sentence

- (3) If Miss Scarlet did it, then either Colonel Mustard was her accomplice or it must be that Professor Plum was involved.

together with

- (4) It's not the case that if Miss Scarlet did it, then Professor Plum was involved.

entails

- (5) If Miss Scarlet did it, then Colonel Mustard was her accomplice.

And axiom **I6** suggests the prediction that

- (6) If Miss Scarlet did it, then either she used the pipe or she might have used the candlestick.

together with

- (7) If Miss Scarlet did it, she didn't use the candlestick.

entails

- (8) If Miss Scarlet did it, then she used the pipe.

Finally, axiom **I7** suggests the prediction that one who *rejects*

- (9) If Miss Scarlet was in the ballroom, then Colonel Mustard is guilty.

should accept

- (10) If Miss Scarlet was in the ballroom, it might be that Colonel Mustard is not guilty.

We will prove Theorem 1.1 using several lemmas. The first lemma follows from the well-known fact that K45 is the logic of the “uniform” relational models mentioned above, plus the equivalence of domain semantics and uniform relational semantics for \mathcal{L} .

Lemma 1.3. The set of \mathcal{L} formulas that are valid according to Yalcin's semantics is the logic K45.

The next two lemmas show that the conditional axioms of the Yalcin logic, which we have presented in their form in Figure 1 to bring out connections with natural language inference as in Example 1.2, can be equivalently repackaged as valid *reduction axioms* in the sense of dynamic epistemic logic [5].

Lemma 1.4. The following axioms are valid according to Yalcin's semantics:

$$\mathbf{A1} \quad (\varphi \Rightarrow \pi) \leftrightarrow \Box(\varphi \rightarrow \pi) \text{ for } \pi \text{ a nonmodal formula;}$$

$$\mathbf{A2} \quad (\varphi \Rightarrow (\alpha \wedge \beta)) \leftrightarrow ((\varphi \Rightarrow \alpha) \wedge (\varphi \Rightarrow \beta));$$

$$\mathbf{A3} \quad (\varphi \Rightarrow (\alpha \vee \Box\beta)) \leftrightarrow ((\varphi \Rightarrow \alpha) \vee (\varphi \Rightarrow \beta));$$

$$\mathbf{A4} \quad (\varphi \Rightarrow (\alpha \vee \Diamond\beta)) \leftrightarrow ((\varphi \Rightarrow \alpha) \vee \neg(\varphi \Rightarrow \neg\beta)).$$

Proof. For **A1**, suppose π is nonmodal. By definition, we have

$$\begin{aligned} \mathcal{M}, w, X \models \varphi \Rightarrow \pi & \text{ iff } \mathcal{M}, w, [\varphi]^{\mathcal{M}, X} \models \Box \pi \\ & \text{ iff } [\varphi]^{\mathcal{M}, X} \subseteq [\pi]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}}. \end{aligned} \quad (1)$$

Since π is nonmodal, we have

$$[\pi]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} = [\varphi]^{\mathcal{M}, X} \cap [\pi]^{\mathcal{M}, X},$$

so the right side of (1) is equivalent to

$$[\varphi]^{\mathcal{M}, X} \subseteq [\pi]^{\mathcal{M}, X},$$

which by definition is equivalent to $\mathcal{M}, w, X \models \Box(\varphi \rightarrow \pi)$.

For **A2**, by definition we have

$$\begin{aligned} \mathcal{M}, w, X \models \varphi \Rightarrow (\alpha_1 \wedge \alpha_2) & \text{ iff } \mathcal{M}, w, [\varphi]^{\mathcal{M}, X} \models \Box(\alpha_1 \wedge \alpha_2) \\ & \text{ iff } \mathcal{M}, w, [\varphi]^{\mathcal{M}, X} \models \Box \alpha_1 \wedge \Box \alpha_2 \\ & \text{ iff } \mathcal{M}, w, [\varphi]^{\mathcal{M}, X} \models \Box \alpha_i, i = 1, 2 \\ & \text{ iff } \mathcal{M}, w, X \models \varphi \Rightarrow \alpha_i, i = 1, 2. \end{aligned}$$

For **A3**, by definition we have

$$\begin{aligned} \mathcal{M}, w, X \models \varphi \Rightarrow (\alpha \vee \Box \beta) & \text{ iff } \mathcal{M}, w, [\varphi]^{\mathcal{M}, X} \models \Box(\alpha \vee \Box \beta) \\ & \text{ iff } [\varphi]^{\mathcal{M}, X} \subseteq [\alpha \vee \Box \beta]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} \\ & \text{ iff } [\varphi]^{\mathcal{M}, X} \subseteq [\alpha]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} \cup [\Box \beta]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}}. \end{aligned} \quad (2)$$

Using (2), we first show that if $\mathcal{M}, w, X \models \varphi \Rightarrow (\alpha \vee \Box \beta)$, then $\mathcal{M}, w, X \models \varphi \Rightarrow \alpha$ or $\mathcal{M}, w, X \models \varphi \Rightarrow \beta$.

Case 1: $[\Box \beta]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} = \emptyset$. Then (2) implies

$$[\varphi]^{\mathcal{M}, X} \subseteq [\alpha]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} \Rightarrow \mathcal{M}, w, X \models \varphi \Rightarrow \alpha. \quad (3)$$

Case 2: $[\Box \beta]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} \neq \emptyset$. Then

$$\exists v \in W : \mathcal{M}, v, [\varphi]^{\mathcal{M}, X} \models \Box \beta,$$

and we have

$$\begin{aligned} \mathcal{M}, v, [\varphi]^{\mathcal{M}, X} \models \Box \beta & \text{ iff } [\varphi]^{\mathcal{M}, X} \subseteq [\beta]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} \\ & \text{ iff } \mathcal{M}, w, X \models \varphi \Rightarrow \beta. \end{aligned}$$

Next, we show that if $\mathcal{M}, w, X \models \varphi \Rightarrow \alpha$ or $\mathcal{M}, w, X \models \varphi \Rightarrow \beta$, then $\mathcal{M}, w, X \models \varphi \Rightarrow (\alpha \vee \Box \beta)$, using (2). If $\mathcal{M}, w, X \models \varphi \Rightarrow \alpha$, then $[\varphi]^{\mathcal{M}, X} \subseteq [\alpha]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}}$, which implies (2). On the other hand, if $\mathcal{M}, w, X \models \varphi \Rightarrow \beta$, then we have

$$\begin{aligned} [\varphi]^{\mathcal{M}, X} \subseteq [\beta]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}} & \Rightarrow \forall v \in X : \mathcal{M}, v, [\varphi]^{\mathcal{M}, X} \models \Box \beta \\ & \Rightarrow \forall v \in [\varphi]^{\mathcal{M}, X} : \mathcal{M}, v, [\varphi]^{\mathcal{M}, X} \models \Box \beta \\ & \Rightarrow [\varphi]^{\mathcal{M}, X} \subseteq [\Box \beta]^{\mathcal{M}, [\varphi]^{\mathcal{M}, X}}, \end{aligned}$$

which implies (2).

