

Multi-agent belief dynamics: bridges between dynamic doxastic and doxastic temporal logics

Johan van Benthem
ILLC Amsterdam
Stanford University
johan@science.uva.nl

Cédric Dégremon^{*}
ILLC Amsterdam
cdegremo@science.uva.nl

Abstract

In this paper, we compare two modal frameworks for multi-agent belief revision: dynamic doxastic logics computing stepwise updates and temporal doxastic logics describing global system evolutions, both based on plausibility pre-orders. We prove representation theorems showing under which conditions a doxastic temporal model can be represented as the stepwise evolution of a doxastic model under successive 'priority updates'. We define these properties in a suitable doxastic-temporal language, discuss their meaning, and raise some related definability issues.

Analyzing the behavior of agents in a dynamic environment requires describing the evolution of their knowledge as they receive new information. Moreover agents entertain beliefs that need to be revised after learning new facts. I might be confident that I will find the shop open, but once I found it closed, I should not crash but rather make a decision on the basis of new consistent beliefs. Such beliefs and information may concern ground-level facts, but also beliefs about other agents. I might be a priori confident that the price of my shares will rise, but if I learn that the market is rather pessimistic (say because the shares fell by 10%), this information should change my higher-order beliefs about what other agents believe.

Tools from modal logic have been successfully applied to analyze knowledge dynamics in multi-agent contexts. Among these, Temporal Epistemic Logic [23], [19]'s Interpreted Systems, and Dynamic Epistemic Logic [2] have been particularly fruitful. A recent line of research [11, 10, 9] compares these alternative frameworks, and [10] presents a representation theorem that shows under which conditions a temporal model can be represented as a dynamic one. Thanks to this link, the two languages also become comparable, and one can merge ideas: for example, a new line

of research explores the introduction of protocols into the logic of public announcements PAL, as a way of modeling informational processes (see [9]).

To the best of our knowledge, there are no similar results yet for multi-agent belief revision. One reason is that dynamic logics of belief revision have only been well-understood recently. But right now, there is work on both dynamic doxastic logics [5, 3] and on temporal frameworks for belief revision, with [14] as a recent example. The exact connection between these two frameworks is not quite like the case of epistemic update. In this paper we make things clear, by viewing belief revision as priority update over plausibility *pre-orders*. This correspondence allows for similar language links as in the knowledge case, with similar precise benefits.

We start in the next section with background about earlier results and basic terminology. In section 2 we give the main new definitions needed in the paper. Section 3 presents the key temporal doxastic properties that we will work with. In section 4 we state and prove our main result linking the temporal and the dynamic frameworks, first for the special case of *total* pre-orders and then in general. We also discuss some variations and extensions. In section 5 we introduce formal languages, providing an axiomatization for our crucial properties, and discussing some related definability issues. We state our conclusions and mention some further applications and open problems in the last section.

1 Introduction: background results

Epistemic temporal trees and dynamic epistemic logics with product update are complementary ways of looking at multi-agent information flow. Representation theorems linking both approaches were proposed for the first time in [6]. A nice presentation of these early results can be found in [21, ch5]. We start with one recent version from [9], referring the reader to that paper for a proof, as well as generalizations and variations.

^{*}The second author was supported by a GLoRiClass fellowship of the European Commission (Research Training Fellowship MEST-CT-2005-020841).

Definition 1.1 [Epistemic Models, Event Models and Product Update]

- An *epistemic model* \mathcal{M} is of the form $\langle W, (\sim_i)_{i \in N}, V \rangle$ where $W \neq \emptyset$, for each $i \in N$, \sim_i is a relation on W , and $V : Prop \rightarrow \wp(H)$.
- An *event model* $\epsilon = \langle E, (\sim_i)_{i \in N}, pre \rangle$ has $E \neq \emptyset$, and for each $i \in N$, \sim_i is a relation on W . Finally, there is a precondition map $pre : E \rightarrow \mathcal{L}_{EL}$, where \mathcal{L}_{EL} is the usual language of epistemic logic.
- The *product update* $\mathcal{M} \otimes \epsilon$ of an epistemic model $\mathcal{M} = \langle W, (\sim_i)_{i \in N}, V \rangle$ with an event model ϵ is the model $\langle E, (\sim_i)_{i \in N}, pre \rangle$, whose worlds are pairs (w, e) with the world w satisfying the precondition of the event e , and accessibilities defined as:

$$(w, e) \sim'_i (w', e') \text{ iff } e \sim_i e', w \sim_i w'$$

◁

Intuitively epistemic models describe what agents currently know while the product update describe the new multi-agent epistemic situation after some epistemic event has taken place. Nice intuitive examples are in [1].

