The Tree of Knowledge in Action: Towards a Common Perspective

Johan van Benthem and Eric Pacuit

abstract. We survey a number of decidablity and undecidablity results concerning epistemic temporal logic. The goal is to provide a general picture which will facilitate the 'sharing of ideas' from a number of different areas concerned with modeling agents in interactive social situations.

Keywords: Epistemic Temporal Logic, Decidability, Undecidability

1 Introduction

When thinking about rational agents facing choices, one appealing mathematical model recurs in the literature. From Borges' short story 'The Garden of Forking Paths' to a large number of technical paradigms sometimes at war, sometimes at peace, all invoke the picture of a branching tree of finite sequences of events with epistemic indistinguishability relations for agents between these sequences, reflecting their limited powers of observation. Indeed, tree models for computation, with branches standing for process evolutions over time, have long been studied in computer science, cf. [23, 24, 6, 1, 11]. Philosophers have studied similar models, now enriched with epistemic relations, for studying the behavior of intelligent human agents facing choices: see Thomason & Gupta [28], Belnap et al. [4] and Horty [15]. Epistemic models of events over time have also been proposed in computer science by various authors, witness Fagin et al. [7] and Parikh & Ramanujam [20, 21]. And finally, 'dynamic logics' of communication and information flow in the tradition of Baltag, Moss & Solecki [3] have tree models of events as their natural broader habitat.

Bringing together knowledge and temporal change is a natural move in modeling, but it is also a potentially dangerous one from a complexity perspective, as has been pointed out forcefully in Halpern & Vardi [12]. The context is clear from the literature cited just now. On the one hand, Rabin's Theorem tells us that the full monadic second-order logic of the tree of events ordered by the relation of 'initial segment', and provided with some finite set of successor functions is decidable [24]. This explains the decidability of purely temporal logics of events such as **CTL**, and others. Likewise, the tree-like nature of models explains the decidability of many modal logics (see [17]). In a slogan, 'Trees are Safe'. But on the other hand, we know that the monadic second-order logic, indeed, even the monadic

Paper submitted to Advances in Modal Logic 2006

Π^1_1-theory of the grid $\mathbb{N} \times \mathbb{N}$ is undecidable (see [13]). A grid is like a tree, but successors meet, and the resulting confluent structure is known to cause undecidability in many areas of modal logic ([16]), witness in particular the work of Gabbay et al. on 'product models' [9]. In one more slogan: 'Grids are Dangerous'.

Now, epistemic temporal logics live at a dangerous edge here. Even though they use Rabin-style tree models, they introduce additional epistemic indistinguishability relations which generate a 'second dimension', and if the language gets too powerful, enough grid structure can be encoded to cause undecidability. Illustrations for this again come from a wide range of papers. E.g., Thomas [27] points out, following Läuchli, how introducing a relation of 'simultaneity' into the Rabin tree makes the monadic second-order logic undecidable. Likewise, Halpern & Vardi show how epistemic-temporal logics of agents with Perfect Recall and No Learning can become undecidable [12]. But the situation is delicate, as small changes in an epistemic temporal language or class of models can affect the complexity of the logic in drastic ways.

In this paper, we position ourselves close to the edge of undecidability in a straightforward system of epistemic temporal logic. We will discuss a number of complexity results, on both sides of the edge, while pointing out how results from all different traditions mentioned here help illuminate the landscape. In doing so, we also have a broader aim. The area that we are describing consists of a number of different frameworks, whose practitioners either do not know about relevant work by others, or are not even on speaking terms. We feel that this is an unfortunate situation, since much is to be gained by seeing the commonality of one area of research here. As we shall see, issues are often the same, and notions and techniques can be borrowed freely. Our paper is one such contribution toward a merge.

2 Epistemic Temporal Logic

This section describes the basic models that we will use in this study. The intended interpretation is of a group of agents interacting in a social situation. For example, a group of agents having a conversation or playing a game. We are interested in how the agents' knowledge about the situation may change over time. Let Σ be a set of **events**. The events are the building blocks of our models, i.e., the primitive descriptions of a social situation. For example, an event might be a move in some game, or an message sent from one agent to another agent or even to a group of agents. The first basic assumption is that not all agents are aware of all events. The other assumptions we make involve the nature of time in our models.

First of all, we assume there is a global discrete clock (which the agents may or may not be aware of). Second we assume that the agents only have a finite capacity to remember events. This capacity may be unbounded, but at every moment the agents can be aware of only a finite number of events. Thus even if an infinite number of events have taken place, at any moment each agent will only be aware of a finite number of them. Therefore, we assume moments are elements of $\mathbb{N}$ and so there is a finite past with a possibly infinite future. Below we give the formal details of our models.

2.1 Our Models

Let Σ be any set. The elements of Σ are called **events**. Given any set X, let X^* denote the set of finite strings over X and X^ω the set of infinite strings over X. Elements of $\Sigma^* \cup \Sigma^\omega$ will be called **histories**. Given $H \in \Sigma^* \cup \Sigma^\omega$, let $\mathsf{len}(H)$ denote the **length** of H, i.e. the number of characters (possibly infinite) in H. Given $H, H' \in \Sigma^* \cup \Sigma^\omega$, we write $H \preceq H'$ if H is a *finite* prefix of H'. If $H \preceq H'$ we say that H is an **initial segment** of H' and H' is an **extension** of H. Let $\mathsf{FinPre}(\mathcal{H}) = \{H \mid \exists H' \in \mathcal{H} \text{ such that } H \preceq H'\}$ be the set of finite prefixes of the elements of $\mathcal{H}$.

DEFINITION 1. Let Σ be any set. A set $\mathcal{H} \subseteq \Sigma^* \cup \Sigma^\omega$ is called a **protocol** provided $\mathsf{FinPre}(\mathcal{H}) \subseteq \mathcal{H}$.

Intuitively, a protocol is the set of all possible ways an interactive situation may evolve. Given a protocol $\mathcal{H}$ and a finite history $H \in \mathcal{H}$, let $\mathsf{Ext}_\mathcal{H}(H) = \{H' \mid H' \in \mathcal{H}, H \preceq H'\}$ be the set of extensions of H from $\mathcal{H}$. If no confusion will arise we will write $\mathsf{Ext}(H)$ instead of $\mathsf{Ext}_\mathcal{H}(H)$. We write $\mathsf{Ext}^{<\omega}(H)$ for the set of **finite extensions** of H and $\mathsf{Ext}^\omega(H)$ for the set of **infinite extensions** of H. Given $t \in \mathbb{N}$ and a history H, we write H_t for the unique initial segment of H of length t.

DEFINITION 2. Suppose $\mathcal{A}$ is a finite set of agents and Σ a finite set of events. A pair $\langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ is an **epistemic temporal frame** based on a set of events Σ if $\mathcal{H} \subseteq \Sigma^* \cup \Sigma^\omega$ is a protocol and each $\sim_i$ is an equivalence relation on $\mathsf{FinPre}(\mathcal{H})$.

Note that in the definition of an epistemic temporal frame, the set of events is assumed to be finite. Typically, with only few exceptions, we will be interested in finitely branching models. Thus, in general, an epistemic temporal frame can be pictured as a forest of finitely branching trees (there may be more than one root) with relations between finite branches of the trees. This corresponds to the situation that at any moment only finitely many events may take place and there may be uncertainty about the initial situation. For this paper, we are typically interested in *trees*. That is, unless otherwise stated, we assume that our protocols have a root. We now move to assumptions about the reasoning capabilities of our agents.