For **A4**, by similar reasoning to that for (2), we have

$$\mathcal{M}, w, X \models \varphi \Rightarrow (\alpha \vee \diamond\beta) \quad \text{iff} \quad \llbracket \varphi \rrbracket^{\mathcal{M}, X} \subseteq \llbracket \alpha \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} \cup \llbracket \diamond\beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}}. \quad (4)$$

Using (4), we show that if $\mathcal{M}, w, X \models \varphi \Rightarrow (\alpha \vee \diamond\beta)$, then $\mathcal{M}, w, X \models \varphi \Rightarrow \alpha$ or $\mathcal{M}, w, X \models \neg(\varphi \Rightarrow \neg\beta)$.

Case 1: $\llbracket \diamond\beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} = \emptyset$. Then as in Case 1 for **A3**, (4) implies $\mathcal{M}, w, X \models \varphi \Rightarrow \alpha$.

Case 2: $\llbracket \diamond\beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} \neq \emptyset$. Then we observe that

$$\begin{aligned} \llbracket \diamond\beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} \neq \emptyset & \quad \text{iff} \quad \exists v \in \llbracket \varphi \rrbracket^{\mathcal{M}, X} : \mathcal{M}, v, \llbracket \varphi \rrbracket^{\mathcal{M}, X} \models \diamond\beta \\ & \quad \text{iff} \quad \llbracket \varphi \rrbracket^{\mathcal{M}, X} \cap \llbracket \beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} \neq \emptyset \\ & \quad \text{iff} \quad \llbracket \varphi \rrbracket^{\mathcal{M}, X} \not\subseteq W \setminus \llbracket \beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} \\ & \quad \text{iff} \quad \llbracket \varphi \rrbracket^{\mathcal{M}, X} \not\subseteq \llbracket \neg\beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} \\ & \quad \text{iff} \quad \mathcal{M}, w, \llbracket \varphi \rrbracket^{\mathcal{M}, X} \not\models \Box\neg\beta \\ & \quad \text{iff} \quad \mathcal{M}, w, X \not\models \varphi \Rightarrow \neg\beta \\ & \quad \text{iff} \quad \mathcal{M}, w, X \models \neg(\varphi \Rightarrow \neg\beta). \end{aligned} \quad (5)$$

Finally, we show that if $\mathcal{M}, w, X \models \varphi \Rightarrow \alpha$ or $\mathcal{M}, w, X \models \neg(\varphi \Rightarrow \neg\beta)$, then $\mathcal{M}, w, X \models \varphi \Rightarrow (\alpha \vee \Box\beta)$, using (4). The argument starting from $\mathcal{M}, w, X \models \varphi \Rightarrow \alpha$ is the same as for **A3**. The argument from $\mathcal{M}, w, X \models \neg(\varphi \Rightarrow \neg\beta)$ uses (5) and the observation that

$$\begin{aligned} \llbracket \varphi \rrbracket^{\mathcal{M}, X} \cap \llbracket \beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}} \neq \emptyset & \quad \Rightarrow \quad \forall v \in X : \mathcal{M}, v, \llbracket \varphi \rrbracket^{\mathcal{M}, X} \models \diamond\beta \\ & \quad \Rightarrow \quad \forall v \in \llbracket \varphi \rrbracket^{\mathcal{M}, X} : \mathcal{M}, v, \llbracket \varphi \rrbracket^{\mathcal{M}, X} \models \diamond\beta \\ & \quad \Rightarrow \quad \llbracket \varphi \rrbracket^{\mathcal{M}, X} \subseteq \llbracket \diamond\beta \rrbracket^{\mathcal{M}, \llbracket \varphi \rrbracket^{\mathcal{M}, X}}, \end{aligned}$$

which implies (4). □

Next we verify that the axioms of Lemma 1.4 are indeed an equivalent repackaging of the axioms of the Yalcin logic.

Lemma 1.5. A formula φ is a theorem of the Yalcin logic iff φ is a theorem of the logic defined in the same way as the Yalcin logic (in Theorem 1.1) but with axioms **A1-A4** in place of **I1-I7**.

Proof. Axioms **I1** and **I2** are just axioms **A1** and **A2**, so we begin with **I3**. As an instance of **A2**, $(\varphi \Rightarrow (\delta \wedge \gamma)) \leftrightarrow ((\varphi \Rightarrow \delta) \wedge (\varphi \Rightarrow \gamma))$, we have

$$(\varphi \Rightarrow ((\alpha \vee \beta) \wedge (\alpha \vee \neg\beta))) \leftrightarrow ((\varphi \Rightarrow (\alpha \vee \beta)) \wedge (\varphi \Rightarrow (\alpha \vee \neg\beta))).$$

Since $(\alpha \vee \beta) \wedge (\alpha \vee \neg\beta)$ is equivalent to α , by replacement of equivalents the left-to-right direction of the biconditional gives us $(\varphi \Rightarrow \alpha) \rightarrow (\varphi \Rightarrow (\alpha \vee \beta))$, which is **I3**.

For **I4**, as an instance of **A3**, $(\varphi \Rightarrow (\alpha \vee \Box\beta)) \leftrightarrow ((\varphi \Rightarrow \alpha) \vee (\varphi \Rightarrow \beta))$, we have

$$(\varphi \Rightarrow (\perp \vee \Box\beta)) \leftrightarrow ((\varphi \Rightarrow \perp) \vee (\varphi \Rightarrow \beta)).$$

Since $\perp \vee \Box\beta$ is equivalent to $\Box\beta$, by replacement of equivalents the right-to-left direction of the biconditional gives us $(\varphi \Rightarrow \beta) \rightarrow (\varphi \Rightarrow \Box\beta)$, which is **I4**.

I5 is a Boolean rewriting of the left-to-right direction of **A3**, and **I6** is a Boolean rewriting of the left-to-right direction of **A4**.

Finally, for **I7**, as an instance of **A4**, $(\varphi \Rightarrow (\delta \vee \Diamond\gamma)) \leftrightarrow ((\varphi \Rightarrow \delta) \vee \neg(\varphi \Rightarrow \neg\gamma))$, we have

$$(\varphi \Rightarrow (\perp \vee \Diamond\neg\beta)) \leftrightarrow ((\varphi \Rightarrow \perp) \vee \neg(\varphi \Rightarrow \neg\neg\beta)),$$

which by replacement of equivalents is equivalent to

$$(\varphi \Rightarrow \Diamond\neg\beta) \leftrightarrow ((\varphi \Rightarrow \perp) \vee \neg(\varphi \Rightarrow \beta)),$$

the right-to-left direction of which gives us $\neg(\varphi \Rightarrow \beta) \rightarrow (\varphi \Rightarrow \Diamond\neg\beta)$, which is **I7**.