Next we turn to the epistemic temporal models introduced by [23]. In what follows, Σ^* is the set of finite sequences on any set Σ , which forms a branching ‘tree’.

Definition 1.2 [Epistemic Temporal Models] An *epistemic temporal model* (ETL model for short) \mathcal{H} is of the form $\langle \Sigma, H, (\sim_i)_{i \in N}, V \rangle$ where Σ is a finite set of events, $H \subseteq \Sigma^*$ and H is closed under non-empty prefixes. For each $i \in N$, \sim_i is a relation on H , and there is a valuation $V : Prop \rightarrow \wp(H)$. ◁

The following epistemic temporal properties drive [9]’s main theorem.

Definition 1.3 Let $\mathcal{H} = \langle \Sigma, H, (\sim_i)_{i \in N}, V \rangle$ be an ETL model. \mathcal{H} satisfies:

- **Propositional stability** if, whenever h is a finite prefix of h' , then h and h' satisfies the same proposition letters.
- **Synchronicity** if, whenever $h \sim h'$, we have $len(h) = len(h')$.

Let \sim^* be the reflexive transitive closure of the relation $\bigcup_{i \in N} \sim_i$:

- **Local Bisimulation Invariance** if, whenever $h \sim^* h'$ and h and h' are epistemically bisimilar¹, we have $h'e \in H$ iff $he \in H$.

¹The reader is referred to Subsection 3.1 for a precise definition of bisimulation invariance.

- **Perfect Recall** if, whenever $ha \sim_i h'b$, we also have $h \sim_i h'$.

- **Local No Miracles** if, whenever $ga \sim_i g'b$ and $g \sim^* h \sim_i h'$, then for every $h'a, hb \in H$, we also have $h'a \sim_i hb$.

◁

These properties describe the idealized epistemic agents that are presupposed in dynamic epistemic logic:

Theorem 1.4 (van Benthem et al. [9]) Let \mathcal{H} be an ETL model, \mathcal{M} an epistemic model, and the ‘protocol’ P a set of finite sequences of pointed events models closed under prefixes. We write \otimes for product update. Let $Forest(\mathcal{M}, P) = \bigcup_{\vec{e} \in P} \mathcal{M} \otimes \vec{e}$ be the ‘epistemic forest generated by’ \mathcal{M} and sequential application of the events in P .² The following are equivalent:

- \mathcal{H} is isomorphic to $Forest(\mathcal{M}, P)$.
- \mathcal{H} satisfies propositional stability, synchronicity, local bisimulation invariance, Perfect Recall, and Local No Miracles.

Thus, epistemic temporal conditions describing idealized epistemic agents characterize just those trees that arise from performing iterated product update governed by some protocol. [9] and [21, ch5] have details.

Our paper extends this analysis to the richer case of belief revision, where plausibility orders of agents evolve as they observe possibly surprising events. We prove two main results, with variations and extensions:

Theorem 1.5 Let \mathcal{H} be a doxastic temporal model, \mathcal{M} a plausibility model, \vec{e} a sequence of event models, and \otimes priority update. The following are equivalent, where the notions will of course be defined later:

1. \mathcal{H} is isomorphic to the forest generated by $\mathcal{M} \otimes \vec{e}$
2. \mathcal{H} satisfies propositional stability, synchronicity, invariance for bisimulation, as well as principles of Preference Propagation, Preference Revelation and Accommodation.

Theorem 1.6 Preference Propagation, Preference Revelation and Accommodation are definable in an extended doxastic modal language.

²For a more precise definition of this notion, see Section 2 below.

2 Definitions

We now turn to the definitions needed for the simplest version of our main representation theorem, postponing matching formal languages to Section 5. In what follows, let $N = \{1, \dots, n\}$ be a finite set of agents.