DEFINITION 3. An epistemic temporal frame $\langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ based on Σ satisfies the property **no learning**, or **no miracles**, provided for all finite histories $H, H' \in \mathcal{H}$ and events $e \in \Sigma$ such that $He \in \mathcal{H}$ and $H'e \in \mathcal{H}$, if $H \sim_i H'$ then $He \sim_i H'e$.

DEFINITION 4. An epistemic temporal frame $\langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ based on Σ satisfies the property **perfect recall** provided for all finite histories $H, H' \in \mathcal{H}$ and events $e \in \Sigma$ such that $He \in \mathcal{H}$ and $H'e \in \mathcal{H}$, if $He \sim_i H'e$ then $H \sim_i H'$.

Intuitively, perfect recall means that the set of histories an agent considers possible can only decrease or remain the same. The no learning property, which we sometimes refer to as the more suggestive 'no miracles' property, says that the uncertainty of the agents cannot be erased by the same event. Epistemic temporal frames that satisfy no learning and perfect recall will play an important role for us in this paper (see Section 4).

DEFINITION 5. An epistemic temporal frame $\langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ is **synchronous** if for all finite histories $H, H' \in \mathcal{H}$, if $H \sim_i H'$ then $\mathsf{len}(H) = \mathsf{len}(H')$.

Intuitively, if a frame satisfies the synchronous property, the value of the global clock is common knowledge. If a frame is not synchronous, then it is said to be **asynchronous**. There are a number of other assumptions that can be made about the interaction between the epistemic relation and time. The reader is referred to [7] for more information.

There have been a number of different modal languages proposed to reason about the above structures (see the Handbook chapter [14] for a discussion). The differences typically center around whether the temporal operators are 'branching' or 'linear'. We will say more about this distinction below. Our strategy will be to first introduce a modal logic with linear temporal operators, then to present a stronger language with features of both branching and linear time operators. Let At be a finite or countable set of atomic propositions. We are interested in language with various combinations of the following modalities: $P\phi$ (ϕ is true *sometime* in the past), $F\phi$ (ϕ is true *sometime* in the future), $Y\phi$ (ϕ is true at *the* previous moment), $N\phi$ (ϕ is true at *the* next moment), $K_i\phi$ (agent i knows ϕ) and $C_B\phi$ (the group $B \subseteq \mathcal{A}$ commonly knows ϕ).

Duals of the above operators are denoted as usual (eg., let $\langle i\rangle\phi$ denote $\neg K_i\neg\phi$). If X is a sequence of modalities from $\{P, F, Y, N\}$ let $\mathcal{L}_n^X$ be the language with n knowledge modalities $K_1, \ldots, K_n$ together with the modalities from X. If X is a sequence of modalities, then $\mathcal{L}_C^X$ will be the language $\mathcal{L}_n^X$ closed under the common knowledge modality C. Let $\mathcal{L}_{ETL}$ be the full epistemic temporal language, i.e., it contains all of the above temporal and knowledge operators.

We now introduce a **PDL**-style language intended to capture features of both linear and branching time language. Let $\mathcal{A}$ be a (finite) set of agents and recall that Σ is a (finite) set of events. Define the language $\mathcal{L}_\Sigma(\mathcal{A})$ inductively as follows:

$$\phi := p \mid \neg\phi \mid \phi \wedge \psi \mid \langle\alpha\rangle\phi$$

$$\alpha := a \mid ?p \mid \alpha;\beta \mid \alpha \cup \beta \mid \alpha^*$$

where $p \in \mathsf{At}$, $a \in \Sigma \cup \mathcal{A}$ and $\sigma \in \Sigma$. Let $\mathcal{L}_\Sigma(\mathcal{A})^-$ be the language $\mathcal{L}_\Sigma(\mathcal{A})$ which allows expressions of the forms $\langle\sigma^-\rangle\phi$.

Regardless of whether the language has branching time or linear time temporal operators, formulas are assumed to express properties about finite histories. The difference lies in format of the satisfaction relation. That is, in a linear temporal logic setting, formulas are interpreted at pairs H, t where H is a 'maximal' (possibly infinite) history and t is an element of

$\mathbb{N}$. The intended interpretation of $H, t \models \phi$ is that *on the branch* H at time t, ϕ is true. In the branching time setting, essentially we only need the moment. That is, formulas can be interpreted at finite histories H. In the interest of a unified approach we will interpret formulas at branch-time pairs. However, it what follows it will be useful to take the branching time interpretation. This will be used to draw parallels with existing results in temporal modal logic and products of modal logics [9].

This move between a branching time interpretation and a linear time interpretation can, in part, be justified by the following observation. Let $\mathcal{M}$ and $\mathcal{N}$ be two bisimular Kripke structures. Let w and v be two bisimular states in $\mathcal{M}$ and $\mathcal{N}$ respectively, and consider the tree unravelings of $\mathcal{M}$ and $\mathcal{N}$, denoted $Tree(w, \mathcal{M})$ and $Tree(v, \mathcal{N})$. Consider the temporal language with only linear time future looking operators (F and N). This language can be interpreted over the tree structures $Tree(w, \mathcal{M})$ and $Tree(v, \mathcal{N})$. Let H and H' be two paths in $Tree(w, \mathcal{M})$ and $Tree(v, \mathcal{N})$ respectively in which matching nodes at each stage along the paths are bisimular. It is not hard to show that for each temporal formula ϕ, ϕ is true at H, t iff ϕ is true at H', t.

An **epistemic temporal model** based on an epistemic temporal frame $\langle \mathcal{H}, \{\sim_i\}_{i \in \mathcal{A}} \rangle$ is a tuple $\langle \mathcal{H}, \{\sim_i\}_{i \in \mathcal{A}}, V \rangle$ where V is a valuation $V : \mathsf{At} \rightarrow 2^{\mathsf{FinPre}(\mathcal{H})}$. Formulas are interpreted at pairs H, t where $t \in \mathbb{N}$ and $H \in \mathcal{H}$ has length longer than t (typically we assume H is infinite). Truth for the languages $\mathcal{L}_n^X$ where X is a sequence of modalities is defined as usual. See [7] and [14] for details. We only remind the reader of the definition of the knowledge and some temporal operators:

- $H, t \models P\phi$ iff there exists $t' \leq t$ such that $H, t' \models \phi$
- $H, t \models F\phi$ iff there exists $t' \geq t$ such that $H, t' \models \phi$
- $H, t \models K_i\phi$ iff for each $H' \in \mathcal{H}$ and $m \geq 0$ if $H_t \sim_i H'_m$ then $H', m \models \phi$

We now turn to the $\mathcal{L}_\Sigma(\mathcal{A})$ language. This language is intended to be (strictly) stronger than the language described above. Before defining truth in a model we introduce a relation R_α on the set $\mathsf{FinPre}(\mathcal{H})$, where α is defined by the above grammar. Let H, H' be finite sequences of events and V a valuation (assigning sets of atomic propositions to finite sequences). Suppose $\sigma \in \Sigma$ and $i \in \mathcal{A}$.