In the other direction, given the observations above, it only remains to show that the right-to-left directions of **A3** and **A4** are theorems of the Yalcin logic, which by Boolean reasoning reduces to showing that the following four formulas are theorems of the Yalcin logic:

$$(\varphi \Rightarrow \alpha) \rightarrow (\varphi \Rightarrow (\alpha \vee \Box\beta)) \quad (6)$$

$$(\varphi \Rightarrow \beta) \rightarrow (\varphi \Rightarrow (\alpha \vee \Box\beta)) \quad (7)$$

$$(\varphi \Rightarrow \alpha) \rightarrow (\varphi \Rightarrow (\alpha \vee \Diamond\beta)) \quad (8)$$

$$\neg(\varphi \Rightarrow \neg\beta) \rightarrow (\varphi \Rightarrow (\alpha \vee \Diamond\beta)). \quad (9)$$

Formula (6) is an instance of **I3**. For (7), as instances of **I4** and **I3**, we have:

$$(\varphi \Rightarrow \beta) \rightarrow (\varphi \Rightarrow \Box\beta)$$

$$(\varphi \Rightarrow \Box\beta) \rightarrow (\varphi \Rightarrow (\Box\beta \vee \alpha)),$$

and we can use replacement of equivalents to replace $\Box\beta \vee \alpha$ by $\alpha \vee \Box\beta$. Next, (8) is an instance of **I3**.

Finally, for (9), as an instance of **I7**, we have

$$\neg(\varphi \Rightarrow \neg\beta) \rightarrow (\varphi \Rightarrow \Diamond\neg\neg\beta)$$

and hence

$$\neg(\varphi \Rightarrow \neg\beta) \rightarrow (\varphi \Rightarrow \Diamond\beta)$$

by replacement of equivalents; and as an instance of **I3**, we have

$$(\varphi \Rightarrow \Diamond\beta) \rightarrow (\varphi \Rightarrow (\Diamond\beta \vee \alpha))$$

and hence

$$(\varphi \Rightarrow \Diamond\beta) \rightarrow (\varphi \Rightarrow (\alpha \vee \Diamond\beta))$$

by replacement of equivalents. Putting the above implications together, we can derive (9). \square

For the next step in our argument, we use the following well-known fact about the modal logic K45, which is exactly the set of theorems of the Yalcin logic restricted to the language \mathcal{L} .

Lemma 1.6. Every formula $\varphi \in \mathcal{L}$ is provably equivalent in K45 to a formula of the form

$$\bigwedge_{1 \leq i \leq n} (\pi^i \vee \Diamond\beta^i \vee \Box\beta_1^i \vee \dots \vee \Box\beta_{m_i}^i),$$

where π^i , β^i , and $\beta_1^i, \dots, \beta_{m_i}^i$ are nonmodal formulas, and to a formula of the form

$$\bigvee_{1 \leq i \leq n} (\pi^i \wedge \Box\beta^i \wedge \Diamond\beta_1^i \wedge \dots \wedge \Diamond\beta_{m_i}^i),$$

where π^i , β^i , and $\beta_1^i, \dots, \beta_{m_i}^i$ are nonmodal formulas.

Now we take advantage of Lemmas 1.5 and 1.6 to show that using the Yalcin logic, every formula containing conditionals can be reduced to an equivalent formula without conditionals.

Lemma 1.7. Every formula of $\mathcal{L}(\Rightarrow)$ is provably equivalent in the Yalcin logic to a formula of \mathcal{L} .

Proof. It suffices to show that any conditional formula *containing no embedded conditionals* is equivalent to a formula of \mathcal{L} . For in any formula containing conditionals, there must be “innermost” occurrences of conditional formulas containing no embedded conditionals, and then we can replace those occurrences of conditional formulas—using replacement of equivalents—with equivalent formulas of \mathcal{L} , repeating this process until we eventually obtain a formula containing no conditionals. For a rigorous treatment of such a reduction argument, see [5, §7.4].

In what follows, we use the fact from Lemma 1.5 that **A1–A4** are derivable in the Yalcin logic.

Given a conditional formula $\varphi \Rightarrow \psi$ with no conditionals in φ or ψ , we use the fact from Lemma 1.6 that ψ is equivalent to a formula ψ' in K45 conjunctive normal form:

$$\psi' := \bigwedge_{1 \leq i \leq n} (\pi^i \vee \diamond \beta^i \vee \square \beta_1^i \vee \dots \vee \square \beta_{m_i}^i).$$

By replacement of equivalents, $\varphi \Rightarrow \psi$ is equivalent to $\varphi \Rightarrow \psi'$. By repeated application of **A2**, $\varphi \Rightarrow \psi'$ is equivalent to

$$\bigwedge_{1 \leq i \leq n} (\varphi \Rightarrow (\pi^i \vee \diamond \beta^i \vee \square \beta_1^i \vee \dots \vee \square \beta_{m_i}^i)).$$

So it suffices to show that each formula of the form

$$\varphi \Rightarrow (\pi \vee \diamond \beta \vee \square \beta_1 \vee \dots \vee \square \beta_m)$$

is equivalent to a formula of \mathcal{L} . Let

$$\alpha := \pi \vee \square \beta_1 \vee \dots \vee \square \beta_m.$$

By replacement of equivalents, we can equivalently write the above conditional as

$$\varphi \Rightarrow (\alpha \vee \diamond \beta),$$

which is equivalent, by **A4**, to

$$(\varphi \Rightarrow \alpha) \vee \neg(\varphi \Rightarrow \neg \beta),$$

which is in turn equivalent, by **A1**, to

$$(\varphi \Rightarrow \alpha) \wedge \neg \square(\varphi \rightarrow \neg \beta),$$

since β is nonmodal. Now it suffices to show that each formula of the form

$$\varphi \Rightarrow (\pi \vee \square \beta_1 \vee \dots \vee \square \beta_m)$$

is equivalent to a formula of \mathcal{L} . We do so by induction on m . Now let

$$\alpha := \pi \vee \square \beta_1 \vee \dots \vee \square \beta_{m-1},$$

so we can write the above conditional as

$$\varphi \Rightarrow (\alpha \vee \square \beta_m),$$

which is equivalent, by **A3**, to

$$(\varphi \Rightarrow \alpha) \wedge (\varphi \Rightarrow \beta_m),$$

which is in turn equivalent, by **A1**, to

$$(\varphi \Rightarrow \alpha) \wedge \Box(\varphi \rightarrow \beta_m),$$

because β_m is nonmodal. By the inductive hypothesis, the first disjunct is equivalent to a formula of \mathcal{L} . This completes the proof that our original formula $\varphi \Rightarrow \psi$ is equivalent to a formula of \mathcal{L} . \square

We can now put everything together to prove Theorem 1.1.

