Dynamic Epistemic Temporal Logic[*]

Bryan Renne[†]   Joshua Sack[‡§]   Audrey Yap[¶]

Abstract

We introduce a new type of arrow in the *update frames* (or "action models") of *Dynamic Epistemic Logic* in a way that enables us to reason about epistemic temporal dynamics in multi-agent systems that *need not be synchronous*. Since van Benthem and Pacuit (later joined by Hoshi and Gerbrandy) showed that standard Dynamic Epistemic Logic necessarily satisfies *synchronicity*, it follows that our arrow type is a new way of extending the domain of applicability of the Dynamic Epistemic Logic approach. Furthermore, our framework provides a new perspective on the van Benthem et al work itself. In particular, while each of our work and their work shows that epistemic temporal models generated by standard update frames necessarily satisfy certain structural properties such as synchronicity, our work clarifies the way in which these structural properties arise as a result of the inherent structure of standard update frames themselves.

1   Introduction

Dynamic Epistemic Logic [1, 2, 3, 4, 8, 12] is a modal-logic approach to reasoning about belief dynamics in multi-agent systems. The characteristic feature of this approach is its use of *update modals*, which are modal operators $[U, s]$ that describe operations on Kripke models. These operations, called *updates*, represent *informational events* in which the agents receive information that may bring about changes in their beliefs. The basic idea is that an update modal $[U, s]$ describes a specific partial function $f_{[U,s]}$ that maps a pointed Kripke model (M, w) in the domain of $f_{[U,s]}$ to another pointed Kripke model that we write as $\big(M[U], (w, s)\big)$. This allows us to view a sequence

$$(M_0, w_0), (M_1, w_1), (M_2, w_2), \ldots, (M_n, w_n) \tag{1}$$

of pointed Kripke models, with (M_{i+1}, w_{i+1}) generated from (M_i, w_i) by the update $f_{[U_{i+1}, s_{i+1}]}$ described by update modal $[U_{i+1}, s_{i+1}]$, as a discrete-time distributed multi-agent system in which the state of the system at time i is described by (M_i, w_i). Defining the *time of a world w in M_i*

within the sequence (1) to be the index i, we obtain a notion of time that is *external* to the pointed Kripke model (M_i, w_i). One consequence of adopting this external notion of time is that all of the worlds that an agent considers possible relative to a world w in M_i have time i. This implies that at every world, every agent knows the current time. Systems in which the current time is known at every world are called *synchronous* [5, 6]. Dynamic Epistemic Logic, which itself adopts this external notion of time, is consequently restricted to the study of *synchronous* multi-agent systems [5, 6].

In this paper, we propose a simple extension to the update modals $[U, s]$ that allows us to reason about discrete-time distributed multi-agent systems that *need not be synchronous*. We achieve this by adapting the methodology of standard Dynamic Epistemic Logic so that it fits naturally within a version of *Epistemic Temporal Logic* [9, 11] whose only temporal modality is a discrete one-step–past operator; this version will be called *Simple Epistemic Temporal Logic*. Simple Epistemic Temporal Logic uses *epistemic temporal models*, which are Kripke models in which one of the relational components is designated as a *time-keeping relation*. When w is related to w' according to the time-keeping relation, the intended interpretation is that w' is a possible way the system might have been one time-step before w. This provides us with an *internal* notion of time, in that the *time of a world w in an epistemic temporal model M* is determined solely based on the time-keeping relation, which is internal to the model M. Diagrammatically, we will represent this relation using arrows labeled by the symbol Y—called Y-*arrows*—where "Y" is a mnemonic for "yesterday" (so having a Y-arrow from world w to world w' is to be thought of as saying that w' is one of the possible ways w might have been "yesterday," meaning one time-step ago). In order to distinguish between Kripke models with and without a Y-*relation* (the time-keeping relation), we adopt the following terminology: *epistemic temporal models* are Kripke models with a designated Y-relation—these have an *internal* notion of time—whereas *epistemic models* are Kripke models without a designated Y-relation—these have an *external* notion of time. Since an epistemic temporal model M uses an *internal* notion of time, the ways in which the system described by M can evolve are determined in advance by the structure of the Y-relation in M; said informally, the *protocol is fixed*. In contrast, the protocol in Dynamic Epistemic Logic is *dynamic*, as it can be changed on-the-fly by using a different update modal to produce the next pointed Kripke model appearing in the sequence (1).

In extending the updates of standard Dynamic Epistemic Logic from the class of epistemic models (having *external* time) to the class of epistemic temporal models (having *internal* time), we stand to gain *dynamic protocols* for systems that *need not be synchronous*. While standard Dynamic Epistemic Logic sets each world in $M[U]$ to be one time-step ahead of any world in M, our new updates on epistemic temporal models allow us greater flexibility in modeling the passage of time. In particular, using the internal notion of time associated with the Y-relation, our updates allow us to let worlds in $M[U]$ have any natural-number time; therefore, in certain updates that embed M into $M[U]$, each world in $M[U]$ can be seen either as a world in M or else as an arbitrarily distant possible future of a world in M. Such flexibility is essential to the study of *asynchronous* systems. To bring about this flexibility, we add a new structural component to update modals: the $\underline{Y}$-*arrow*. We use $\underline{Y}$-arrows to specify exact positions in which the update $f_{[U,s]}$ is to insert Y-arrows in the updated model $M[U]$. We then identify sufficient conditions on our new update modals $[U, s]$ that will guarantee that the update $f_{[U,s]}$ embeds M into $M[U]$ or preserves properties such as synchronicity in the resulting epistemic temporal model. We use these conditions to show that epistemic temporal models that result from sequentially applying a proper subclass of our new kinds of updates are *isomorphic* to the generated sequences of epistemic models from standard Dynamic Epistemic Logic that have been studied by a number of authors

[5, 6, 10, 14, 15]. While [5, 6] showed that properties such as synchronicity are necessary of sequences generated in standard Dynamic Epistemic Logic, our isomorphism result demonstrates that the necessity of these properties stems from the inherent structure of standard Dynamic Epistemic Logic update modals $[U, s]$ themselves. This provides a *new perspective* on the results of [5, 6].

In the next section, we introduce the language L_{DETL} and the theory T_{DETL} of *Dynamic Epistemic Temporal Logic*. It is this theory that we use in reasoning about our new kinds of updates on epistemic temporal models. Due to space constraints, we will omit the proofs of our results; the interested reader can find full details in [13], an extended version of this paper.