- $HR_\sigma H'$ iff $H' = H\sigma$ if $\sigma \in \Sigma$
- $HR_i H'$ iff $H \sim_i H'$
- $HR_{\sigma^-} H'$ iff $\mathsf{len}(H) > 1$ and $H = H'\sigma$
- $HR_{?p} H'$ iff $H = H'$ and $p \in V(H)$.
- $HR_{\alpha;\beta} H'$ iff there exist H'' such that $HR_\alpha H''$ and $H'' R_\beta H'$
- $HR_{\alpha^*} H'$ iff for some n, $HR_\alpha^n H'$ (i.e. $R_{\alpha^*} = R_\alpha^*$)

Truth is defined as usual for atomic propositions and boolean connectives, we only give the definition of the modal operator:

- $H, t \models \langle\alpha\rangle\phi$ iff there exists $H' \in \mathcal{H}$ and $m \in \mathbb{N}$ such that $H_t R_\alpha H'_m$ and $H', m \models \phi$

In what follows, if H is a finite history, we may write $H \models \phi$ to mean $H', \mathsf{len}(H) \models \phi$ where H' is *any* infinite extension of H.

3 Living at the Edge

Having set up our basic framework, we now want to demonstrate some key facts about the border line between decidable and undecidable epistemic temporal logics. We will mostly assume that we have tree models with finitely many events only, starting from a single root. Instead of setting up a huge grid of possible model classes and languages, which tends to make the total picture somewhat diffuse, we high-light a few major stages, including one new highly undecidable epistemic tree logic. The main line of our observations is not all that new by itself, but our presentation and variety of sources is. We can only sketch our proofs here: the technical Appendix has further details.

3.1 $\mathcal{L}_{ETL}$ over arbitrary models

First, consider our complete language on arbitrary epistemic tree models, without constraints on the interaction of epistemic relations and events.

THEOREM 6. *The satisfiability problem for the language $\mathcal{L}_{ETL}$ over arbitrary models is RE.*

Proof. This logic is the fusion of multi-agent epistemic S5 with common knowledge, plus some complete axiomatization of our basic temporal language, while the tree structure adds one more axiom:

$$\langle i\rangle\phi \to F\phi \vee P(H\bot \wedge F\phi)$$

This says that any epistemic alternative is reachable in the tree by going up, or by going down to the root (where $H\bot$ holds) and then moving up again. Here the operator 'F' refers to the branching future in the above sense. One can show that this logic is complete for multi-**S5** models without special conditions on event relations R_e. Such a model $\mathcal{M}$ can be unraveled in the standard modal style to an epistemic tree model, where epistemic relations between nodes (finite sequences of worlds in $\mathcal{M}$) are just copied from those in $\mathcal{M}$ for their last members. ■

The special axiom in the preceding formulation reflects the assumption that agents know the current protocol. One could also give up on this, allowing models with different trees, while agents are not sure which one they are in.

We do not know if this general logic is decidable, though we suspect that it is, by the general results on transfer of decidability for fusions of modal logics surveyed in Gabbay, Wolter et al. [9]. Of course, behaviour of specific agents will take place in models satisfying additional epistemic-temporal constraints. But before going there, let's first consider the purely temporal part of our language over a still more specific structure.

3.2 Purely temporal languages on standard tree models

The Rabin Tree ([27]) consists, in our terms, of all finite sequences of events from some given finite set, with the binary relation of 'initial subsequence' plus successor functions taking a sequence H to He, for each $e \in \Sigma$.

THEOREM 7 (Rabin [24]). *The monadic second-order logic of the Rabin Tree is decidable.*

This landmark result explains the decidability of many modal and temporal logics, as first pointed out by Gabbay [8]. It applies particularly well to our setting here, since the Rabin Tree has both points and branches, represented as special sets of points.

THEOREM 8. *The satisfiability problem for $\mathcal{L}_{ETL}$ with respect to tree models is decidable.*

Proof. Validity for us is not on one model, but on all protocol models. However, this variety can be encoded as follows in a single Rabin tree. A formula ϕ involving finitely many events e is true in all protocol models if $\forall A(\text{'}subtree(A)\text{'} \Rightarrow (\phi)_A)$ is true on the Rabin Tree over that finite event set. Here $(\phi)_A$ is the obvious syntactic relativization of the formula ϕ to the unary predicate variable A, and '$subtree(A)$' expresses that A is closed under taking initial segments. ■

3.3 Idealized epistemic agents have a highly undecidable tree logic

Let us now consider the usual idealizations of epistemic logic. Agents have perfect memory, and seeing new events will not confuse them: that is, we have the above Perfect Recall, and No Learning properties. The resulting interaction of temporal and epistemic structure makes trees look more like grids, and indeed, undecidability strikes. We highlight this result, because it is indicative of the 'danger zone' that we are in. The results to follow are reminiscent of those by Halpern & Vardi [12], but our models and language are slightly different, while we also use a by now more standard tiling technique.

THEOREM 9. *The satisfiability problem of $\mathcal{L}_\Sigma(\mathcal{A})$ with respect to epistemic temporal frames that satisfy no learning is Σ^1_1-complete.*

Of course, this also implies that if both no-learning and perfect recall are satisfied, then satisfiability problem is Σ^1_1-complete. For perfect recall alone, we have the following two theorems.

THEOREM 10. *The satisfiability problem with respect to synchronous epistemic temporal frames that satisfy perfect recall is Σ^1_1-complete.*

THEOREM 11. *The satisfiability problem with respect to the language $\mathcal{L}_\Sigma(\mathcal{A})^-$ with respect to epistemic temporal frames that satisfy perfect recall is Σ^1_1 complete.*

We first remind the reader of a few relevant facts about tiling arguments. Let $\mathcal{T}$ be a finite set of tile types and for $T \in \mathcal{T}$, let $right(T)$, $left(T)$, $up(T)$ and $down(T)$ be the colors of T. The tiling problem (for the first quadrant) asks is there a function $t : \mathbb{N} \times \mathbb{N} \to \mathcal{T}$ such that for each $n, m \in \mathbb{N}$

$$\begin{aligned} right(t(n,m)) &= left(t(n+1,m)) \\ up(t(n,m)) &= down(t(n,m+1)) \end{aligned}$$

The function t is called a tiling of $\mathbb{N} \times \mathbb{N}$. The recurrent tiling problem asks, given a set of tiles $\mathcal{T}$ with a distinguished tile $T_1 \in \mathcal{T}$, if there is a tiling t such that T_1 occurs infinitely often in the first row. We will make use of the following Theorem of Harel.

THEOREM 12 ([13]). *The recurrent tiling problem is Σ^1_1-complete.*

It turns out that for frames satisfying perfect recall alone, we must appeal to the following Lemma. This Lemma states that if we can tile every finite square of the plane then we can tile the entire plane. We first need some notation. By a $(n \times n)$**-tiling** of the plane, we mean a function $t^{(n)} : \{(i,j) \mid 0 \le i \le n,\ 0 \le j \le n\} \to \mathcal{T}$ that satisfies the conditions described above (i.e., the tiles match vertically and horizontally). We say two tilings $t^{(n)}$ and $t^{(m)}$ are **consistent** if one tiling extends the other tiling. Thus each $(n \times n)$-tiling can be thought of as a sequence of partial *consistent* tilings.

LEMMA 13. *Suppose that for each $n > 0$, there is at least one (but only finitely many) partial tilings $t^{(n)}$. Then there is a tiling of the entire plane.*

The proof uses König's Lemma and can be found in the appendix.