Proof. For soundness, if φ is derivable in the Yalcin logic, then by Lemma 1.5, φ is derivable in the logic defined with the axioms of Lemma 1.4, which we have shown to be valid. Since the rules also preserve validity, φ is valid. For completeness, if φ is valid, then so is its equivalent φ' in the modal language \mathcal{L} provided by Lemma 1.7, given soundness. By the completeness of K45 in Lemma 1.3, φ' is provable in K45 and hence in the Yalcin logic, and by Lemma 1.7, the Yalcin logic proves $\varphi \leftrightarrow \varphi'$. Thus, the Yalcin logic proves φ . Finally, it is easy to see that φ is an informational consequence of $\{\sigma_1, \dots, \sigma_n\}$ iff $(\Box\sigma_1 \wedge \dots \wedge \Box\sigma_n) \rightarrow \Box\varphi$ is valid according to Yalcin's semantics and hence iff $(\Box\sigma_1 \wedge \dots \wedge \Box\sigma_n) \rightarrow \Box\varphi$ is a theorem of the Yalcin logic by our previous reasoning. \square

2 From Formal Semantics to DEL

In this section, we consider an alternative semantics for the indicative conditional, suggested in the formal semantics literature. This semantics is also of interest purely from the perspective of dynamic epistemic logic, but it has not been previously considered in the DEL literature.

The semantics for the indicative \Rightarrow in §1 was intended for the case where the antecedent is *nonmodal*. If the antecedent is modal, things become trickier, for the well-known reason in dynamic epistemic logic [5, 6] that updating with the antecedent φ may fail to bring about an information state that accepts φ . The desire is to update the information state in such a way that the antecedent φ is accepted and then check whether the consequent ψ is accepted. However, we cannot say that the information state X should be updated to the largest $X' \subseteq X$ that accepts the antecedent φ , because even if there is such a subset, there might fail to be a unique *largest* one. A solution from Kolodny and MacFarlane [9] is the following:

- $\mathcal{M}, w, X \models \varphi \Rightarrow \psi$ iff $\mathcal{M}, w, X' \models \Box\psi$ for all X' such that
 (i) $X' \subseteq X$, (ii) $X' \subseteq \llbracket \varphi \rrbracket^{\mathcal{M}, X'}$, and (iii) there is no X'' satisfying (i) and (ii) such that $X' \subsetneq X''$.

Although this semantics is equivalent to that of §1 for nonmodal φ , they are not equivalent in general.

Example 2.1. Consider a two-world model \mathcal{M} with w and v such that p is true only at w and q is true only at v . According to the semantics of §1, the formula $(\Box p \vee \Box \neg p) \Rightarrow q$ is trivially true at $\mathcal{M}, w, \{w, v\}$, because $\llbracket \Box p \vee \Box \neg p \rrbracket^{\mathcal{M}, \{w, v\}} = \emptyset$. By contrast, according to the Kolodny-MacFarlane semantics, the formula $(\Box p \vee \Box \neg p) \Rightarrow q$ is false at $\mathcal{M}, w, \{w, v\}$, because $\{w\}$ is an $X' \subseteq \{w, v\}$ satisfying conditions (i), (ii), and (iii), and yet $\mathcal{M}, w, \{w\} \not\models \Box q$.

Example 2.2. According to the Kolodny-MacFarlane semantics, $(p \wedge \Diamond \neg p) \Rightarrow \perp$ is valid, because there can be no nonempty X' such that $X' \subseteq \llbracket p \wedge \Diamond \neg p \rrbracket^{\mathcal{M}, X'}$. By contrast, according to the semantics of §1, $(p \wedge \Diamond \neg p) \Rightarrow \perp$ is invalid, because on that semantics $(p \wedge \Diamond \neg p) \Rightarrow \varphi$ is equivalent to $p \Rightarrow \varphi$ for any φ .

From the point of view of dynamic epistemic logic, the above semantics can be seen as interpreting $\varphi \Rightarrow \psi$ as a very natural statement: “every minimal epistemic change yielding knowledge of φ also yields knowledge of ψ .” As Example 2.1 shows, there may be more than one minimal epistemic change yielding knowledge of φ , and as Example 2.2 shows, there may be no epistemic change yielding knowledge of φ . The standard interpretation of $[\varphi]\Box\psi$ in dynamic epistemic logic, as “information update with φ yields knowledge of ψ ,” is quite different, due to the phenomenon of *unsuccessful update* alluded to above, wherein updating with φ fails to produce an epistemic state in which φ is known. The difference between $[\varphi]\Box\psi$ under the standard semantics and $\varphi \Rightarrow \psi$ under the Kolodny-MacFarlane semantics is that in the first case, φ specifies the *mechanism of epistemic change*—remove from the information state all worlds that satisfied $\neg\varphi$ —whereas in the second case, φ specifies the *desired result of epistemic change*—an epistemic state in which φ is known. We think that both ways of reasoning are important for dynamic epistemic logicians interested in information update and learning. There are also natural generalizations of the Kolodny-MacFarlane semantics to the multi-agent setting in which dynamic epistemic logicians typically work (e.g., every minimal epistemic change yielding common knowledge of φ also yields common knowledge of ψ), but in this preliminary analysis we restrict attention to the single-agent case.

Among the interesting application of the Kolodny-MacFarlane conditional is to succinctly express *dependence* as in modal dependence logic [13] or *supervenience* as in philosophy [11].

Definition 2.3. In an information state X , the (truth value of the) propositional variable q *depends on* or *supervenes on* the (truth values of the) propositional variables p_1, \dots, p_n iff any two worlds in X that agree on the truth values of p_1, \dots, p_n also agree on the truth value of q .

The dependence of q on p_1, \dots, p_n can be expressed in the language \mathcal{L} by the formula

$$\bigwedge_{s \in \text{state}(p_1, \dots, p_n)} (\Box(s \rightarrow q) \vee \Box(s \rightarrow \neg q)), \quad (10)$$

where $\text{state}(p_1, \dots, p_n)$ is the set of all conjunctions of the form $\pm_1 p_1 \wedge \dots \wedge \pm_n p_n$ with \pm_i being either \neg or empty. But whereas the formula (10) grows exponentially in the size of n , the Kolodny-MacFarlane conditional allows us to express dependence with a formula that grows only linearly in the size of n .

Proposition 2.4. In an information state X , the propositional variable q depends on p_1, \dots, p_n iff the following formula is true relative to X according to the Kolodny-MacFarlane semantics for \Rightarrow :

$$\left(\bigwedge_{1 \leq i \leq n} (\Box p_i \vee \Box \neg p_i) \right) \Rightarrow (\Box q \vee \Box \neg q). \quad (11)$$

Proof. It is easy to see that if q depends on p_1, \dots, p_n in X , then (11) is true relative to X . Conversely, suppose (11) is true relative to X , and x and y are worlds in X that agree on the truth values of p_1, \dots, p_n . Then we claim that the set X' of all worlds in X that agree with x on the truth values of p_1, \dots, p_n meets conditions (i), (ii), and (iii) above. For (ii), by the definition of X' , we have $\mathcal{M}, w, X' \models \Box p_i \vee \Box \neg p_i$ for each i . For (iii), note that if $X' \subsetneq X'' \subseteq X$, then there is a $z \in X'' \setminus X'$ that disagrees with x on the truth value of some p_i , which with $x, z \in X''$ implies that $\mathcal{M}, w, X'' \not\models \Box p_i \vee \Box \neg p_i$, so X'' does not satisfy (ii). Now since X' satisfies (i), (ii), and (iii), and (11) is true relative to X , it follows that $\mathcal{M}, w, X' \models \Box q \vee \Box \neg q$, which implies that all worlds in X' , and in particular x and y , agree on the truth value of q . \square

Below we will provide a computable translation $(\cdot)^\dagger$ from the language $\mathcal{L}(\Rightarrow)$ into the basic modal language \mathcal{L} . As shown in Lemma 2.6, this translation preserves (in)validity according to Kolodny and MacFarlane’s semantics. By Lemma 1.3 and the fact that we have not changed the semantics from §1 for formulas without conditionals, this provides a full and faithful translation from the logic with epistemic

modals and Kolodny and MacFarlane's indicative conditional to the logic K45. As is well-known, K45 is decidable, so it follows that Kolodny and MacFarlane's logic is also decidable.