2 Syntax

Notation 2.1 (A, Y, $\underline{Y}$). A is a finite nonempty set of symbols not containing the symbols Y and $\underline{Y}$. The members of A will be called *agents*.

To define L_{DETL}, we must first define the internal structure of update modals $[U, s]$. This structure is built on top of finite Kripke frames. If S is a nonempty set of symbols, then a *Kripke frame F (for S)* is a pair (W^F, R^F) consisting of a nonempty set W^F whose members are called *worlds* and a function $R^F : S \to (W^F \to 2^{W^F})$ mapping each symbol $m \in S$ to a function $R_m : W^F \to 2^{W^F}$; to say that F is *finite* means that W^F is finite.[1] The internal structure of update modals $[U, s]$ is given by the structure of the object U, called an *update frame*.

Definition 2.2. For a language L, whose formulas we call *L-formulas*, an *L-update frame* is a tuple $U = (W, R, \mathsf{p})$ satisfying the following: (W, R) is a finite Kripke frame for $A \cup \{Y, \underline{Y}\}$ that will be called the Kripke frame *underlying U*, and $\mathsf{p} : W \to L$ is a function mapping each world $s \in W$ to an *L*-formula $\mathsf{p}(s)$. A *state* in U is just a world in the Kripke frame underlying U. Notation: for an *L*-update frame U, we write W^U to denote the first element of the tuple U, we write R^U to denote the second element of the tuple U, and we write p^U to denote the third element of the tuple U. A *pointed L-update frame* is a pair (U, s) consisting of an *L*-update frame U and a state $s \in W^U$ that will be called the *point* of (U, s).

Update frames are also called "action models" (or "event models") in the Dynamic Epistemic Logic literature [1, 2, 3, 4, 8, 12]. For an update frame U, a state $s \in W^U$ represents the communication of the formula $\mathsf{p}^U(s)$. For an agent $a \in A$, the relation R_a^U represents agent a's conditional uncertainty as to which formula is communicated: if $s' \in R_a^U(s)$ and the formula $\mathsf{p}^U(s)$ was in fact communicated, then agent a will think that the formula $\mathsf{p}^U(s')$ is one of the formulas that might have be communicated.

We now define our language L_{DETL} as an extension of the language L_{ETL} of Simple Epistemic Temporal Logic.

Definition 2.3 (L_{ETL}). L_{ETL}, the *Language of Simple Epistemic Temporal Logic*, consists of the formulas formed by the following grammar.

$$\varphi \ ::= \ \bot \mid \top \mid p_k \mid \varphi \star \varphi \mid \neg \varphi \mid [a]\varphi$$
$$k \in \mathbb{N}, \star \in \{\to, \vee, \wedge, \equiv\}, a \in A \cup \{Y\}$$

---

[1]The function R_a^F gives rise to a binary relation $\bar{R}_a^F := \{(x, y) \in W^F \times W^F \mid y \in R_a^F(x)\}$ on W^F. We will conflate R_a^F and $\bar{R}_a^F$ whenever it is convenient. We will often refer to the members of $\bar{R}_a^F$ as *a-arrows*.

$$
\begin{array}{ll}
\text{CL.} & \text{Schemes for Classical Propositional Logic} \\[4pt]
K_a. & [a](\varphi \to \psi) \to ([a]\varphi \to [a]\psi) \text{ for } a \in A \\[4pt]
K_Y. & [Y](\varphi \to \psi) \to ([Y]\varphi \to [Y]\psi) \\[4pt]
\text{UA.} & [U,s]q \equiv \big(\mathsf{p}^U(s) \to q\big) \text{ for } q \in \{p_k, \bot, \top\} \\[4pt]
\text{U}\star. & [U,s](\varphi \star \psi) \equiv \big([U,s]\varphi \star [U,s]\psi\big) \text{ for } \star \in \{\to, \vee, \wedge, \equiv\} \\[4pt]
\text{U}\neg. & [U,s]\neg\varphi \equiv \big(\mathsf{p}^U(s) \to \neg[U,s]\varphi\big) \\[4pt]
\text{U}[a]. & [U,s][a]\varphi \equiv \big(\mathsf{p}^U(s) \to \bigwedge_{s' \in R_a^U(s)}[a][U,s']\varphi\big) \text{ for } a \in A \\[4pt]
\text{U}[Y]. & [U,s][Y]\varphi \equiv \big(\mathsf{p}^U(s) \to \bigwedge_{s' \in R_Y^U(s)}[Y][U,s']\varphi\big) \wedge \\
& \qquad\qquad \big(\mathsf{p}^U(s) \to \bigwedge_{s' \in R_{\underline{Y}}^U(s)}[U,s']\varphi\big)
\end{array}
$$

$$
\frac{\vdash \varphi \to \psi \quad \vdash \varphi}{\vdash \psi}\,(\text{MP}) \qquad
\frac{a \in A \cup \{Y\} \quad \vdash \varphi}{\vdash [a]\varphi}\,(\text{MN}) \qquad
\frac{\vdash \varphi}{\vdash [U,s]\varphi}\,(\text{UN})
$$

Figure 1: The theory T_{DETL}

Terminology: we call $[Y]$ the *yesterday modal*. For each agent $a \in A$, we read the formula $[a]\varphi$ as "agent a believes that φ is true." We read the formula $[Y]\varphi$ as "φ is true in all possible yesterdays." Notation: for each $a \in A \cup \{Y\}$, we let $\langle a \rangle$ abbreviate $\neg[a]\neg$; we define for each $i \in \mathbb{N}$ the formula $[a]^i\varphi$ by setting $[a]^0\varphi := \varphi$ and $[a]^{i+1}\varphi := [a]([a]^i\varphi)$; for $i \in \mathbb{N}$, the formula $\langle a \rangle^i\varphi$ is defined analogously.