2.1 Plausibility models, event models and priority update

As for the epistemic case, we first introduce static models that encode the current prior (conditional) beliefs of agents. These carry a pre-order \leq between worlds encoding a plausibility relation. Often this relation is taken to be total, but when we think of elicited beliefs as *multi-criteria decisions*, a pre-order allowing for incomparable situations may be all we get [18]. We will therefore assume reflexivity and transitivity, but not totality.

As for notation: we write $a \simeq b$ ('indifference') if $a \leq b$ and $b \leq a$, and $a < b$ if $a \leq b$ and $b \not\leq a$.

Definition 2.1 [Doxastic Plausibility Models] A *doxastic plausibility model* $\mathcal{M} = \langle W, (\preceq_i)_{i \in N}, V \rangle$ has $W \neq \emptyset$, for each $i \in N$, \preceq_i is a pre-order on W , and $V : Prop \rightarrow \wp H$. \triangleleft

We now consider how such models evolve as agents observe events.

Definition 2.2 [Plausibility Event Model] A *plausibility event model* (event model, for short) ϵ is a tuple $\langle E, (\preceq_i)_{i \in N}, \text{pre} \rangle$ with $E \neq \emptyset$, each \preceq_i a pre-order on E , and $\text{pre} : E \rightarrow \mathcal{L}$, where \mathcal{L} is a doxastic language.³ \triangleleft

Definition 2.3 [Priority Update; [3]]

Priority update of a plausibility model $\mathcal{M} = \langle W, (\preceq_i)_{i \in N}, V \rangle$ and an event model $\epsilon = \langle E, (\preceq_i)_{i \in N}, \text{pre} \rangle$ is the plausibility model $\mathcal{M} \otimes \epsilon = \langle W', (\preceq'_i)_{i \in N}, V' \rangle$ defined as follows:

- $W' = \{(w, e) \in W \times E \mid \mathcal{M}, w \Vdash \text{pre}(e)\}$
- $(w, e) \preceq'_i (w', e')$ iff either $e \prec_i e'$, or $e \simeq_i e'$ and $w \preceq_i w'$
- $V'((s, e)) = V(s)$

Here, the new plausibility relation is still a pre-order. \triangleleft

The idea behind priority update is that beliefs about the last event override prior beliefs. If the agent is indifferent, however, the old plausibility order applies. More motivation can be found in [3, 8].

³This definition is incomplete without specifying the relevant language, but all that follows can be understood by considering the formal language as a 'parameter'.

2.2 Doxastic Temporal Models

Definition 2.4 [Doxastic Temporal Models] A *doxastic temporal model* (*DoTL* model for short) \mathcal{H} is of the form $\langle \Sigma, H, (\leq_i)_{i \in N}, V \rangle$, where Σ is a finite set of events, $H \subseteq \Sigma^*$ is closed under non-empty prefixes, for each $i \in N$, \leq_i is a pre-order on H , and $V : Prop \rightarrow \wp H$. \triangleleft

Our task is to identify just when a doxastic temporal model is isomorphic to the 'forest' generated by a sequence of priority updates:

2.3 Dynamic Models Generate Doxastic Temporal Models

Definition 2.5 [*DoTL* model generated by updates]

Each initial plausibility model $\mathcal{M} = \langle W, (\preceq_i)_{i \in N}, V \rangle$ and sequence of event models $\epsilon_j = \langle E_j, (\preceq_i^j)_{i \in N}, \text{pre}_j \rangle$ yields a *generated DoTL plausibility model* $\langle \Sigma, H, (\leq_i)_{i \in N}, V \rangle$ as follows:

- Let $\Sigma := \bigcup_{i=1}^m e_i$.
- Let $H_1 := W$ and for any $1 < n \leq m$ let $H_{n+1} := \{(we_1 \dots e_n) \mid (we_1 \dots e_{n-1}) \in H_n \text{ and } \mathcal{M} \otimes \epsilon_1 \otimes \dots \otimes \epsilon_{n-1} \Vdash \text{pre}_n(e_n)\}$. Finally let $H = \bigcup_{1 \leq k \leq m} H_k$.
- If $h, h' \in H_1$, then $h \leq_i h'$ iff $h \preceq_i^{\mathcal{M}} h'$.
- For $1 < k \leq m$, $he \leq_i h'e'$ iff 1. $he, h'e' \in H_k$, and 2. either $e \prec_i^k e'$, or $e \simeq_i^k e'$ and $h \leq_i h'$.
- Let $wh \in V(p)$ iff $w \in V(p)$.