The strategy for our translation is as follows. We first define the translation λ^* for a conditional formula λ with no embedded conditionals, using K45 normal forms (Lemma 1.6). We then extend this to a translation from the full language by induction.

Suppose we are given a conditional formula

$$\lambda := \Theta \Rightarrow \Omega$$

where Θ and Ω are in K45 *disjunctive* normal form, so that

$$\Theta := \bigvee_{i \in I} \theta_i \quad \text{and} \quad \Omega := \bigvee_{j \in J} \omega_j,$$

where for nonmodal formulas φ_i , ψ_i , χ_n , α_j , β_j , and γ_m ,

$$\theta_i := \varphi_i \wedge \Box \psi_i \wedge \bigwedge_{n \in D_i} \Diamond \chi_n \quad \text{and} \quad \omega_j := \alpha_j \wedge \Box \beta_j \wedge \bigwedge_{m \in D_j} \Diamond \gamma_m.$$

We would like our translation λ^* to express in the basic modal language that every maximal set of worlds making some of the θ_i formulas true also makes at least one of the ω_j formulas true. Where K indexes some subset of the θ_i formulas, info_K below gives the nonmodal information that each world in an information state must satisfy in order for the information state to accept $\bigvee_{i \in K} \theta_i$, while good_K asserts that $\bigvee_{i \in K} \theta_i$ will indeed be accepted when we restrict the current information state to the set of worlds satisfying info_K , thanks to sufficient witnesses for the $\Diamond \chi_n$ formulas. Meanwhile, max_K adds that K is maximal, with respect to the set of worlds satisfying info_K , among such subsets. For $K \subseteq I$, let

$$\text{info}_K := \left(\bigvee_{k \in K} \varphi_k \right) \wedge \bigwedge_{k \in K} \psi_k,$$

let

$$\text{good}_K := \bigwedge_{k \in K} \bigwedge_{n \in D_k} \Diamond (\text{info}_K \wedge \chi_n),$$

and let

$$\text{max}_K := \text{good}_K \wedge \bigwedge_{L \subseteq I} \left(\left(\Box (\text{info}_K \rightarrow \text{info}_L) \wedge \Diamond (\neg \text{info}_K \wedge \text{info}_L) \right) \rightarrow \neg \text{good}_L \right).$$

Concerning the consequent formula Ω : for $S \subseteq J$, let

$$\text{states}_S := \bigwedge_{s \in S} \alpha_s \wedge \bigwedge_{s \in J \setminus S} \neg \alpha_s.$$

Finally, given our λ above, we define

$$\lambda^* := \bigwedge_{K \subseteq I} \left(\text{max}_K \rightarrow \Box \left(\text{info}_K \rightarrow \bigwedge_{S \subseteq J} \left(\text{states}_S \rightarrow \bigvee_{s \in S} \left(\Box (\text{info}_K \rightarrow \beta_s) \wedge \bigwedge_{m \in D_s} \Diamond (\text{info}_K \wedge \gamma_m) \right) \right) \right) \right).$$

As we show in Lemma 2.6.3 below, if max_K is true, then (the truth set of) info_K picks out a maximal Θ -accepting subset of the current information state. The translation thus guarantees that if we restrict

the current information state to such a subset, then at least one of the disjuncts of Ω will be true at each world in the restricted information state. In particular, for at least one such formula ω_s we have that β_s is true throughout the subset, so $\Box\beta_s$ is true, and each of the $\Diamond\gamma_m$ conjuncts of ω_s is witnessed by some world making γ_m true. We also need to know that these formulas info_K for which max_K holds pick out all of the maximal Θ -accepting subsets, which is the content of Lemma 2.6.2.

We regard $(\cdot)^*$ as a partial function from $\mathcal{L}(\Rightarrow)$ to \mathcal{L} such that φ^* is defined iff φ is of the form $\Theta \Rightarrow \Omega$ as above. We then define a partial function $(\cdot)^\dagger$ from $\mathcal{L}(\Rightarrow)$ to \mathcal{L} as follows:

- $p^\dagger = p$; $(\neg\varphi)^\dagger = \neg\varphi^\dagger$; $(\varphi \wedge \psi)^\dagger = \varphi^\dagger \wedge \psi^\dagger$; $(\Box\varphi)^\dagger = \Box\varphi^\dagger$;
- $(\varphi \Rightarrow \psi)^\dagger = \begin{cases} ((\varphi^\dagger)^{NF} \Rightarrow (\psi^\dagger)^{NF})^* & \text{if this is defined} \\ \text{undefined} & \text{otherwise} \end{cases}$,

where $(\chi)^{NF}$ is the K45 disjunctive normal form of χ . An easy induction shows that $(\cdot)^\dagger$ is in fact a total function, so

- $(\varphi \Rightarrow \psi)^\dagger = ((\varphi^\dagger)^{NF} \Rightarrow (\psi^\dagger)^{NF})^*$,

and $\varphi^\dagger \in \mathcal{L}$ for every $\varphi \in \mathcal{L}(\Rightarrow)$.

Theorem 2.5 (Reduction of $\mathcal{L}(\Rightarrow)$ to \mathcal{L}). For every $\delta \in \mathcal{L}(\Rightarrow)$:

1. $\delta^\dagger \in \mathcal{L}$;
2. for every pointed model \mathcal{M}, w, X : $\mathcal{M}, w, X \models \delta$ iff $\mathcal{M}, w, X \models \delta^\dagger$.