We restrict attention to the class of models to generated by two events l and r. That is, suppose that $\Sigma = \{l, r\}$ and consider epistemic temporal frames whose protocol is a subset of $\Sigma^* \cup \Sigma^\omega$. The proof of the theorems above proceeds by encoding a recurrent tiling problem using the language $\mathcal{L}_\Sigma(\mathcal{A})$. Let $\mathcal{T} = \{T_1, \ldots, T_k\}$ be a finite collection of tiles and let $t_1, \ldots, t_k$ be propositional letters corresponding to the k different tile types. Our goal is to find a formula $\phi_\mathcal{T}$ that is satisfiable iff there is a recurrent tiling of $\mathbb{N} \times \mathbb{N}$ using the tiles from $\mathcal{T}$. We begin by describing the formula $\phi_\mathcal{T}$. The details can be found in the appendix. The formula $\phi_\mathcal{T}$ consists of three parts: 1. a formula which forces the extensions of a finite history to have a particular structure, 2. a formula which forces a grid structure and 3. a formula which places tiles on the grid.

That is, we define a formula $\phi_\mathcal{T}$ such that if $H, t \models \phi_\mathcal{T}$ then the extensions of H_t can be pictured as follows (details are found in the appendix).

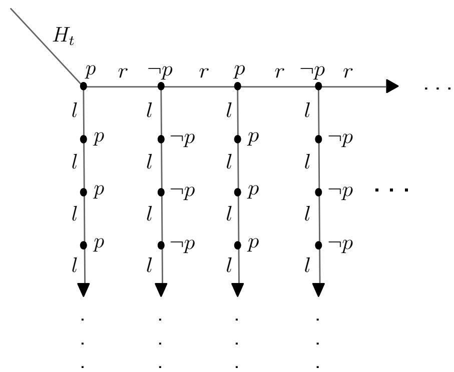

We can think of the above model as representing half of the $\mathbb{N}\times\mathbb{N}$ grid. The idea is to think of the infinite r-path as the y-axis and the first infinite l-path as the x-axis (the fact that the truth value of p alternates between the paths will be used below). We now show how force the second half of the grid. That is, we need a formula that will be satisfied if there are infinitely many infinite "up" paths. The trick will be to consider the following program

$$\alpha_u := ?p; l; 1; ?\neg p; l; 2; ?p$$

Making a step of the above program corresponds to making a 'zig-zag' move through the tree between points which will cross between different branches.

3.4 Decidable versions without interaction of epistemic and temporal parts

The preceding results show where we have oversteppend the bounds of decidability on tree models with additional epistemic relations. However, we can step back a little, and look at a slightly weaker language where we have common knowledge in the epistemic component, and its analogues of P and F in the temporal side. What we do not allow, however, in this language is the sort of mixing of temporal and epistemic steps that was crucial to the encoding of tiling cf. the above zig-zag move of α_u.

It might seem that this does not help, as epistemic relations of this neat sort behave like relations of simultaneity. And Läuchli proved that the first-order theory of the Rabin tree expanded with a binary 'equilevel' predicate for nodes is undecidable. But Thomas [27] provides a more fine-grained perspective: he shows that the monadic second-order theory of the Rabin Tree with an 'equilevel' predicate remains decidable provided that we let the second-order quantification run over linear chains, rather than arbitrary subsets. More succinctly: 'Path Logic' over the Rabin Tree with an equilevel predicate is decidable. Note that Path Logic extends our temporal languages, since these talk about initial segments and extensions of the current finite history.

THEOREM 14. *The logic of the language $\mathcal{L}_C^{P,F}$ over synchronous tree models is decidable.*

Proof. Thomas shows how Path Logic can be embedded into the monadic second-order theory of the Rabin tree, by sending chains to pairs of subsets (A, B) where A encodes the left-most branch on which the chain lies, while B encodes which nodes are on the chain. More precisely, B 'goes left' at levels not represented on the chain, and it 'goes right' at levels where the chain has a node. The equilevel predicate for two nodes is then expressed by saying that they are one-element chains, whose B-sequences go right at the same place. Now, our epistemic relations are subrelations of the 'equilevel' predicate , which latter corresponds roughly to the transitive closure of their union. But we can encode this epistemic language into that of Path Logic by modifying the chain representation. ■

3.5 Epistemic temporal logics over non-tree models

Against this background, here is how one can think of earlier work on epistemic-temporal models. Halpern & Vardi consider models where either the initial model may be infinite, or there may be infinite branching. In particular, in this case, even the 'unmixed' language of 3.4 above leads to undecidability with the assumptions of No Learning or Perfect Recall, because the grid structure now arises as follows. One starts from an infinite model in one direction, and lets the temporal axis form the other one. Note that this was not available to us in the above, since the levels of our event trees were finite. In addition, Halpern & Vardi have many further relevant results, which we do not survey here.

This concludes our statement of some typical results on decidability and undecidability over epistemic temporal tree models. Not surprisingly, the boundary has to do with the transition from mere trees to grid encoding using the additional epistemic structure. The epistemic setting adds some special flavor, however, in that the small differences which affect the complexity represent very concrete assumptions about agents' capabilities, and what we can say about these. Moreover, we have shown how one can learn about relevant results here from traditions that look prima facie quite different: earlier work on epistemic temporal logic, tree languages in the foundations of computation, and current work on products of modal logics.

4 Dynamic Epistemic Logic

Our take on epistemic temporal logic is mostly within the broad tradition of Fagin et al. [7] and Parikh & Ramanujam [20, 21]. One current family of logics which diverges slightly from these, though still clearly within the same spirit, is that of 'dynamic epistemic logic' ('DEL', [10, 22, 2]). Here, epistemic actions are described per se, such as announcing a true proposition, or performing some more complex type of communication. These are encoded explicitly in 'action models' $\mathbb{A}$ consisting of the relevant events and the preconditions for their occurrence, plus agents' epistemic relations over these, representing their partial powers of observation. In the strict version of DEL, preconditions for actions are defined by purely epistemic formulas. What agents learn from such epistemic events, given some current epistemic model $\mathbb{M}$, is encoded by a new 'product model' $\mathbb{M} \times \mathbb{A}$, where agents are uncertain between worlds (s, e) and (t, f) iff they were uncertain between both the old worlds s, t and the observed events e, f. The language for these models has the usual epistemic operators, plus modalities $\langle \mathbb{A}, e \rangle \phi$ interpreted as follows:

$$\mathbb{M}, s \models \langle \mathbb{A}, e \rangle \phi \ \text{ iff } \ \mathbb{M} \times \mathbb{A}, (s, e) \models \phi$$

The resulting logic is decidable, and it revolves around so-called 'reduction axioms' allowing for compositional analysis of the effects of epistemic axioms. E.g., a typical reduction axiom is that for agents' knowledge after a public announcement:

$$[!P]K_i\phi \leftrightarrow (P \rightarrow K_i[!P]\phi)$$

We refer to the literature for more precise definitions and more elaborate results on complete dynamic-epistemic logics (cf. van Benthem [29], Baltag & Moss [3], van Benthem [30], van Eijck & Kooi [32], and van Ditmarsch, van der Hoek & Kooi [33]).

Prima facie, DEL does not look like epistemic-temporal logic at all: since there is no explicit mention of time. And vice versa, ETL does not look much like dynamic-epistemic logic, since it does not explicitly describe the epistemic events that lead to the construction of successive models by agents. What we want to illustrate now, as a sample of our 'convergence' view on the whole area of epistemic temporal reasoning, is that these appearances are misleading. The two approaches have much to offer to one another, precisely, because they are so close that borrowing is easy and natural. We develop this theme in a number of separate topics.