Definition 2.4 (L_{DETL}, T_{DETL})**.** L_{DETL} is the *Language of Dynamic Epistemic Temporal Logic*. The L_{DETL}-*formulas* are the formulas that may be formed by the grammar obtained from that in Definition 2.3 by adding the following formula-formation rule: if φ is an L_{DETL}-formula and (U,s) is a pointed L-update frame with $\emptyset \neq L \subseteq L_{\mathsf{DETL}}$, then $[U,s]\varphi$ is an L_{DETL}-formula. L_{DETL} consists of the L_{DETL}-formulas along with the L-update frames for which $\emptyset \neq L \subseteq L_{\mathsf{DETL}}$. Terminology: we call $[U,s]$ an *update modal*. Notation: we let $\langle U,s \rangle$ abbreviate $\neg[U,s]\neg$. We read the formula $[U,s]\varphi$ as "after update (U,s), φ is true." An *update frame* is an L_{DETL}-update frame. A *formula* is a L_{DETL}-formula. T_{DETL}, the *Theory of Dynamic Epistemic Temporal Logic*, is defined in Figure 1.

Since our interest here is in implementing update mechanisms on Kripke models with a designated Y-relation, we do not impose any of the usual properties on belief or on time that one might expect [5, 6, 9, 10, 14, 15]. So T_{DETL} should be viewed as the *minimal* theory that brings update mechanisms to Simple Epistemic Temporal Logic. Future work will investigate extensions of this theory that include familiar restrictions on belief and on time, though we do address the preservation of certain time-related properties in Section 5.

3 Semantics

Having defined the language L_{DETL} and theory T_{DETL} of Dynamic Epistemic Temporal Logic, we now define the semantics of L_{DETL}. A *Kripke model* M is a tuple (W^M, R^M, V^M) consisting of

a Kripke frame (W^M, R^M) and a function $V^M : \{p_k \mid k \in \mathbb{N}\} \to 2^{W^M}$ called a *(propositional)* *valuation*. A *pointed Kripke model* is a pair (M, w) consisting of a Kripke model M and a world $w \in W^M$. The notion of L_{DETL}-truth extends the standard semantics for Dynamic Epistemic Logic [1, 2, 3, 4, 8, 12] in the following way.

Definition 3.1 (L_{DETL}-Truth, L_{DETL}-Validity). For a pointed Kripke model (M, w) and a formula φ, we write $M, w \models_{L_{\mathsf{DETL}}} \varphi$ to mean that φ is *true at* (M, w), and we write $M, w \not\models_{L_{\mathsf{DETL}}} \varphi$ to mean that φ is not true at (or *false at*) (M, w). The notion of truth of a formula at a pointed Kripke model is defined by an induction on formula construction; we omit the Boolean cases.

- $M, w \models_{L_{\mathsf{DETL}}} [a]\varphi$ means that $M, x \models_{L_{\mathsf{DETL}}} \varphi$ for each $x \in R_a^M(w)$.

- $M, w \models_{L_{\mathsf{DETL}}} [U, s]\varphi$ means that if $M, w \models_{L_{\mathsf{DETL}}} \mathsf{p}^U(s)$, then $M[U], (w, s) \models_{L_{\mathsf{DETL}}} \varphi$, where the model $M[U]$ is defined as follows.

$$
\begin{aligned}
W^{M[U]} &:= \{(x, t) \in W^M \times W^U \mid M, x \models_{L_{\mathsf{DETL}}} \mathsf{p}^U(t)\} \\
\text{For } a \in A : R_a^{M[U]}(x, t) &:= \{(y, u) \in W^{M[U]} \mid y \in R_a^M(x) \text{ and } u \in R_a^U(t)\} \\
R_Y^{M[U]}(x, t) &:= \{(y, u) \in W^{M[U]} \mid y \in R_Y^M(x) \text{ and } u \in R_Y^U(t)\} \cup \\
&\quad\; \{(y, u) \in W^{M[U]} \mid y = x \text{ and } u \in R_{\underline{Y}}^U(t)\} \\
V^{M[U]}(p_k) &:= \{(x, t) \in W^{M[U]} \mid M, x \models_{L_{\mathsf{DETL}}} p_k\}
\end{aligned}
$$

To say that a formula φ is *valid in* a Kripke model M, written $M \models_{L_{\mathsf{DETL}}} \varphi$, means that $M, w \models_{L_{\mathsf{DETL}}} \varphi$ for each world $w \in W^M$. To say that a formula φ is *valid*, written $\models_{L_{\mathsf{DETL}}} \varphi$, means that $M \models_{L_{\mathsf{DETL}}} \varphi$ for each Kripke model M. When it ought not cause confusion, we may omit the subscript "L_{DETL}" when writing $\models_{L_{\mathsf{DETL}}}$.

Given a pointed Kripke model (M, w) representing a multi-agent situation and a pointed update frame (U, s) with $M, w \models \mathsf{p}^U(s)$, the pointed Kripke model $\big(M[U], (w, s)\big)$ represents the situation after the occurrence of the update described by $[U, s]$. According to Definition 3.1, a world (x, t) must satisfy the property that $M, x \models \mathsf{p}^U(t)$. The set $\{x \in W^M \mid M, x \models \mathsf{p}^U(t)\}$ of worlds x in M that satisfy $\mathsf{p}^U(t)$ intuitively represents the set of worlds in M at which the formula $\mathsf{p}^U(t)$ can truthfully be communicated—these are the worlds at which t can take place.

For each $a \in A$, Definition 3.1 tells us that the relation $R_a^{M[U]}$ is determined by two factors: agent a's uncertainty as to which world was the case before the communication (represented by R_a^M) and agent a's uncertainty as to which communication has occurred (represented by R_a^U). In particular, suppose $(x', t') \in R_a^{M[U]}(x, t)$. Then if the communication corresponding to t actually occurred at world x, then agent a will think it possible that the communication corresponding to t' occurred at world x'.

According to Definition 3.1, the relation $R_Y^{M[U]}$ is determined by two factors. The first is the interaction between the relations R_Y^U and R_Y^M, which adds pairs to $R_Y^{M[U]}$ just as the interaction between R_a^U and R_a^M did to $R_a^{M[U]}$ for $a \in A$. The second factor is the relation $R_{\underline{Y}}^U$: if there is a $\underline{Y}$-arrow from state t to state t' in U, then there will be a Y-arrow from world (x, t) to world (x, t') in $M[U]$. The presence of a $\underline{Y}$-arrow from t to t' in U thus says that the communication corresponding to t' is to be thought of as occurring one time-step before the communication corresponding to t. This addition to the standard definition of updates in Dynamic Epistemic Logic [1, 2, 3, 4, 8, 12] allows us to control how an update affects the time of worlds in the model $M[U]$.