This is a temporal doxastic model as above. \triangleleft

Now come the key doxastic temporal properties of our idealized agents.

3 Frame Properties for Priority Updaters

We first introduce the notion of bisimulation, modulo a choice of language.

3.1 Bisimulation Invariance

Definition 3.1 [\leq -Bisimulation]

Let \mathcal{H} and \mathcal{H}' be two *DoTL* plausibility models $\langle H, (\leq_1, \dots, \leq_n), V \rangle$ and $\langle H', (\leq'_1, \dots, \leq'_n), V' \rangle$ (for simplicity, assume they are based on the same alphabet Σ). A relation $Z \subseteq H \times H'$ is a \leq -Bisimulation if, for all $h \in H, h' \in H'$, and all \leq_i in (\leq_1, \dots, \leq_n) ,

(prop) h and h' satisfy the same proposition letters,

(zig) If hZh' and $h \leq_i j$, then there exists $j' \in H'$ such that jZj' and $h' \leq'_i j'$,

(zag) If hZh' and $h' \leq'_i j'$, then there exists $j \in H$ such that jZj' and $h \leq_i j$.

If Z is a \leq_n -bisimulation and hZh' , we will say that h and h' are \leq -bisimilar. \triangleleft

Definition 3.2 [\leq -Bisimulation Invariance] A *DoTL* model \mathcal{H} satisfies \leq -bisimulation invariance if, for all \leq -bisimilar histories $h, h' \in H$, and all events $e, h'e \in H$ iff $he \in H$. \triangleleft

3.2 Agent-Oriented Frame Properties

In the following we drop agent labels and the “for each $i \in N$ ” for the sake of clarity. Also, when we write ha we will always assume that $ha \in H$. We will make heavy use of the following notion:

Definition 3.3 [Accommodating Events]

Two events $a, b \in \Sigma$ are *accommodating* if, for all $ga, g'b$, ($g \leq g' \leftrightarrow ga \leq g'b$) and similarly for \geq , i.e., a, b preserve and anti-preserve plausibility. \triangleleft

Definition 3.4 Let $\mathcal{H} = \langle \Sigma, H, (\leq_i)_{i \in N}, V \rangle$ be a *DoTL* model. \mathcal{H} satisfies:

- **Propositional stability** if, whenever h is a finite prefix of h' , then h and h' satisfy the same proposition letters.
- **Synchronicity** if, whenever $h \leq h'$, we have $\text{len}(h) = \text{len}(h')$.

The following three properties trace the belief revising behavior of agents in doxastic trees.

- **Preference Propagation** if, whenever $ja \leq j'b$, then $h \leq h'$ implies $ha \leq h'b$.
- **Preference Revelation** if, whenever $jb \leq j'a$, then $ha \leq h'b$ implies $h \leq h'$.
- **Accommodation** if, a and b are *accommodating* whenever both $ja \leq j'b$ and $ha \not\leq h'b$.

\triangleleft

These properties - and in particular the last one - are somewhat trickier than in the epistemic case, reflecting the peculiarities of priority update in settings where incomparability is allowed. But we do have:

Fact 3.5 If \leq is a total pre-order and \mathcal{H} satisfies *Preference Propagation* and *Preference Revelation*, then \mathcal{H} satisfies *Accommodation*.

Proof. From left to right. Assume that $g \leq g'$ and $ja \leq j'b$. By *Preference Propagation*, $ga \leq g'b$. Now assume that $ha \not\leq h'b$. Then by totality, $h'b \leq ha$. Since $g \leq g'$, it follows by *Preference Propagation* that $gb \leq g'a$.

From right to left, assume that $gb \leq g'a$ and that $ja \leq j'b$. It follows by *Preference Revelation* that $g \leq g'$. Now assume that $ga \leq g'b$ (1) and $ha \not\leq h'b$ (2). From (2), it follows by totality that $h'b \leq ha$ (3). But if (3) and (1), then by *Preference Revelation* we have $g \leq g'$. QED

We can also prove a partial converse without totality:

Fact 3.6 If \mathcal{H} satisfies *Accommodation*, it satisfies *Preference Propagation*.