4.1 Representing DEL models inside ETL models

Product update involves three major ingredients with a logical 'reflection', as was first observed in van Benthem 2001:

(a) Product update implies Perfect Recall: , $(x, a) \sim_i (y, a)$ implies $x \sim_i y$

(b) Product also propagates uncertainty (there are 'No Miracles'): if $x \sim_i y$, then after performance of a in both cases, $(x, a) \sim_i (y, a)$

(c) Moreover, 'no miracles' holds uniformly: Actions are either always distinguishable, or never: if $(x, a) \sim_i (y, b)$, then, whenever $u \sim_i v$, also $(u, a) \sim_i (v, b)$ if the latter moves can be performed at all.

These principles can be translated into axioms of our epistemic temporal logic in a straightforward manner. But here, we are rather concerned in the immediate semantic connection. Given any initial model $\mathbb{M}$ and action model $\mathbb{A}$ representing all possible events plus agents' observational powers over these, we can form a natural epistemic tree model $Tree(\mathbb{M}, \mathbb{A})$ as follows. Nodes are finite sequences of events, and the successive epistemic models of the DEL-style process are the horizontal levels of the tree. Events only take place when their precondition is satisfied. Note that this is a tree model whose initial epistemic model at the root can be arbitrary: there is no special restriction to finiteness (though van Benthem [30], Sadzik [26] do explore this special case, looking for epistemic bisimulations between different finite levels).

In particular, the epistemic decoration of the tree models $Tree(\mathbb{M}, \mathbb{A})$ is rather special, since it obeys the above three constraints. Indeed, van Benthem & Liu [31] prove the following representation result:

THEOREM 15. *An epistemic tree model $\mathcal{M}$ is bisimilar to a model of the form $Tree(\mathbb{M}, \mathbb{A})$ if and only if it satisfies (a) Perfect Recall, (b) Uniform No Miracles, and (c) for any event e, the set of nodes where e can take place is closed under epistemic bisimulations inside $\mathbb{M}$.*

Thus, product update corresponds to a special epistemic temporal logic. As stated before, we can unpack the content of the conditions to the usual axioms for Perfect Recall ($K_i[a]\phi \to [a]K_i\phi$) and No Learning, where the uniform version of the latter would require the use of universal modalities $E\phi$, $U\phi$ stating that ϕ holds at some world, at all worlds, resp. (cf. Blackburn, de Rijke & Venema [5]):

$$E(\langle a^- \rangle \top \wedge \neg K_i \neg \langle b^- \rangle \top) \to U(\langle a^- \rangle \neg K_i \neg \phi \to K_i [b^-]\phi)$$

Thus, in a sense, dynamic-epistemic logic describes the behaviour of special idealized agents on epistemic temporal trees, much as we discussed in preceding sections. But this is not all to the story, since DEL also has some further special features. In particular, in the light of Section 3, its decidability calls for explanation!

More can be said here. E.g., 'No Miracles' is a much more plausible version of propagating ignorance than the usual formulations of No Learning, which seem to say that the passage of time never helps increase knowledge. This improvement reflects the fact that DEL gives a deeper analysis of the processes that drive information change, instead of merely describing the Grand Stage where all agents live in time.

4.2 Why is DEL decidable?

Using the above observations, one can analyze the principles of DEL-style calculi in our epistemic temporal logic. But when we do, we end up in a very restricted fragment of our languages:

THEOREM 16. *Over models of the form $Tree(\mathbb{M}, \mathbb{A})$, DEL corresponds to the epistemic-temporal language of epistemic logic plus one-step future operators $\langle e \rangle$.*

Thus, we do not reach the expressive power needed for undecidability in our earlier arguments. For DEL proper, this can be seen as follows. The reduction axioms of the standard calculi translate every formula into an equivalent one without action modalities, and one then uses the decidability of the purely epistemic language. This is a sort of 'one-dimensional reduction' of a two-dimensional system. But actually, something stronger holds. Suppose that events can have preconditions in the epistemic temporal language, as happens in many concrete communications scenarios which refer, e.g. to the past of the current conversation. In its most blunt form, the precondition for e to occur is then just this: $\langle e \rangle \top$.

THEOREM 17. *The logic of the epistemic temporal language with only operators $\langle e \rangle$ on models with Perfect Recall and No Miracles is decidable.*

Proof. This can be derived from the decidability result for the modal 'product logic' $\mathbf{PDL} \times \mathbf{K}_m$ in Gabbay, Wolter et al. [9], by embedding our language into it. The models for this logic may be viewed as grids with an 'epistemic' **PDL** direction and a 'temporal' one-step **K**-direction. Still, no embedding of tiling problems is possible, because the language does not contain a true universal modality or transitive closure modality accessing all points of the grid (cf. also Marx & Mikulas [16]). ■

4.3 Program structure, true future, and undecidability

DEL style logics do become undecidable when the complete future is added. This would happen, e.g., if one adds sequential structure to action models, modeling, say, conversational processes involving composition and iteration. The landmark paper Miller & Moss [18] proves many relevant results, including this surprising effect of combining two decidable logics:

THEOREM 18 (Miller & Moss). *The dynamic-epistemic logic of public announcement with program iterations is undecidable.*

This shows that the undecidability phenomena already noted in Halpern & Vardi for the language with common knowledge and true future even occur in very restricted settings, where events are just announcements. Indeed, Miller & Moss show that iterated announcement of one single proposition $\Diamond \top$ suffices. However, in the light of our ETL-based Section 3, their analysis also leaves open questions. One of them has to do, again, with our assumption of finite levels. It is unknown whether their undecidability results hold when initial models are assumed to be finite.

4.4 Decidable fragments; tomorrow and yesterday

Our epistemic temporal analysis also suggests various additions to the standard language of DEL which still remain decidable. One typical illustration is the addition of the temporal past of an epistemic process (van Benthem [30]). Yap [34] and Sack [25] analyze such additions, and propose valid axioms. Here is the general setting.

THEOREM 19. *The logic of the epistemic temporal language with one-step future operators $\langle e \rangle$ and one-step past operators $\langle e^{-} \rangle$ on models with Perfect Recall and No Miracles is decidable.*

Proof. This follows from a simple modification of the above-mentioned decidability proof for $\mathbf{PDL} \times \mathbf{K}$ given in Gabbay, Wolter et al. [9]. ■

We conjecture that decidability still holds when we add many-step Past operators, at least, on our rooted finite-event trees. Logics such as this can discuss what agents knew at the previous stage, but they can also include preconditions for events that reach back in time, such as "saying P if you have not already done so".