Finally, we see that the valuation $V^{M[U]}$ after the update simply inherits its truth conditions from the valuation V^M before the update, making our updates *purely temporal-epistemic*.
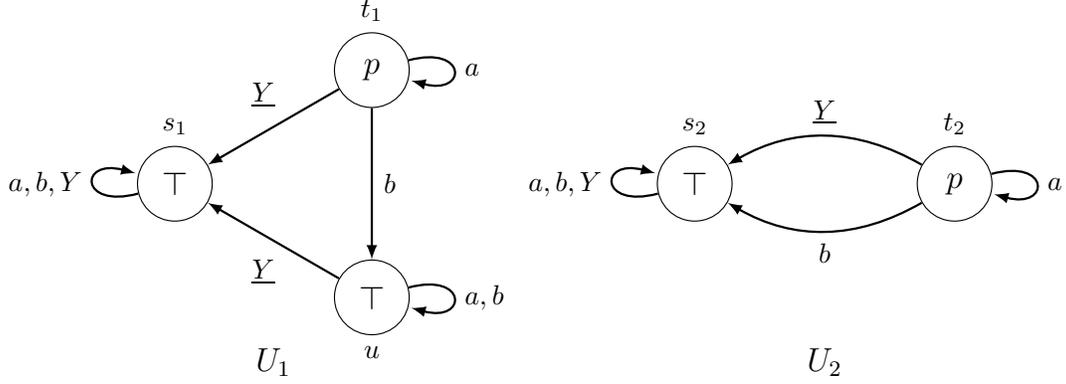
Figure 2: Update frames for the synchronous (U_1, t_1) and asynchronous (U_2, t_2) private announcement of p to a.

Theorem 3.2 (Correctness; [13]). *For each formula φ, we have $\vdash \varphi$ if and only if $\models \varphi$.*

4   A Simple Example

Suppose Passengers a and b are traveling together by train in China. Further, suppose Passenger a understands Mandarin but that Passenger b does not, though Passenger b mistakenly believes that they are both equally ignorant of the language. Now consider two scenarios in which an announcement in Mandarin about a delay in arrival is made over the loudspeaker.

1. Passengers a and b are both awake and alert during the announcement.

2. Passenger a is awake and alert, but Passenger b, who is sleepy, dozes off and sleeps through the announcement. Waking up a few minutes later without knowing that the announcement occurred, Passenger b mistakenly thinks that instead of sleeping for a few minutes, he merely blinked.

Taking p to be a propositional letter denoting the statement about late arrival, we represent the first and the second scenarios in our framework using update frames (U_1, t_1) and (U_2, t_2), respectively pictured on the left and on the right in Figure 2.

In the first scenario, Passenger b knows that an announcement has taken place, but it provides him with no new (non-temporal) information—nor does he believe that a gained any (non-temporal) information. In effect, this is a *synchronous private announcement to a*; after all, both a and b know that an announcement occurred—so the event is *synchronous*—but only a knows the content of the announcement—so the event is *private to a*. In Figure 2, s_1 and u are states in which no new (non-temporal) information is conveyed (since $\top$ is always true and thus conveys no new non-temporal information), while t_1 is a state in which the message p is communicated. Since t_1 and u are each connected to s_1 using a $\underline{Y}$-arrow, the communications they represent occur one time-step after the communication represented by s_1.

Since s_1 is labeled by $\top$, has a reflexive x-arrow for every $x \in \{a, b, Y\}$, and has no exiting $\underline{Y}$-arrows, we see by the definition of truth (Definition 3.1) that any Kripke model M is embedded into the Kripke model $M[U_1]$ by the mapping taking each world $y \in W^M$ to the world $(y, s_1) \in W^{M[U_1]}$. This embedding preserves a copy of the "past situation" M within the "current situation" $M[U_1]$,

6

which leads us to call s_1 a "past state." So the role of the past state s_1 is to preserve a copy of a given situation M. The states t_1 and u then represent communications that occur one time-step after the situation M. At state t_1, Passenger a believes that t_1 represents the only possible communication, while Passenger b believes that u represents the only possible communication. Since both u and t_1 are one time-step after the past state s_1, the update $f_{[U_1, t_1]}$ describes the private communication of p to Passenger a in which it is common knowledge that one time-step occurs. So we see that

$$\models (\neg\langle Y\rangle\top \wedge \neg[b]p) \rightarrow [U_1, t_1]\big([a]\langle Y\rangle\top \wedge [a]p \wedge [b]\langle Y\rangle\top \wedge \neg[b]p\big).$$

That is, if no event has yet occurred and Passenger b does not believe p, then, after the occurrence of $f_{[U_1, t_1]}$, Passenger a believes that an event occurred and that p is true, whereas Passenger b believes that an event occurred but does not believe that p is true.

In contrast, the second scenario is in effect an *asynchronous private announcement to a*. After all, while Passenger a knows that an announcement occurred and she knows its content, Passenger b has two mistaken beliefs: first, that no announcement occurred, and second, that the amount of time between closing and later opening his eyes is essentially negligible. b thus *does not even think it possible* that an event has occurred. Since the announcement results in b having a mistaken belief about the number of events that have occurred, the announcement event is *asynchronous*. At state t_2 in Figure 2, Passenger a knows that p is communicated, but Passenger b mistakenly believes that *no event took place* because the only state he considers possible is the past state s_2. Accordingly, we see that

$$\models (\neg\langle Y\rangle\top \wedge \neg[b]p) \rightarrow [U_2, t_2]\big([a]\langle Y\rangle\top \wedge [a]p \wedge \neg[b]\langle Y\rangle\top \wedge \neg[b]p\big).$$

That is, if no event has yet occurred and Passenger b does not believe that p is true, then, after the occurrence of $f_{[U_2, t_2]}$, Passenger a believes that an event occurred and that p is true, whereas Passenger b believes neither that an event occurred nor that p is true.

These scenarios demonstrate the way in which our framework uses $\underline{Y}$-arrows to describe *synchronous* and *asynchronous* private communications. In particular, we see that $\underline{Y}$-arrows can be used to describe updates that *need not preserve synchronicity*, as is the case with the asynchronous private announcement.

5   Properties and Preservation

In this section, we define several properties of Kripke models and update frames and then study sufficient conditions for the preservation of these properties after the occurrence of an update.