Proof. Let $ja \leq j'b$ (1) and $h \leq h'$ (2). Assume that $ha \not\leq h'b$. Then by *Accommodation*, for every $ga, g'b, g \leq g' \leftrightarrow ga \leq g'b$. So, in particular, $h \leq h' \leftrightarrow ha \leq h'b$. But since $h \leq h'$, we get $ha \leq h'b$: a contradiction. QED

No similar result holds for *Preference Revelation*. An easy counter-example shows that, even when \leq is total:

Fact 3.7 *Accommodation* does not imply *Preference Revelation*.

4 The Main Representation Theorem

We start with a warm-up case, taking plausibility to be a total pre-order.

4.1 Total pre-orders

Theorem 4.1 Let \mathcal{H} be a total doxastic-temporal model, \mathcal{M} a total plausibility model, \vec{e} a sequence of total event models, and let \otimes stand for priority update. The following are equivalent:

- \mathcal{H} is isomorphic to the forest generated by $\mathcal{M} \otimes \vec{e}$.
- \mathcal{H} satisfies *propositional stability*, *synchronicity*, *bisimulation invariance*, *Preference Propagation*, and *Preference Revelation*.

Proof.

Necessity We first show that the given conditions are indeed satisfied by any *DoTL* model generated through successive priority updates along some given protocol sequence. Here, *Propositional stability* and *Synchronicity* are straightforward from the definition of generated forests.

Preference Propagation Assume that $ja \leq j'b$ (1). It follows from (1) plus the definition of priority update that $a \leq b$ (2). Now assume that $h \leq h'$ (3). It follows from (2), (3) and priority update that $ha \leq h'b$.

Preference Revelation Assume that $jb \leq j'a$ (1). It follows from (1) and the definition of priority update that $b \leq a$ (2). Now assume $ha \leq h'b$ (3). By the definition of priority update, (3) can happen in two ways. Case 1: $a < b$ (4). It follows from (4) by the definition of $<$ that $b \not\leq a$ (5). But (5) contradicts (2). We are therefore in Case 2: $a \simeq b$ (6) and $h \leq h'$ (7). But (7) is precisely what we wanted to show.

Note that we did not make use of totality here.

Sufficiency Given a *DoTL* model \mathcal{M} , we first show how to construct a *DDL* model, i.e., a plausibility model and a sequence of event models.

Construction Here is the initial plausibility model $\mathcal{M} = \langle W, (\preceq_i)_{i \in N}, \hat{V} \rangle$:

- $W := \{h \in H \mid \text{len}(h) = 1\}$.
- Set $h \preceq_i h'$ iff $\hat{V}_i(h) \leq \hat{V}_i(h')$.
- For every $p \in \text{Prop}$, $\hat{V}(p) = V(p) \cap W$.

Now we construct the j -th event model $\epsilon_j = \langle E_j, (\preceq_i^j)_{i \in N}, \text{pre}_j \rangle$:

- $E_j := \{e \in \Sigma \mid \text{there is a history } he \in H \text{ with } \text{len}(h) = j\}$
- For each $i \in N$, set $a \preceq_i^j b$ iff there are $ha, h'b \in H$ such that $\text{len}(h) = \text{len}(h') = j$ and $ha \preceq_i h'b$.
- For each $e \in E_j$, let $\text{pre}_j(e)$ be the formula that characterizes the set $\{h \mid he \in H \text{ and } \text{len}(h) = j\}$. By general modal logic, bisimulation invariance guarantees that there is such a formula, though it may be an infinitary one in general.

Now we show that the construction is correct in the following sense:

Claim 4.2 (Correctness) *Let \leq be the plausibility relation in the given doxastic temporal model. Let \preceq_{DDL}^F be the plausibility relation in the forest induced by priority update over the just constructed plausibility model and matching sequence of event models. We have:*

$$h \leq h' \text{ iff } h \preceq_{DDL}^F h'.$$

Proof of the claim The proof is by induction on the length of histories. The base case is obvious from the construction of our initial model \mathcal{M} . Now for the induction step. As for notation we will write $a \leq b$ for $a \preceq_i^n b$ with n the length for which the claim has been proved, and i an agent.