4.5 Protocols and model constructions

But there are still further features to the comparison of DEL and ETL. First, it is sometimes claimed that DEL lacks an essential resource available in our epistemic temporal models, viz. the choice of a *protocol*, i.e., a set of 'relevant histories'. Now, this is not true, since the above models $Tree(\mathbb{M}, \mathbb{A})$ do have explicit restrictions on their available runs, since events can only occur when their precondition is satisfied. Thus, DEL has an explicit calculus for preconditions, as these are encoded in the action models, and these are again available inside the formal language through the modalities $\langle \mathbb{A}, e\rangle\phi$. On the other hand, given the special epistemic format for these preconditions, one can only define special protocols, via local restrictions, that must be stated in a purely epistemic language. A more general approach here would merge the two ideas. On the one hand, it seems a good idea to make the protocols explicit in the language, as DEL does. On the other hand, one needs a richer repertoire of definitions for realistic protocols, including temporal operators in their formulation. This can be achieved in the following 'Logic of Protocols':

The language is similar to our $\mathcal{L}_\Sigma(\mathcal{A})$ language. We first introduce a **PDL** style language for representing protocols in the object language. A protocol can have the following syntactic form

$$e \mid \alpha;\beta \mid \alpha \cup \beta \mid \alpha^* \mid \mathsf{skip} \mid \phi?$$

where $e \in \Sigma$ is an event and ϕ is a formula of $\mathcal{L}_{ETL}$. For example, the protocol $(e \cup \mathsf{skip})^*$ represents the set of histories that contain the event e. A test-free protocol α is simply a regular expression; and so it represents the set of histories that match α. Another more interesting example is a 'Liar Protocol'. Let $\mathsf{send}(i, B, p)$ be the event agent i sends the message p to the group of agents B, i.e., i announces p to the set of agents B. Then $((K_i\neg p?;\mathsf{send}(i, \mathcal{A}, p)) \cup (K_i p?;\mathsf{send}(i, \mathcal{A}, \neg p)) \cup \mathsf{skip})^*$ represents a liar protocol. That is, if i knows p then i publically announces $\neg p$, if i knows $\neg p$ then i publically announces p, or i does not say anything.

For each protocol α introduce a modal operator N_α to the language. The intended interpretation of $N_\alpha\phi$ is that ϕ is true at the next moment in *all* extensions of the current history compatible with the protocol α. Thus truth is defined as $H, t \models N_\alpha\phi$ if $H' \in \mathcal{H}, H \preceq_\alpha H'$ and $H', t \models \phi$, where $\preceq_\alpha$ is an extension relation much like the previously define R_α relations. This addition, though very useful in practice, is arguably a matter of convenience:

THEOREM 20. *ETL with explicit protocols is no more expressive than ETL by an effective translation.*

Introducing explicit protocols is also akin to the use of 'knowledge programs' in Fagin et al. [7]. We forgo the precise connection here.

Our conclusion is that older and newer approaches to dynamic actions and epistemic logic all meet in the same arena of epistemic logic, and that insights can be transferred in illuminating ways.

5 Conclusions

This paper has tried to show that epistemic temporal models are a natural meeting place for logical studies from many different directions. In Section 2, we have defined some basic structures that seem to recur in most major studies of agents' interaction and information. In Section 3, we discussed the decidable/undecidable boundary where many of the interesting issues live concerning agent behaviour. We found that these issues also lead to a natural combination of insights from a number of different traditions: epistemic temporal logics in computers science, but also logics of computation, modal logics of products, and dynamic-epistemic logics.

Concerning the relation between these frameworks, our view is this. Epistemic logics in the style of Fagin et al, and Parikh et al. are largely the same, even up to mutual mathematical representation (cf. [19]). The situation with these logics vis-à-vis dynamic-epistemic logics is a bit more complex, but Section 4 has shown some natural merges that combine ideas from both sides. In particular, the tendency one sometime finds to play up differences between these approaches as different 'world views' seems both pointless from a mathematical viewpoint, and harmful from a conceptual or a practical point of view, as it impedes mutual flow of ideas.

Indeed, many further examples of such traffic can be found, which we had to leave out for reasons of space. E.g., we also have new results (not reported on here) on less-than-ideal agents with *bounded memory* in the same combined setting. In particular, one can show that the epistemic-temporal logic of memory-free agents reduces to a decidable purely temporal logic over trees (cf. van Benthem [29], van Benthem & Liu [31]): this is the opposite, so to speak, of the reduction found in dynamic-epistemic logic. But perhaps the more exciting perspective is when bounded agents and idea agents meet. In that setting, speaking about the knowledge of bounded agents introduces global universal and existential modalities, which may interact with the one-step future logic for the ideal ones. We conjecture that the resulting logics can get undecidable again, reflecting the difficulties of the interplay in a society with a diversity of agents. But there are many other possible confluences. E.g., the explicit treatment of model constructions in dynamic-epistemic logic is also reminiscent of that in *process algebra*, and again, in this way, one more connection would be made in the broad area of process theories for agents that can know things and display rational intelligent behavior.

Still more broadly, our approach would suggest a turn in perspective, from competition between frameworks to cooperation. Compare the situation in the 1930s, when many different models were proposed for computation. Instead of creating different churches, logicians started looking for similarities and equivalences (at some appropriate level), and the result was Church's Thesis, usually taken to mean that the field had a stable and mathematically respectable topic. Likewise, convergence, if not downright equivalence, between epistemic temporal logics of agents might signal to a broader world that there is a core notion of genuine interest here concerning 'intelligent interaction', rather than a set of warring religions. Seeing the differences may make for short-term gains in terms of project funds and reputation, seeing the analogies leads to the long-term gain of a common cause.

BIBLIOGRAPHY

[1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 2002.

[2] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In *Proceedings of TARK 1998*, 1998.

[3] Alexandrau Baltag and Larry Moss. Logics for epistemic programs. *Synthese: Knowledge, Rationality, and Action*, 2:165 – 224, 2004.

[4] N. Belnap, M. Perloff, and M. Xu. *Facing the Future*. Oxford University Press, 2001.

[5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, Cambridge, 2002.

[6] E. A. Emerson. Temporal and modal logics. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*. 1990.

[7] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. The MIT Press, Boston, 1995.

[8] D. Gabbay. *Investigations into modal and tense logics, with applications to problems in linguistics and philosophy*. Reidel Dordrecth, 1976.

[9] D.M. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyaschev. *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier, 2003.

[10] J. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, ILLC, 1998.

[11] V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, to appear 2006.

[12] Joseph Halpern and Moshe Vardi. The complexity of reasoning about knowledge and time. *J. Computer and System Sciences*, 38:195 – 237, 1989.

[13] David Harel. Recurring dominoes: Making the highly undecidable highly understandable. *Annals fo Discrete Mathematics*, 24:51 — 72, 1985.

[14] I. Hodkinson and M. Reynolds. Temporal logic. In *Handbook of Modal Logic*. forthcoming.

[15] J. Horty. *Agency and Deontic Logic*. Oxford University Press, 2001.

[16] M. Marx and S. Mikulas. Products, or how to create modal logics of high complexity. *Logic Journal of the IGPL*, pages 77 –88, 2001.

[17] M. Marx and Y. Venema. Local variations on a loose theme: modal logic and decidability. In M. Vardi and S. Weinstein, editors, *Finite-Model Theory and Its Applications*, Texts in Theoretical Computer Science. Springer, 2005.

[18] J. Miller and L. Moss. The undecidability of iterated modal relativization. *Studia Logica*, 79(3), 2005.

[19] E. Pacuit. Some comments on history based structures. *Journal of Applied Logic*, forthcoming.

[20] Rohit Parikh and R. Ramanujam. Distributed processes and the logic of knowledge. In *Logic of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 256 – 268. Springer, 1985.