Definition 5.1 (*T-Runs, T-Histories, T-Depth*). Fix a symbol $T \in \{Y, \underline{Y}\}$ and let $F = (W, R)$ be a Kripke frame for $A \cup \{Y, T\}$. A *T-run* (in F) is a finite nonempty sequence $\{w_i\}_{i=0}^n$ of worlds in F satisfying the property that $n \in \mathbb{N}$ and for each $i \in \mathbb{N}$ with $i < n$, we have that $w_{i+1} \in R_T^F(w_i)$. We say that a *T*-run $\{w_i\}_{i=0}^n$ *begins at* w_0 and *ends at* w_n. The *length* of a *T*-run $\{w_i\}_{i=0}^n$ is defined as the number n. (Observe that the length of a *T*-run is one less than the number of worlds that make up the *T*-run.) To say that a *T*-run σ' *end-extends* a *T*-run σ means that σ is a (not necessarily proper) prefix of σ'. (Note that each *T*-run end-extends itself.) To say that a *T*-run σ is *end-maximal* (in F) means that no *T*-run in F end-extends σ. A *T-history* (in F) is a *T*-run in F that is end-maximal. (Note that a suffix of a *T*-history is itself a *T*-history.) A world appearing at the end of a *T*-history in F is said to be *T-terminal* (in F). We define a function

$d_T^F : W^F \to \mathbb{N} \cup \{\infty\}$ as follows: if there is a maximum $n \in \mathbb{N}$ such that there is a T-history in F of length n that begins at w, then $d_T^F(w)$ is n; otherwise, if no such maximum $n \in \mathbb{N}$ exists, then $d_T^F(w)$ is ∞. We will call $d_T^F(w)$ the *T-depth of w*.

Definition 5.2. Fix $T \in \{Y, \underline{Y}\}$ and let $F = (W, R)$ be a Kripke frame for $A \cup \{Y, T\}$.

- *T-Depth–Defined (T-DD).* To say that F is *T-depth–defined (T-DD)* means that for each world w in F, we have that $d_T^F(w) \neq \infty$.[2]

- *Non–T-Branching.* To say that F is *non–T-branching* means that for each $w \in W^F$, the set $R_T^F(w)$ has at most one member.

- *T-Synchronous.* If F is T-DD, then to say that F is *T-synchronous* means that for each $a \in A$, each $w \in W^F$, and each $w' \in R_a^F(w)$, we have that $d_T^F(w') = d_T^F(w)$. The negation of "T-synchronous" is *T-asynchronous*.

Convention: for tuples J having a Kripke frame (W^J, R^J) underlying J, any use of a property or concept from Definition 5.1 or Definition 5.2 in reference to J is meant to be a use of that property or concept in reference to the Kripke frame (W^J, R^J) underlying J. Example: for an update frame U, the expression "$\underline{Y}$-run in U" is to be identified with the expression "$\underline{Y}$-run in (W^U, R^U)."

Definition 5.3 (Kripke Model Properties)**.** Let M be a Kripke model.

- *Synchronicity (under Y-DD).* If M is Y-DD, then to say that M is *synchronous* means that M is Y-synchronous. The negation of "synchronous" is *asynchronous*.

- *Non–Past-Branching.* To say that M is *non–past-branching* means that M is non–Y-branching.

- *Forest-like.* To say that M is *forest-like* means that M is Y-DD and non–past-branching.

Definition 5.4 (Update Frame Properties and Concepts)**.** Let U be an update frame.

- *Path-Preserving.* A *path-preserving run* (in U) is a $\underline{Y}$-run $\{s_i\}_{i=0}^n$ in U satisfying the property that for each $i \in \mathbb{N}$ with $i < n$, we have $\models \mathsf{p}^U(s_i) \to \mathsf{p}^U(s_{i+1})$. To say that U is a *path-preserving update frame* means that each $\underline{Y}$-run in U is path-preserving.

- *Depth-Respecting (under $\underline{Y}$-DD).* If U is $\underline{Y}$-DD, then to say that U is *depth-respecting* means that for each $s \in W^U$ and each $s' \in R_{\underline{Y}}^U(s)$, we have that $d_{\underline{Y}}^U(s') \leq d_{\underline{Y}}^U(s)$.

- *Past State, Past-Preserving.* A *past state* is a state $s \in W^U$ satisfying the property that $\mathsf{p}^U(s) = \top$, that $R_{\underline{Y}}^U(s) = \emptyset$, and that $R_a^U(s) = \{s\}$ for each $a \in A \cup \{Y\}$. To say that U is *past-preserving* means U is $\underline{Y}$-DD and path-preserving and that every $\underline{Y}$-run in U can be end-extended to a $\underline{Y}$-history in U that ends at a past state.

- *Non–Past-Splitting.* To say that U is *non–past-splitting* means that for each $s \in W^U$, we have that $R_{\underline{Y}}^U(s) \cup R_Y^U(s)$ has at most one element and that $R_{\underline{Y}}^U(s) \cap R_Y^U(s) = \emptyset$.

---

[2]We observe that if F is T-depth–defined, then F is *T-converse well-founded* (that is, for every nonempty set S of worlds in F, there is a nonempty subset $S' \subseteq S$ such that for each $w \in S'$, the unique T-run in F that begins at w has length zero). However, if F is T-converse well-founded, it need not be the case that F is also T-depth–defined. So the notion of T-depth–definedness is strictly stronger than the notion of T-converse well-foundedness.

Having defined these properties, we investigate their preservation under the presence of updates in the following two theorems. Theorem 5.5 concerns the behavior of past states in update frames, and Theorem 5.6 concerns the preservation of properties in Kripke models.

Theorem 5.5 (Past State Theorem; [13]). Let U be an update frame and M be a Kripke model.

- If s is a past state in U, then for each $\varphi \in L_{\mathsf{DETL}}$ and each $w \in W^M$, we have that $M[U], (w, s) \models \varphi$ if and only if $M, w \models \varphi$.

- If U is past-preserving and non–past-spliting, $s \in W^U$ has $\mathrm{d}_Y^U(s) = n$, and $w \in W^M$ satisfies $M, w \models \mathsf{p}^U(s)$, then for each $\varphi \in L_{\mathsf{DETL}}$, we have that $M[U], (w, s) \models \langle Y \rangle^n \varphi$ if and only if $M, w \models \varphi$.