Proof. The proof is by induction on δ . The only nontrivial case is where δ is of the form $\varphi \Rightarrow \psi$, so we must show $\mathcal{M}, w, X \models \varphi \Rightarrow \psi$ iff $\mathcal{M}, w, X \models ((\varphi^\dagger)^{NF} \Rightarrow (\psi^\dagger)^{NF})^*$. By the inductive hypothesis, $\llbracket \varphi \rrbracket^{\mathcal{M}, X} = \llbracket \varphi^\dagger \rrbracket^{\mathcal{M}, X}$, and by Lemma 1.6, since $\varphi^\dagger \in \mathcal{L}$, $\llbracket \varphi^\dagger \rrbracket^{\mathcal{M}, X} = \llbracket (\varphi^\dagger)^{NF} \rrbracket^{\mathcal{M}, X}$. Thus, $\llbracket \varphi \rrbracket^{\mathcal{M}, X} = \llbracket (\varphi^\dagger)^{NF} \rrbracket^{\mathcal{M}, X}$, and similarly $\llbracket \psi \rrbracket^{\mathcal{M}, X} = \llbracket (\psi^\dagger)^{NF} \rrbracket^{\mathcal{M}, X}$. It follows that $\mathcal{M}, w, X \models \varphi \Rightarrow \psi$ iff $\mathcal{M}, w, X \models (\varphi^\dagger)^{NF} \Rightarrow (\psi^\dagger)^{NF}$. To complete the proof, it only remains to show that for any $\Theta \Rightarrow \Omega$ where $\Theta, \Omega \in \mathcal{L}$ are in normal form, we have $\mathcal{M}, w, X \models \Theta \Rightarrow \Omega$ iff $\mathcal{M}, w, X \models (\Theta \Rightarrow \Omega)^*$, as in Proposition 2.7 below. \square

As the proof of Theorem 2.5 shows, the key task is now to prove Proposition 2.7, for which we need a preliminary lemma. From now on, we say that $Y \subseteq X$ is a Θ -subset of X if $Y \subseteq \llbracket \Theta \rrbracket^{\mathcal{M}, Y}$.

Lemma 2.6. For any pointed model \mathcal{M}, w, X :

1. if $\mathcal{M}, w, X \models \text{good}_L$, then $\llbracket \text{info}_L \rrbracket^{\mathcal{M}, X}$ is a Θ -subset of X ;
2. if Y is a maximal Θ -subset of X , there is a $K \subseteq I$ such that $Y = \llbracket \text{info}_K \rrbracket^{\mathcal{M}, X}$ and $\mathcal{M}, w, X \models \text{max}_K$;
3. if $\mathcal{M}, w, X \models \text{max}_L$, then $\llbracket \text{info}_L \rrbracket^{\mathcal{M}, X}$ is a maximal Θ -subset of X .

Proof. For part 1, suppose $\mathcal{M}, w, X \models \text{good}_L$. We must show that $\llbracket \text{info}_L \rrbracket^{\mathcal{M}, X} \subseteq \llbracket \Theta \rrbracket^{\mathcal{M}, \llbracket \text{info}_L \rrbracket^{\mathcal{M}, X}}$. So suppose $v \in \llbracket \text{info}_L \rrbracket^{\mathcal{M}, X}$. Then there is some $p \in L$ such that $v \in \llbracket \varphi_p \rrbracket^{\mathcal{M}, X}$. We claim that $v \in \llbracket \theta_p \rrbracket^{\mathcal{M}, \llbracket \text{info}_L \rrbracket^{\mathcal{M}, X}}$. By definition of info_L , we have $\llbracket \text{info}_L \rrbracket^{\mathcal{M}, X} \subseteq \llbracket \bigwedge_{l \in L} \psi_l \rrbracket^{\mathcal{M}, X}$ and hence $\llbracket \text{info}_L \rrbracket^{\mathcal{M}, X} \subseteq \llbracket \bigwedge_{l \in L} \psi_l \rrbracket^{\mathcal{M}, \llbracket \text{info}_L \rrbracket^{\mathcal{M}, X}}$ since the ψ_l 's are nonmodal, so

$$\llbracket \text{info}_L \rrbracket^{\mathcal{M}, X} \subseteq \llbracket \bigwedge_{l \in L} \Box \psi_l \rrbracket^{\mathcal{M}, \llbracket \text{info}_L \rrbracket^{\mathcal{M}, X}}. \quad (12)$$

Then since $\mathcal{M}, w, X \models \text{good}_L$, we also have

$$\mathcal{M}, w, X \models \bigwedge_{n \in D_p} \diamond(\text{info}_L \wedge \chi_n),$$

which with the fact that the χ_n 's are nonmodal implies

$$[\text{info}_L]^{\mathcal{M}, X} \subseteq [\bigwedge_{n \in D_p} \diamond \chi_n]^{\mathcal{M}, [\text{info}_L]^{\mathcal{M}, X}}. \quad (13)$$

By $v \in [\varphi_p]^{\mathcal{M}, X}$, (12), and (13), we have $v \in [\theta_p]^{\mathcal{M}, [\text{info}_L]^{\mathcal{M}, X}}$. Hence $[\text{info}_L]^{\mathcal{M}, X} \subseteq [\Theta]^{\mathcal{M}, [\text{info}_L]^{\mathcal{M}, X}}$.

For part 2, suppose Y is a maximal Θ -subset of X , so $Y \subseteq [\Theta]^{\mathcal{M}, Y}$ and there is no Z such that $Y \subsetneq Z \subseteq X$ and $Z \subseteq [\Theta]^{\mathcal{M}, Z}$. Let

$$K = \{k \in I \mid [\theta_k]^{\mathcal{M}, Y} \neq \emptyset\}. \quad (14)$$

We will show that $Y \subseteq [\text{info}_K]^{\mathcal{M}, X}$. First, observe that for each $k \in K$, since $[\theta_k]^{\mathcal{M}, Y} \neq \emptyset$, we have $Y \subseteq [\psi_k]^{\mathcal{M}, X}$, so $Y \subseteq [\bigwedge_{k \in K} \psi_k]^{\mathcal{M}, X}$. Moreover, since $Y \subseteq [\Theta]^{\mathcal{M}, Y}$, for every $y \in Y$, there is a $k_y \in K$ with $y \in [\theta_{k_y}]^{\mathcal{M}, Y}$, so $y \in [\varphi_{k_y}]^{\mathcal{M}, X}$. Thus, for every $y \in Y$, $y \in [\varphi_{k_y} \wedge \bigwedge_{k \in K} \psi_k]^{\mathcal{M}, X}$. Hence $Y \subseteq [\text{info}_K]^{\mathcal{M}, X}$.

Next we show that $Y \supseteq [\text{info}_K]^{\mathcal{M}, X}$. Suppose not, so there is a $w \in [\text{info}_K]^{\mathcal{M}, X}$ such that $w \notin Y$. Since $w \in [\text{info}_K]^{\mathcal{M}, X}$, we have $w \in [\bigwedge_{k \in K} \psi_k]^{\mathcal{M}, X}$. Now we claim that $Y \cup \{w\} \subseteq [\Theta]^{\mathcal{M}, Y \cup \{w\}}$. Consider a $y \in Y$ and a disjunct θ_r of Θ such that $y \in [\theta_r]^{\mathcal{M}, Y}$, which exists since $Y \subseteq [\Theta]^{\mathcal{M}, Y}$. We claim that $y \in [\theta_r]^{\mathcal{M}, Y \cup \{w\}}$. First, clearly $y \in [\varphi_r \wedge \bigwedge_{n \in D_r} \diamond \chi_n]^{\mathcal{M}, Y}$ implies $y \in [\varphi_r \wedge \bigwedge_{n \in D_r} \diamond \chi_n]^{\mathcal{M}, Y \cup \{w\}}$. Then since $w \in [\bigwedge_{k \in K} \psi_k]^{\mathcal{M}, X}$, $y \in [\square \psi_r]^{\mathcal{M}, Y}$ implies $y \in [\square \psi_r]^{\mathcal{M}, Y \cup \{w\}}$. Thus, $y \in [\theta_r]^{\mathcal{M}, Y \cup \{w\}}$ and hence $y \in [\Theta]^{\mathcal{M}, Y \cup \{w\}}$, so we have shown that $Y \subseteq [\Theta]^{\mathcal{M}, Y \cup \{w\}}$. It only remains to observe that $w \in [\Theta]^{\mathcal{M}, Y \cup \{w\}}$. Since $w \in [\text{info}_K]^{\mathcal{M}, X}$, there is a $t \in K$ such that $w \in [\varphi_t]^{\mathcal{M}, X}$. It is then easy to see that $w \in [\theta_t]^{\mathcal{M}, Y \cup \{w\}}$. Hence $Y \cup \{w\} \subseteq [\Theta]^{\mathcal{M}, Y \cup \{w\}}$, contradicting the assumption that Y is a maximal Θ -subset of X .