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

[22] J. Plaza. Logics of public communications. In *Proceedings, 4th International Symposium on Methodolgies for Intelligent Systems*, 1989.
[23] A. Pnueli. The temporal logic of programs. In *Proc. 18th Symp. Foundations of Computer Science*, pages 46 — 57, 1977.
[24] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1 – 35, 1969.
[25] J. Sack. Temporal language for epistemic programs. unpublished manuscript, Indiana at Bloomington.
[26] T. Sadzick. Exploring the update universe. unpublished manuscript, Stanford University.
[27] W. Thomas. Infinite trees and automaton definable relations over omega-words. *Theoretical Computer Science*, 103(1):143 – 159, 1992.
[28] Richmond Thomason and Anil Gupta. A theory of conditionals in the context of branching time. *The Philosophical Review*, 80:65–90, 1980.
[29] J. van Benthem. Games in dynamic epistemic logic. *Bulletin of Economic Research*, 53:216 – 248, 2001.
[30] J. van Benthem. One is a lonely number. Technical report, ILLC, University of Amsterdam, 2002.
[31] J. van Benthem and F. Liu. Diversity of logical agents in games. *Philosophia Scietiae*, 8(2):163 – 178, 2004.
[32] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. In *Proceedings of TARK 2005*, 2005.
[33] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, forthcoming.
[34] A. Yap. Product update and looking backward. unpublished manuscript, Stanford University.

Appendix

A Technical Proofs for Section 3.3

Definition of $\phi_{\mathcal{T}}$: Let ϕ_S be the conjunction of the following formulas.

1. Only $r^* - l^*$ paths: $[r^*; l; l^*]\neg\langle r\rangle\top$
2. Infinitely many infinite l-paths: $[r^*; l^*]\langle l\rangle\top$
3. Infinite r-path: $[r^*]\langle r\rangle\top$
4. Even p paths: $[(r;r)^*][l^*]p$
5. Odd $\neg p$ paths: $[r;(r;r)^*][l^*]\neg p$

Let ϕ_u be the conjunction of the following two formulas

1. "Up moves" are always possible: $[r^*; l^*]((p \to \langle 1\rangle\neg p) \wedge (\neg p \to \langle 2\rangle p))$
2. There are infinitely many "Up moves": $[l^*][\alpha_u^*]\langle\alpha_u\rangle\top$

Finally, we place the tiles on our tree.

1. Exactly one tile is placed at each node: $\phi_1 := [(r;r)^*; l^*](\bigvee_{i=1}^{k} t_i \wedge \bigwedge_{1\leq i<j\leq k} \neg(t_i \wedge t_j))$
2. Place tiles going across: $\phi_2 := [(r;r)^*; l^*](\bigvee_{right(T_i)=left(T_j)}(t_i \wedge \langle l\rangle t_j))$

3. Place tiles going up: $\phi_3 := [(r;r)^*; l^*](\bigvee_{up(T_i)=down(T_j)}(t_i \wedge \langle \alpha_u \rangle t_j))$

4. And only enough tiles are placed: $\phi_4 := [(r;r)^*; l^*](\bigwedge_{up(T_i)\neq down(T_j)}(t_i \rightarrow \neg\langle \alpha_u \rangle t_j))$

Let $\phi_{\mathcal{T}} := \phi_S \wedge \phi_u \wedge \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4$.

We first note that given a tiling of $\mathbb{N} \times \mathbb{N}$, we can build a satisfying model of $\phi_{\mathcal{T}}$. In the interest of space we will not give the full proof. The key idea is to remove all vertical lines from the grid and tree the remaining structure as a tree rooted at $(0,0)$.

LEMMA 21. *Suppose that t is a tiling of $\mathbb{N} \times \mathbb{N}$ using tiles from $\mathcal{T}$. Then $\phi_{\mathcal{T}}$ is satisfiable.*

Before proceeding to the proof of Theorem 9 we state some facts which are immediate consequences of the definition of truth in a model.

LEMMA 22. *Suppose that t_i and t_j are tiling propositions and $H, t \models \phi_{\mathcal{T}}$. For for each $H', H'' \in \mathsf{Ext}^{<\omega}(H_t)$ if $H' \models t_i$, $H'' \models t_j$ and $H' R_{\alpha_u} H''$ then $up(t_i) = down(t_j)$.*

LEMMA 23. *Suppose that t_i and t_j are tiling propositions and $H, t \models \phi_{\mathcal{T}}$. For for each $H', H'' \in \mathsf{Ext}^{<\omega}(H_t)$ if $H' \models t_i$, $H'' \models t_j$ and $H' R_l H''$ then $right(t_i) = left(t_j)$.*

LEMMA 24. *Suppose that the epistemic temporal frame $\langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ satisfies the no learning property. Then for all finite histories $H, H' \in \mathcal{H}$, if $H R_{\alpha_u} H'$ then $Hl R_{\alpha_u} H'l$.*

LEMMA 25. *Suppose that $\mathcal{F} = \langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ satisfies perfect recall, then for each $H, H' \in \mathcal{H}$ such that $Hl, H'l \in \mathcal{H}$, if $Hl R_{\alpha_u} H'l$ then $H R_{\alpha_u} H'$.*

Proof of Theorem 9 Our goal is to show, under the assumption of no learning, that we can generate a tiling of the plane from satisfying model of $\phi_{\mathcal{T}}$. The strategy is to show that there is a function f from $\mathbb{N} \times \mathbb{N}$ into $\mathsf{Ext}^{<\omega}(H_t)$ such that the following function

- $t(n,m) = T_i$ iff$_{\text{def}}$ $f(n,m) \models t_i$.

is a tiling of $\mathbb{N} \times \mathbb{N}$. We will show that such a function can be extracted from a satisfying model. Note that we are not providing an explicit definition of the function, but rather proving that such a function exists. First of all, since $H, t \models \phi_1$, f is well-defined. Start by defining $f(0,0) = T$ where $H, t \models t$. We now show that $f(n,m)$ has been defined then we can define both $f(n+1,m)$ and $f(n,m+1)$.

LEMMA 26. *Suppose that $H, t \models \phi_{\mathcal{T}}$, $(n,m) \in \mathbb{N} \times \mathbb{N}$, $f(n,m) = H_t H'$ and $H_t H' \models t_i$. Then there are finite histories H^r and H^u such that*

1. *$H_t H^r, H_t H^u \in \mathsf{Ext}^{<\omega}(H_t)$ and $H^r, H^u \in \mathsf{Lang}((rr)^\star \cdot l^\star)$,*

2. *$H_t H^r \models t_j$ where $right(T_i) = left(T_j)$, and*

3. *$H_tH^u \models t_k$ where $up(T_i) = down(T_k)$.*

Proof. Suppose that $H, t \models \phi_{\mathcal{T}}$ and let $(n, m) \in \mathbb{N}\times\mathbb{N}$ with $f(n, m) = H_tH'$ and $H_tH' \models t_i$. First of all, note that by the definition of truth in a model, H' is of the form an even number of rs followed by some number of ls. Since $H, t \models \phi_2$ there is a t_j such that $H_tH'l \models t_j$. Furthermore, this t_j is unique because of the structure of the frame. Hence $right(T_i) = T_j$. Thus the history $H^r = H'l$ satisfies the appropriate properties. The proof of the existence of H^u is analogous. Since $H, t \models \phi_3$ and $H_tH' \models t_i$, there is some H'' such that $H_tH'R_{\alpha_u}H_tH'H''$ and $H_tH'H'' \models t_k$ with $up(T_i) = down(T_k)$. Then $H^u = H''$ satisfies the appropriate properties. ■