Theorem 5.5 tells us that past states play the role of "maintaining a link to the past" within past-preserving, non–past-splitting update frames. In particular, if s is a past state, then the submodel of $M[U]$ consisting of the worlds of the form (w, s) for some world $w \in W^M$ is L_{DETL}-indistinguishable from the Kripke model M itself. So the operation $(M, w) \mapsto (M[U], (w, s))$ retains a copy of the "past" state of affairs (M, w). Furthermore, if U is past-preserving, then from any world in $W^{M[U]}$, there is a finite sequence of Y-arrows that leads back to this "past" state of affairs, thereby "maintaining a link to the past."

Let us now examine the preservation of properties of the Kripke model M in the presence of the operation $M \mapsto M[U]$.

Theorem 5.6 (Preservation Theorem; [13]). Let (U, s) be a pointed update frame and (M, w) be a pointed Kripke model such that $M, w \models \mathsf{p}^U(s)$.

- *Y-DD.* If M is Y-DD and U is $\underline{Y}$-DD and depth-respecting, then $M[U]$ is Y-DD.

- *Synchronicity.* If M is synchronous (and Y-DD) and U is $\underline{Y}$-DD, depth-respecting, past-preserving, and $\underline{Y}$-synchronous, then $M[U]$ is synchronous.

- *Non–Past-Branching.* If M is non–past-branching and U is non–past-splitting, then $M[U]$ is non–past-branching.

- *Forest-likeness.* If M is forest-like and U is $\underline{Y}$-DD and non–past-splitting, then $M[U]$ is forest-like.

6   Embedding Standard DEL

In this section, we show that standard (Temporal) Dynamic Epistemic Logic, whose update modals contain neither $\underline{Y}$- nor Y-arrows, can be embedded in our framework in a natural way. This provides clear connections between our work and the work in [5, 6, 10, 14, 15] on (Temporal) Dynamic Epistemic Logic, which will be described at the end of this section.

Definition 6.1 (Standard). Choose $T \in \{Y, \underline{Y}\}$. To say that a Kripke frame F for $A \cup \{Y, T\}$ is *standard* means that for each $s \in W^F$ and each $m \in \{Y, T\}$, we have $R_m^F(s) = \emptyset$. To say that a Kripke model or an L-update frame is *standard* means that the Kripke frame underlying that model or L-update frame is standard. To say that a pointed Kripke model or a pointed L-update frame is *standard* means that the Kripke model or L-update frame making up the first component of the pair is standard.

Definition 6.2 (L_{TDEL}; [10]). L_{TDEL} is the *Language of Temporal Dynamic Epistemic Logic*. The L_{TDEL}-*formulas* are the formulas that may be formed by the grammar obtained from that in Definition 2.3 (the definition of L_{ETL}) by adding the following formula-formation rule: if φ is an L_{TDEL}-formula and (U, s) is a standard pointed L-update frame with $\emptyset \neq L \subseteq L_{\mathsf{TDEL}}$, then $[U, s]\varphi$ is an L_{TDEL}-formula. L_{TDEL} consists of the L_{TDEL}-formulas along with the L-update frames for which $\emptyset \neq L \subseteq L_{\mathsf{TDEL}}$.

Notation 6.3 (Sequences). Let τ be a finite possibly empty sequence. We write $\tau \cdot x$ to denote the sequence obtained from τ by adding x at the end. $|\tau|$ denotes the number of elements in τ.

Definition 6.4 (Adapted from [5, 6, 11, 14, 15]). A *run* is a nonempty finite sequence $\{M_i\}_{i=0}^n$ of Kripke models satisfying the property that for each $i \in \mathbb{N}$ with $0 < i \leq n$ and each $w \in W^{M_i}$, we have that w is of the form $(\pi(w), s)$ for some world $\pi(w) \in W^{M_{i-1}}$. A *pointed run* is a pair $(r \cdot M, w)$ consisting of a run $r \cdot M$ and a world $w \in W^M$; the world w is called the *point* of $(r \cdot M, w)$. A *standard (pointed) run* is a (pointed) run whose constituent pointed Kripke models are all standard. An *L event-run* is a finite possibly empty sequence of pointed L-update frames. A *standard L event-run* is an L event-run whose constituent pointed L-update frames are all standard.

Definition 6.5 (L_{TDEL}-Truth; [5, 6, 10]). We define a notion of truth for L_{TDEL}-formulas at standard runs r by an induction on the construction of L_{TDEL}-formulas; we consider only the non-Boolean cases.

- For $a \in A$: $r \cdot M, w \models_{L_{\mathsf{TDEL}}} [a]\varphi$ means that $r \cdot M, x \models_{L_{\mathsf{TDEL}}} \varphi$ for each $x \in R_a^M(w)$.

- $r \cdot M, w \models_{L_{\mathsf{TDEL}}} [Y]\varphi$ means that if $|r| > 0$, then $r, \pi(w) \models_{L_{\mathsf{TDEL}}} \varphi$.

- $r \cdot M, w \models_{L_{\mathsf{TDEL}}} [U, s]\varphi$ means that if we have $r \cdot M, w \models_{L_{\mathsf{TDEL}}} \mathsf{p}^U(s)$, then, letting $r' := r \cdot M$, it follows that $r' \cdot r'[U], (w, s) \models_{L_{\mathsf{TDEL}}} \varphi$, where $r'[U]$ is the standard Kripke model defined as follows.

$$
\begin{aligned}
W^{r'[U]} &:= \{(x, t) \in W^M \times W^U \mid r \cdot M, x \models_{L_{\mathsf{TDEL}}} \mathsf{p}^U(t)\} \\
\text{For } a \in A: \quad R_a^{r'[U]}(x, t) &:= \{(y, u) \in W^{r'[U]} \mid y \in R_a^M(x) \text{ and } u \in R_a^U(t)\} \\
R_Y^{r'[U]}(x, t) &:= \emptyset \\
V^{r'[U]}(p_k) &:= \{(x, t) \in W^{r'[U]} \mid r \cdot M, x \models_{L_{\mathsf{TDEL}}} p_k\}
\end{aligned}
$$

When it ought not cause confusion, we may omit the subscript "L_{TDEL}" in writing $\models_{L_{\mathsf{TDEL}}}$.

Definition 6.6 (Generated Structures). Let (M, w) be a standard pointed Kripke model.