Thus, we have shown that $Y = [\text{info}_K]^{\mathcal{M}, X}$. Finally, we must show that $\mathcal{M}, w, X \models \text{max}_K$. To show that $\mathcal{M}, w, X \models \text{good}_K$, it follows from (14) that for every $k \in K$ and $n \in D_k$, $Y \cap [\gamma_n]^{\mathcal{M}, X} \neq \emptyset$, which with $Y = [\text{info}_K]^{\mathcal{M}, X}$ gives us $[\text{info}_K]^{\mathcal{M}, X} \cap [\gamma_n]^{\mathcal{M}, X} \neq \emptyset$, so $\mathcal{M}, w, X \models \text{good}_K$. Finally, suppose $\mathcal{M}, w, X \models \square(\text{info}_K \rightarrow \text{info}_L) \wedge \diamond(\neg \text{info}_K \wedge \text{info}_L)$, so $[\text{info}_L]^{\mathcal{M}, X} \supsetneq [\text{info}_K]^{\mathcal{M}, X}$. Then since $Y = [\text{info}_K]^{\mathcal{M}, X}$ is a maximal Θ -subset of X , $[\text{info}_L]^{\mathcal{M}, X}$ is not a Θ -subset of X . Hence by part 1, we have $\mathcal{M}, w, X \models \neg \text{good}_L$. This completes the proof that $\mathcal{M}, w, X \models \text{max}_K$.

For part 3, if $\mathcal{M}, w, X \models \text{max}_L$, then $\mathcal{M}, w, X \models \text{good}_L$, so by part 1, $[\text{info}_L]^{\mathcal{M}, X}$ is a Θ -subset of X . Now suppose $[\text{info}_L]^{\mathcal{M}, X}$ is not a maximal Θ -subset, so there is a Θ -subset $Y \supsetneq [\text{info}_L]^{\mathcal{M}, X}$. Then by part 2, there is a $K \subseteq I$ such that $Y = [\text{info}_K]^{\mathcal{M}, X}$ and $\mathcal{M}, w, X \models \text{max}_K$, so $\mathcal{M}, w, X \models \text{good}_K$. Since $Y \supsetneq [\text{info}_L]^{\mathcal{M}, X}$ and $Y = [\text{info}_K]^{\mathcal{M}, X}$, we have $[\text{info}_K]^{\mathcal{M}, X} \supsetneq [\text{info}_L]^{\mathcal{M}, X}$, so $\mathcal{M}, w, X \models \square(\text{info}_L \rightarrow \text{info}_K) \wedge \diamond(\neg \text{info}_L \wedge \text{info}_K)$. Then since $\mathcal{M}, w, X \models \text{max}_L$, it follows that $\mathcal{M}, w, X \models \neg \text{good}_K$. From this contradiction we conclude that $[\text{info}_L]^{\mathcal{M}, X}$ is a maximal Θ -subset of X . \square

We are now ready to establish the key proposition used in the proof of Theorem 2.5, namely the semantic equivalence of a conditional λ (without embedded conditionals) and its translation λ^* .

Proposition 2.7. For any pointed model \mathcal{M}, w, X :

$$\mathcal{M}, w, X \models \lambda \text{ iff } \mathcal{M}, w, X \models \lambda^*.$$

Proof. Since $\lambda := \Theta \Rightarrow \Omega$, we have $\mathcal{M}, w, X \models \lambda$ iff for all maximal Θ -subsets Y of X , $Y \subseteq \llbracket \Omega \rrbracket^{\mathcal{M}, Y}$. By Lemma 2.6, Y is a maximal Θ -subset of X iff there is a $K \subseteq I$ such that $Y = \llbracket \text{info}_K \rrbracket^{\mathcal{M}, X}$ and $\mathcal{M}, w, X \models \max_K$. Thus, the condition that $\mathcal{M}, w, X \models \lambda^*$, i.e.,

$$\mathcal{M}, w, X \models \bigwedge_{K \subseteq I} \left(\max_K \rightarrow \square \left(\text{info}_K \rightarrow \bigwedge_{S \subseteq J} \left(\text{states}_S \rightarrow \bigvee_{s \in S} \left(\square(\text{info}_K \rightarrow \beta_s) \wedge \bigwedge_{m \in D_s} \diamond(\text{info}_K \wedge \gamma_m) \right) \right) \right) \right)$$

is equivalent to:

$$\text{for all maximal } \Theta\text{-subsets } Y \text{ of } X, \text{ there is a } K \subseteq I \text{ such that } Y = \llbracket \text{info}_K \rrbracket^{\mathcal{M}, X} \text{ and} \\ \mathcal{M}, w, X \models \square \left(\text{info}_K \rightarrow \bigwedge_{S \subseteq J} \left(\text{states}_S \rightarrow \bigvee_{s \in S} \left(\square(\text{info}_K \rightarrow \beta_s) \wedge \bigwedge_{m \in D_s} \diamond(\text{info}_K \wedge \gamma_m) \right) \right) \right). \quad (15)$$

Below we will show that if $Y = \llbracket \text{info}_K \rrbracket^{\mathcal{M}, X}$, then

$$\mathcal{M}, w, X \models \square \left(\text{info}_K \rightarrow \bigwedge_{S \subseteq J} \left(\text{states}_S \rightarrow \bigvee_{s \in S} \left(\square(\text{info}_K \rightarrow \beta_s) \wedge \bigwedge_{m \in D_s} \diamond(\text{info}_K \wedge \gamma_m) \right) \right) \right) \quad (16)$$

iff $Y \subseteq \llbracket \Omega \rrbracket^{\mathcal{M}, Y}$. Thus, (15) is equivalent to: for all maximal Θ -subsets Y of X , $Y \subseteq \llbracket \Omega \rrbracket^{\mathcal{M}, Y}$. Given the other equivalences above, this establishes $\mathcal{M}, w, X \models \lambda^*$ iff $\mathcal{M}, w, X \models \lambda$.