From DoTL to Forest(DDL) Assume that $h_1a \leq h_2b$ (1). It follows that in the constructed event model $a \leq b$ (2). Case 1: $a < b$. By priority update we have $h_1a \preceq_{DDL}^F h_2b$. Case 2: $b \leq a$ (3). This means that there are h_3b, h_4a such that $h_3b \leq h_4a$. But then by *Preference Revelation* and (1) we have $h_1 \leq h_2$ (in the doxastic temporal model). It follows by the inductive hypothesis that $h_1 \preceq_{DDL}^F h_2$. But then by priority update, since by (2) and (3) a and b are indifferent, we have $h_1a \preceq_{DDL}^F h_2b$.

From Forest(DDL) to DoTL Next let $h_1a \preceq_{DDL}^F h_2b$. The definition of priority update has two clauses. Case 1: $a < b$. By definition, this implies that $b \not\leq a$. But then by the above construction, for all histories $h_3, h_4 \in H$ we have $h_3b \not\leq h_4a$. In particular we have $h_2b \not\leq h_1a$. But then by *totality*⁴, $h_1a \leq h_2b$. Case 2: $a \simeq b$ (4) and $h_1 \preceq_{DDL}^F h_2$. For a start, by the inductive hypothesis, $h_1 \leq h_2$ (5). By (4) and our construction, there are h_3a, h_4b with $h_3a \leq h_4b$ (6). But then by *Preference Propagation*, (5) and (6) imply that we have $h_1a \leq h_2b$. QED

Next, we turn to the general case of pre-orders, allowing incomparability.

4.2 The general case

While the argument went smoothly for *total* pre-orders, it gets somewhat more interesting when incomparability enters the stage. In the case of pre-orders we need the additional axiom of Accommodation as stated below:

Theorem 4.3 *Let \mathcal{H} be a doxastic-temporal model, \mathcal{M} a plausibility model, \vec{e} be a sequence of event models while \otimes is priority update. The following assertions are equivalent:*

- \mathcal{H} is isomorphic to the forest generated by $\mathcal{M} \otimes \vec{e}$,
- \mathcal{H} satisfies bisimulation invariance, propositional stability, synchronicity, Preference Revelation and Accommodation.

By Fact 3.6, requiring Accommodation also gives us Preference Propagation.

Proof.

Necessity of the conditions The verification of the conditions in the preceding subsection did not use totality. So we concentrate on the new condition:

⁴Note that this is the only place in which we make use of totality.

Accommodation Assume that $ja \leq j'b$ (1). It follows by the definition of priority update that $a \leq b$ (2). Now let $ha \not\leq h'b$ (3). This implies by priority update that $a \not\leq b$ (4). By definition, (2) and (4) means that $a \simeq b$ (5). Now assume that $g \leq g'$ (6). It follows from (5), (6) and priority update that $ga \leq g'b$. For the other direction of the consequent assume instead that $g \not\leq g'$ (7). It follows from (5), (7) and priority update that $ga \not\leq g'b$.

Sufficiency of the conditions Given a *DoTL* model, we again construct a *DDL* plausibility model plus sequence of event models:

Construction The plausibility model $\mathcal{M} = \langle W, (\preceq_i)_{i \in N}, \hat{V} \rangle$ is as follows:

- $W := \{h \in H \mid \text{len}(h) = 1\}$,
- Set $h \preceq_i h'$ whenever \leq_i ,
- For every $p \in \text{Prop}$, $\hat{V}(p) = V(p) \cap W$.

We construct the j -th event model $\epsilon_j = \langle E_j, (\preceq_i^j)_{i \in N}, \text{pre}_j \rangle$ as follows:

- $E_j := \{e \in \Sigma \mid \text{there is a history of the form } he \in H \text{ with } \text{len}(h) = j\}$
- For each $i \in N$, define $a \preceq_i^j b$ iff either (a) there are $ha, h'b \in H$ such that $\text{len}(h) = \text{len}(h') = j$ and $ha \leq_i h'b$, or (b) [a new case] a and b are accommodating, and we put $a \simeq b$ (i.e. $a \leq b$ and $b \leq a$).
- For each $e \in E_j$, let $\text{pre}_j(e)$ be the formula that characterizes the set $\{h \mid he \in H \text{ and } \text{len}(h) = j\}$. *Bisimulation invariance* guarantees that there is always such a formula (maybe involving an infinitary syntax).