- If $\sigma = \{(U_i, s_i)\}_{i=1}^n$ is an L_{DETL} event-run, then $(M, w) *^{\flat} \sigma$, the *pointed Kripke model that is point-generated from (M, w) by σ*, is the pointed Kripke model (M_m, w_m) appearing at the end of the sequence $\{(M_i, w_i)\}_{i=0}^m$ having the largest integer $m \leq n$ subject to the following restrictions: $(M_0, w_0) = (M, w)$ and for each $j \in \mathbb{N}$ with $j < m$, we have

  - $M_j, w_j \models_{L_{\mathsf{DETL}}} \mathsf{p}^{U_{j+1}}(s_{j+1})$ and
  - $(M_{j+1}, w_{j+1}) = (M_j[U_{j+1}], (w_j, s_{j+1}))$.

  Note: "$\models_{L_{\mathsf{DETL}}}$" and $M_j[U_{j+1}]$ are given by L_{DETL}-truth (Definition 3.1).

$$
\begin{aligned}
q^{\sharp n} &:= q \text{ if } q \in \{p_k, \bot, \top\} \\
(\varphi \star \psi)^{\sharp n} &:= \varphi^{\sharp n} \star \psi^{\sharp n} \\
(\neg \varphi)^{\sharp n} &:= \neg(\varphi^{\sharp n}) \\
([a]\varphi)^{\sharp n} &:= [a](\varphi^{\sharp n}) \text{ if } a \in A \text{ or } (a = Y \text{ and } n = 0) \\
([Y]\varphi)^{\sharp n} &:= [Y]\varphi^{\sharp(n-1)} \text{ if } n > 0 \\
([U,s]\varphi)^{\sharp n} &:= [U^{\sharp n}, s](\varphi^{\sharp(n+1)})
\end{aligned}
$$

$$
W^{U^{\sharp n}} := U^W \uplus \{\flat\} \quad \text{(disjoint union)}
$$

$$
\text{for } a \in A \cup \{Y, \underline{Y}\}, \qquad
R_a^{U^{\sharp n}}(s) :=
\begin{cases}
R_a^U(s) & \text{if } a \neq \underline{Y} \text{ and } s \neq \flat, \\
\{\flat\} & \text{if } a \neq \underline{Y} \text{ and } s = \flat, \\
\{\flat\} & \text{if } a = \underline{Y} \text{ and } s \neq \flat, \\
\emptyset & \text{if } a = \underline{Y} \text{ and } s = \flat.
\end{cases}
$$

$$
\mathsf{p}^{U^{\sharp n}}(s) :=
\begin{cases}
\left(\mathsf{p}^U(s)\right)^{\sharp n} \wedge \langle Y \rangle^n [Y]\bot & \text{if } s \neq \flat, \\
\top & \text{if } s = \flat.
\end{cases}
$$

Figure 3: Definition of $\sharp n : L_{\mathsf{TDEL}} \to L_{\mathsf{DETL}}$ for $n \in \mathbb{N}$

- If $\sigma = \{(U_i, s_i)\}_{i=1}^n$ is a standard L_{TDEL} event-run, then $(M, w) *^{\mathsf{s}} \sigma$, the *pointed run that is sequence-generated from (M, w) by σ*, is the pointed run $(\{M_i\}_{i=0}^m, w_m)$ obtained from the sequence $\{(M_i, w_i)\}_{i=0}^m$ of pointed Kripke models having the largest integer $m \leq n$ subject to the following restrictions: $(M_0, w_0) = (M, w)$ and for each $j \in \mathbb{N}$ with $j < m$, we have

  - $\{M_i\}_{i=0}^j, w_j \models_{L_{\mathsf{TDEL}}} \mathsf{p}^{U_{j+1}}(s_{j+1})$ and
  - $(M_{j+1}, w_{j+1}) = \left(\{M_i\}_{i=0}^j[U_{j+1}], (w_j, s_{j+1})\right)$.

  Note: "$\models_{L_{\mathsf{TDEL}}}$" and $\{M_i\}_{i=0}^j[U_{j+1}]$ are given by L_{TDEL}-truth (Definition 6.5).

Definition 6.7 ($\downarrow$). Let $(r, w) = (\{M_i\}_{i=0}^n, w)$ be the standard pointed run sequence-generated by a standard L_{TDEL} event-run from a standard pointed Kripke model. We write $(r, w)\downarrow$ to denote the pointed Kripke model (M, w) defined in the following way.

$$
\begin{aligned}
W^M &:= \bigcup_{i=0}^n W^{M_i} \\
R_a^M(v) &:= R_a^{M_i}(v) \text{ if } i \in \mathbb{N} \text{ and } v \in W^{M_i} \\
R_Y^M(v) &:=
\begin{cases}
\{v'\} & \text{if } v = (v', s) \in W^{M_i} \text{ and } i > 0 \\
\emptyset & \text{otherwise}
\end{cases}
\end{aligned}
$$

Definition 6.8 ($\sharp n, \sharp$). For $n \in \mathbb{N}$, we define the function $\sharp n : L_{\mathsf{TDEL}} \to L_{\mathsf{DETL}}$ in Figure 3. If $\sigma = \{(U_i, s_i)\}_{i=1}^n$ is a standard L_{TDEL} event-run, then we define $\sigma^\sharp := \{(U_i^{\sharp(i-1)}, s_i)\}_{i=1}^n$.

Theorem 6.9 (Isomorphism Theorem; [13]). Let (M, w) be a standard pointed Kripke model and let σ be a standard L_{TDEL} event-run. Defining $m := |(M, w) *^{\mathsf{p}} \sigma| - 1$, we have each of the following.

1. For $\varphi \in L_{\mathsf{TDEL}}$: $(M, w) *^{\mathfrak{s}} \sigma \models_{L_{\mathsf{TDEL}}} \varphi$ if and only if $(M, w) *^{\mathfrak{p}} \sigma^{\sharp} \models_{L_{\mathsf{DETL}}} \varphi^{\sharp m}$.