Thus if $f(n, m)$ is defined, we can define $f(n + 1, m)$ and $f(n, m + 1)$. All that remains is to show that we can "complete the square". Suppose that $f(n, m)$, $f(n, m+1)$ and $f(n+1, m)$ have all been assigned finite extensions of H_t as described in the Lemma 26. Call these histories $H_{(n,m)}$, $H_{(n,m+1)}$ and $H_{(n+1,m)}$ respectively, where

- $t(n, m) = T_i$, $t(n, m + 1) = T_j$ and $t(n + 1, m) = T_k$; and
- $up(t(n, m)) = down(t(n, m + 1))$ and $right(t(n, m)) = left(t(n + 1, m))$; and

Then we must show that there is a finite extension $H_{(n+1,m+1)}$ of H_t that has the required properties. That is such that $H' \models t_l$ where

LEMMA 27. *Suppose that $H_{(n,m)}, H_{(n,m+1)}$ and $H_{(n+1,m)}$ have been defined as above. Then there is a finite history $H_{(n+1,m+1)} \in \mathsf{Ext}^{<\omega}(H_t)$ such that there is a unique tile proposition with $H_{(n+1,m+1)} \models t$ and*

1. *$right(T_j) = left(T)$ and*
2. *$up(T_k) = down(T)$.*

Proof. Let $H_{(n,m)}, H_{(n+1,m)}$ and $H_{(n,m+1)}$ be defined as above. And define $H_{(n+1,m+1)} = H_{(n,m+1)}l$. First note that $H_{(n+1,m+1)}$ exists since $H, t \models \phi_S$ (specifically because of the second conjunct in ϕ_S). By Fact 2, $right(T_j) = left(T)$. Thus we need only show that $up(T_k) = down(T)$. Since $H_{(n,m)}R_{\alpha_u}H_{n,m+1}$, by Lemma 24, $H_{(n+1,m)} = H_{(n,m)}lR_{\alpha_u}H_{(n,m+1)}l = H_{(n+1,m+1)}$. Thus $H_{(n+1,m)} \models t_k \wedge \langle\alpha_u\langle t$. Hence by Fact 1, $up(T_k) = down(T)$. ■

Finally, we must show that there is a formula such that, if satisfiable, implies that a particular tile occurs infinitely often along the x-axis. Let $T_0 \in \mathcal{T}$ be a tile and t_0 the corresponding propositional variable. Consider the formula $\phi_{t_0} := [l^*]\langle l; l^*\rangle t_0$Note that $H, t \models \phi_{t_0}$ implies that t_0 is true infinitely often on the first branch extending H_t. This proves Theorem 9.

Proof of Lemma 13 Suppose that for each $n > 0$, there is at least one (but only finitely many) partial tilings $t^{(n)}$. We can use this fact to define a tree where the nodes of the tree are $(n \times n)$-tilings of the plane. The tree is defined as follows. For each $n \geq 0$, there is a finite set of $(n \times n)$-tilings $t_1^{(n)}, \ldots, t_k^{(n)}$. Let T_n be the tree with a root r (which is not a tiling) and nodes which are all the $(n' \times n')$-tilings for $0 \leq n' \leq n$. For each $1 \leq l \leq n$, put an edge from $t^{(l-1)}$ to $t^{(l)}$ if $t^{(l-1)}$ and $t^{(l)}$ are consistent. Finally put edges from the root node r to all $(0,0)$-tilings. Then it is easy to see that the tree $T = \cup_{n<\omega} T_n$ (where the union of two trees is defined to be the union of the nodes and the branches) is finitely branching and has infinitely many nodes. Thus by König's Lemma there is an infinite branch. This infinite branch is a tiling of the plane.

Proof of Theorem 10 and Theorem 11 In light of Lemma 13 we need only show that we tile every finite square of the plane.

DEFINITION 28. Call a sequence $H_1, \ldots, H_{n+1}$ of elements of $\mathsf{Ext}(H_t)$ a α_u **sequence of length** n if

1. $H_1 R_{\alpha_u} H_2 R_{\alpha_u} H_3 \cdots R_{\alpha_u} H_{n+1}$

If, in addition, $H_1 = H_t l^n$, then we call the sequence a n-α_u **sequence of length** n.

The idea is that given any such sequence of histories, we can read off a $(n \times n)$-tiling. That is we must define a function $t^{(n)} : \{(i,j) \mid 0 \leq i \leq n, 0 \leq j \leq n\} \to \mathcal{T}$ that satisfies the tiling properties. We begin by defining for each $i = 1, \ldots, n+1$, $t^{(n)}(n, i-1) = T$ iff $H_i \models t$.

We first check that for $i = 0, \ldots, n-1$, $up(t^{(n)}(n,i)) = down(t^{(n)}(n, i+1))$. This is a consequence of Fact 1. Thus with any α_u sequence of n histories can be used to tile a vertical line of length n. We now need to show that from a n-α_u sequence of length n, we can find at least n α_u-sequences of length n. If we can find n such sequences, then we can tile the vertical lines in the $n \times n$ square. Say that two α_u sequences $H_1, \ldots, H_{n+1}$ and $H'_1, \ldots, H'_{n+1}$ are R_l**-connected** if for each $i = 1, \ldots n+1$, $H_i R_l H'_i$. By the Lemma **??**, the vertical lines generated from an n-α_u sequence will be R_l-connected. I.e., for $m = 1, \ldots, n$, the m-α_u sequence will be R_l-connected to the $(m+1)$-α_u sequence.

Let $H_1, \ldots, H_{n+1}$ be a n-α_u sequence of length n. By the definition of truth in a model, for each $j = 1, \ldots n+1$, there is a $k_j \geq 0$ such that $H_j = H' l^{k_j}$ and $H' \in \mathsf{Ext}^{<\omega}(H_t)$. For each $H \in \mathsf{Ext}^{<\omega} H_t$ let H^- be $H' l^m$ if $H = H' l^{m+1}$ and H otherwise. The perfect recall property ensures that if for each $j = 1, \ldots, n+1$, $k_j > 0$, then $H_1^-, \ldots H_{n+1}^-$ is a $(n-1)$-α_u sequence of length n. Thus if for each $j = 1 \ldots, n+1$, $k_j \geq n$, then we are done.
However, in general, there may be some $k_j < n$. Thus we must force the model to be such that for each $j = 1, \ldots, n+1$, $k_j \geq n$. First of all note that if we assume synchronicity then for each $j = 1, \ldots, n+1$, $k_j = n$. Hence we have proved Theorem 10. In fact, the synchronous property is stronger

than what is needed. Essentially, if we can force that if $H_t r^m l^n R_{\alpha_u} H_t H''$ then $H'' = r^{m'} l^{n'}$ where $n' \geq n$, then we can prove the result. It is easy to see that this property will hold for any extension if H_t if the following formulas is satisfied at H, t

$$\phi_l := [r^*; l^*](\langle l^- \rangle \top \rightarrow [\alpha_u]\langle l^- \rangle \top)$$

Note the use of the converse operator. Thus we have shown Theorem 11.

Johan van Benthem
Plantage Muidergracht 24
Institute of Logic, Language and Computation
University of Amsterdam
johan@science.uva.nl

Eric Pacuit
Plantage Muidergracht 24
Institute of Logic, Language and Computation
University of Amsterdam
epacuit@staff.science.uva.nl