Now suppose $Y = \llbracket \text{info}_K \rrbracket^{\mathcal{M}, X}$. Then (16) is equivalent

$$Y \subseteq \left[\bigwedge_{S \subseteq J} \left(\text{states}_S \rightarrow \bigvee_{s \in S} \left(\square(\text{info}_K \rightarrow \beta_s) \wedge \bigwedge_{m \in D_s} \diamond(\text{info}_K \wedge \gamma_m) \right) \right) \right]^{\mathcal{M}, X}.$$

Thus, to show $Y \subseteq \llbracket \Omega \rrbracket^{\mathcal{M}, Y}$, it suffices to show

$$\left[\bigwedge_{S \subseteq J} \left(\text{states}_S \rightarrow \bigvee_{s \in S} \left(\square(\text{info}_K \rightarrow \beta_s) \wedge \bigwedge_{m \in D_s} \diamond(\text{info}_K \wedge \gamma_m) \right) \right) \right]^{\mathcal{M}, X} \subseteq \llbracket \Omega \rrbracket^{\mathcal{M}, Y}. \quad (17)$$

So suppose v is in the left-hand side of (17). By definition of states_S , there is exactly one $S \subseteq J$ such that $v \in \llbracket \text{states}_S \rrbracket^{\mathcal{M}, X}$. Moreover, $S \neq \emptyset$, for if $S = \emptyset$, then the empty disjunction in (17) is \perp , so we would have $v \in \llbracket \text{states}_S \rrbracket^{\mathcal{M}, X}$ and $v \in \llbracket \text{states}_S \rightarrow \perp \rrbracket^{\mathcal{M}, X} = \llbracket \neg \text{states}_S \rrbracket^{\mathcal{M}, X}$, i.e., $v \notin \llbracket \text{states}_S \rrbracket^{\mathcal{M}, X}$, a contradiction. Since $S \neq \emptyset$, there is an $s \in S$ such that

$$v \in \llbracket \alpha_s \wedge \square(\text{info}_K \rightarrow \beta_s) \wedge \bigwedge_{m \in D_s} \diamond(\text{info}_K \wedge \gamma_m) \rrbracket^{\mathcal{M}, X} \quad (18)$$

which with $Y = \llbracket \text{info}_K \rrbracket^{\mathcal{M}, X}$ implies

$$v \in \llbracket \alpha_s \wedge \square \beta_s \wedge \bigwedge_{m \in D_s} \diamond \gamma_m \rrbracket^{\mathcal{M}, Y}, \quad (19)$$

so $v \in \llbracket \omega_s \rrbracket^{\mathcal{M}, Y} \subseteq \llbracket \Omega \rrbracket^{\mathcal{M}, Y}$. This establishes (17), which completes the proof. \square

Although we have now shown that any formula of $\mathcal{L}(\Rightarrow)$ can be effectively translated into a modal formula that is semantically equivalent according to Kolodny and MacFarlane's semantics, there is clearly a huge blowup in formula size. It is reasonable to conjecture that according to this semantics, $\mathcal{L}(\Rightarrow)$ is *exponentially more succinct* than the basic modal language \mathcal{L} in the sense of [10].

3 Conclusion

In §1, we presented an example of how techniques from dynamic epistemic logic can be fruitfully applied to the formal semantics of modals and conditionals, by providing a complete axiomatization of the inferences validated by a formal semantics. For arguments that knowing such a complete axiomatization is of value for formal semantics, see [8]. In the other direction, in §2, we presented an example of how ideas from the formal semantics of modals and conditionals can be profitably imported into dynamic epistemic logic, by enabling a natural kind of reasoning about epistemic change—focusing not on the mechanism of epistemic change but rather on the desired result of epistemic change. We hope that these examples might provide some stimulus for further cross-pollination between these two overlapping fields.

References

- [1] Johan van Benthem (2008): *Logical Dynamics Meets Logical Pluralism?* *Australasian Journal of Logic* 6, pp. 182–209.
- [2] Johan van Benthem (2011): *Logical Dynamics of Information and Interaction*. Cambridge University Press, doi:10.1017/CB09780511974533.
- [3] Patrick Blackburn, Maarten de Rijke & Yde Venema (2001): *Modal Logic*. Cambridge University Press, doi:10.1017/CB09781107050884.
- [4] Justin Bledin (2014): *Logic Informed*. *Mind* 123(490), pp. 277–316, doi:10.1093/mind/fzu073.
- [5] Hans van Ditmarsch, Wiebe van der Hoek & Barteld Kooi (2008): *Dynamic Epistemic Logic*. Springer, doi:10.1007/978-1-4020-5839-4.
- [6] Wesley H. Holliday & Thomas F. Icard, III (2010): *Moorean Phenomena in Epistemic Logic*. In L. Beklemishev, V. Goranko & V. Shehtman, editors: *Advances in Modal Logic*, 8, College Publications, pp. 178–199.
- [7] Wesley H. Holliday & Thomas F. Icard, III (2013): *Measure semantics and qualitative semantics for epistemic modals*. In T. Snider, editor: *Proceedings of SALT 23*, LSA and CLC Publications, pp. 514–534, doi:10.3765/salt.v23i0.2670.
- [8] Wesley H. Holliday & Thomas F. Icard, III (Forthcoming): *Axiomatization in the Meaning Sciences*. In D. Ball & B. Rabern, editors: *The Science of Meaning*, Oxford University Press.
- [9] Niko Kolodny & John MacFarlane (2010): *Ifs and Oughts*. *Journal of Philosophy* 107, pp. 115–143, doi:10.5840/jphil2010107310.
- [10] Carsten Lutz (2006): *Complexity and succinctness of public announcement logic*. In P. Stone & G. Weiss, editors: *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 06)*, ACM, pp. 137–143, doi:10.1145/1160633.1160657.
- [11] Brian McLaughlin & Karen Bennett (2014): *Supervenience*. In E. N. Zalta, editor: *The Stanford Encyclopedia of Philosophy*, spring 2014 edition, Metaphysics Research Lab, Stanford University.
- [12] Moritz Schulz (2010): *Epistemic modals and informational consequence*. *Synthese* 174(3), pp. 385–395, doi:10.1007/s11229-009-9461-8.
- [13] Jouko Väänänen (2008): *Modal Dependence Logic*. In K. R. Apt & R. van Rooij, editors: *New Perspectives on Games and Interaction, Texts in Logic and Games* 4, Amsterdam University Press, pp. 237–254, doi:10.5117/9789089640574.
- [14] Malte Willer (2012): *A Remark on Iffy Oughts*. *The Journal of Philosophy* 109(7), pp. 449–461, doi:10.5840/jphil2012109719.
- [15] Seth Yalcin (2007): *Epistemic Modals*. *Mind* 116(464), pp. 983–1026, doi:10.1093/mind/fzm983.
- [16] Seth Yalcin (2012): *A Counterexample to Modus Tollens*. *Journal of Philosophical Logic* 41(6), pp. 1001–1024, doi:10.1007/s10992-012-9228-4.