2. $\big((M, w) *^{\mathfrak{s}} \sigma\big)\!\downarrow$ and $(M, w) *^{\mathfrak{p}} \sigma^{\sharp}$ are isomorphic.[3]

The Isomorphism Theorem (Theorem 6.9) allows us to view results about Kripke models that have been sequence-generated by standard L_{TDEL} event-runs as results about (Temporal) Dynamic Epistemic Logic—and the other way around. In particular, [5, 6] studies certain structural properties of the forest structure given by a run $(M, w) *^{\mathfrak{s}} \sigma$ that has been sequence-generated from a standard pointed Kripke model (M, w) by a standard L_{ETL} event-run σ. In [5, 6], the authors define what it means for the run $(M, w) *^{\mathfrak{s}} \sigma$ to be *synchronous* (among other properties) and then show that every run sequence-generated from a standard pointed Kripke model by a standard L_{ETL} event-run is synchronous.[4] Our Preservation Theorem (Theorem 5.6) works together with the Isomorphism Theorem (Theorem 6.9) to provide a different perspective on this synchronicity result. In particular, our work shows that the results of [5, 6] can be viewed as a consequence of the structural properties that are present in an update frame $U^{\sharp n}$, produced from a standard update frame U, thereby pinpointing the source of the synchronicity result in the structure of standard update frames themselves.

7 Conclusion

In this paper, we showed how to extend the updates of Dynamic Epistemic Logic so that they operate not just on *epistemic models* but also on *epistemic temporal models* in a way that allowed us to control how an update affects the time of worlds in the model $M[U]$. This enabled us to extend the domain of applicability of the Dynamic Epistemic Logic approach to discrete-time multi-agent distributed systems that *need not be synchronous*. We then studied sufficient conditions for the preservation of various properties of Kripke models, such as synchronicity. Identifying an isomorphism that connects epistemic temporal models generated in our framework with epistemic temporal models generated by standard updates as in [5, 6], we saw that the necessity of synchronicity in standardly generated epistemic temporal models stems from the structure of standard updates themselves. We then presented two scenarios contrasting synchronous and asynchronous private announcements.

In its technical essence, this paper is about adding a new type of arrow—the $\underline{Y}$-arrow—to update frames and then studying what we can do when the operation $M \mapsto M[U]$ described on epistemic models in [1, 2] is extended by the $\underline{Y}$-arrow mechanism to epistemic temporal models in a way that allows us to control how the update affects the time of worlds in the model $M[U]$. Essentially, the $\underline{Y}$-arrow describes a *sufficient condition for the creation of Y-arrows* in the model $M[U]$ resulting from the occurrence of an update. Namely, when there is a $\underline{Y}$-arrow from state s to state s' in update frame U, then there should be a Y-arrow from state (x, s) to state (x, s') in $M[U]$. While this is one possible sufficient condition for the creation of a certain kind of arrow,

---

[3]To say that two (pointed) Kripke models are *isomorphic* means that there exists an isomorphism between them. An *isomorphism between Kripke models* M and M' is a bijection $f : W^M \to W^{M'}$ satisfying each of the following: (i) $v \in V^M(p_k)$ if and only if $f(v) \in V^{M'}(p_k)$ for each $k \in \mathbb{N}$, and (ii) $u \in R_a^M(v)$ if and only if $f(u) \in R_a^{M'}(f(v))$ for each $a \in A \cup \{Y\}$. An *isomorphism between pointed Kripke models* (M, w) and (M', w') is an isomorphism f between M and M' for which $f(w) = w'$. See [7] for more information.

[4]If $(M, w) *^{\mathfrak{s}} \sigma$ is a run sequence-generated from a standard pointed Kripke model (M, w) by a standard L_{ETL} event-run σ, then the definition in [5, 6] would have us say that $(M, w) *^{\mathfrak{s}} \sigma$ satisfies *synchronicity* if and only if $\big((M, w) *^{\mathfrak{s}} \sigma\big)\!\downarrow$ is synchronous (according to our Definition 5.3).

there other conditions we may wish to consider. In particular, examining the hybrid scheme

$$[U,s][a]\varphi \equiv \mathsf{p}^U(s) \to \bigwedge_{s' \in W^U} \forall z.\big(\mathsf{a}_a^U(s,s') \to @_z(\mathsf{p}^U(s') \to [U,s']\varphi)\big) \tag{2}$$

in which a_a^U is a function mapping pairs (s,s') of states in U to a formula (possibly containing z), we see that the function a_a^U allows us to express a *precondition for the creation of a-arrows* in the model $M[U]$ produced by a *generalized update frame* $U = (W, \mathsf{p}, \mathsf{a})$. The hybrid language of such generalized update frames, called the *arrow-precondition language*, allows us to describe a wide variety of arrow-creation conditions, including all of those mentioned in this paper [13]. Though there is much to be studied about this generalization, it may prove useful in extending Dynamic Epistemic Logic to a much wider class of applications.

References

[1] Baltag, A., van Ditmarsch, H.P., Moss, L.S.: Epistemic logic and information update. In Adriaans, P., van Benthem, J., eds.: Handbook on the Philosophy of Information. Elsevier (2008) 369–463

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

[3] Baltag, A., Moss, L., Solecki, S.: The logic of public announcements, common knowledge and private suspicions. In Gilboa, I., ed.: TARK 1998. (1998) 43–56

[4] van Benthem, J., van Eijck, J., Kooi, B.: Logics of communication and change. Information and Computation **204**(11) (2006) 1620–1662

[5] van Benthem, J., Gerbrandy, J., Hoshi, T., Pacuit, E.: Merging frameworks for interaction. Journal of Philosophical Logic (2009) In press.

[6] van Benthem, J., Gerbrandy, J., Pacuit, E.: Merging frameworks for interaction: DEL and ETL. In: Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK XI). (2007) 72–81

[7] Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. Cambridge University Press (2001)

[8] van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Synthese Library. Springer (2007)

[9] Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning about Knowledge. The MIT Press (1995)

[10] Hoshi, T., Yap, A.: Dynamic epistemic logic with branching temporal structures. Synthese (2009) In press.

[11] Parikh, R., Ramanujam, R.: A knowledge based semantics of messages. Journal of Logic, Language, and Information **12** (2003) 453–467

[12] Renne, B.: A survey of Dynamic Epistemic Logic. Manuscript (July 2008)

[13] Renne, B., Sack, J., Yap, A.: Dynamic Epistemic Temporal Logic. Extended Manuscript (June 2009)

[14] Sack, J.: Temporal languages for epistemic programs. Journal of Logic, Language, and Information **17**(2) (2008) 183–216

[15] Yap, A.: Dynamic epistemic logic and temporal modality. Forthcoming in Proceedings of Dynamic Logic Montréal (2007)