Again we show that the construction is correct in the following sense:

Claim 4.4 (Correctness) *Let \leq be the plausibility relation in the doxastic temporal model. Let \preceq_{DDL}^F be the plausibility relation in the forest induced by successive priority updates of the plausibility model by the sequence of event models we constructed. We have:*

$$h \leq h' \text{ iff } h \preceq_{DDL}^F h'.$$

Proof of the claim We proceed by induction on the length of histories. The base case is clear from our construction of the initial model \mathcal{M} . Now for the induction step, with the same simplified notation as earlier.

From DoTL to Forest(DEL) There are two cases:

Case 1. $ha \leq h'b, h \leq h'$. By the inductive hypothesis, $h \leq h'$ implies $h \preceq_{DDL}^F h'$ (1). Since $ha \leq h'b$, it follows by construction that $a \leq b$ (2). It follows from (1) and (2) that by priority update $ha \preceq_{DDL}^F h'b$.

Case 2. $ha \leq h'b, h \not\leq h'$. Clearly, then, a and b are not accommodating and thus the special clause has not been used to build the event model, though we do have $a \leq b$ (1). By the contrapositive of Preference Revelation, we also conclude that for all $ja, j'b \in H$, we have $j'b \not\leq ja$ (2). Therefore, our construction gives $b \not\leq a$ (3), and we conclude that $a < b$ (4). But then by priority update, we get $ha \preceq_{DDL}^F h'b$.

From Forest(DEL) to DoTL We distinguish again two relevant cases.

Case 1. $ha \preceq_{DDL}^F h'b, h \preceq_{DDL}^F h'$. By definition of priority update, $ha \preceq_{DDL}^F h'b$ implies that $a \leq b$ (1). There are two possibilities. Case 1: The special clause of the construction has been used, and a, b are accommodating (2). By the inductive hypothesis, $h \preceq_{DDL}^F h'$ implies $h \leq h'$ (3). But (2) and (3) imply that $ha \leq h'b$. Case 2: Clause (1) holds because for some $ja, j'b \in H$, in the *DoTL* model, $ja \leq j'b$ (4). By the inductive hypothesis, $h \preceq_{DDL}^F h'$ implies $h \leq h'$ (5). Now, it follows from (4), (5) and Preference Propagation that $ha \leq h'b$.

Case 2. $ha \preceq_{DDL}^F h'b, h \not\preceq_{DDL}^F h'$. Here is where we put our new accommodation clause to work. Let us label our assertions: $h \not\preceq_{DDL}^F h'$ (1) and $ha \preceq_{DDL}^F h'b$ (2). It follows from (1) and (2) by the definition of priority update that $a < b$ (3), and hence, by definition $b \not\leq a$ (4). Clearly, a and b are not accommodating (5): for otherwise, we would have had $a \simeq b$, and hence $b \leq a$, contradicting (4). Therefore, (3) implies that there are $ja, j'b \in H$ with $ja \leq j'b$ (6). Now assume for *contradictio* that (in the *DoTL* model) $ha \not\leq h'b$ (7). It follows from (6) and (7) by Accommodation that a and b are accommodating, contradicting (5). Thus we have $ha \leq h'b$. QED

Given a doxastic temporal model describing the evolution of the beliefs of a group of agents, we have determined whether it could have been generated by successive ‘local’ priority updates of a plausibility model. Of course, further scenarios are possible, e.g., bringing in knowledge as well. We discuss some extensions in the next subsection.

4.3 Extensions and variations

4.3.1 Unified plausibility models

There are two roads to merging epistemic indistinguishability and doxastic plausibility. The first works with a plau-

sibility order *and* an epistemic indistinguishability relation, explaining the notion of *belief* with a mixture of the two. Baltag and Smets [3] apply product update to epistemic indistinguishability and priority update to the plausibility relation. A characterization for the doxastic epistemic temporal models induced in this way follows from van Benthem et al. [9] Theorem 1.4 plus Theorem 4.3 of previous subsection (or its simpler counterpart for total orders). All this has the flavor of working with *prior* beliefs and information partitions, taking the *posteriors* to be computed from them.

However there are also reasons for working with (*posterior*) beliefs only (see e.g. [22]). Indeed, Baltag and Smets [3] take this second road, using *unified* ‘local’ plausibility models with just one explicit relation \trianglelefteq . We briefly show how our earlier results transform to this setting. In what follows, we write $a \cong b$ iff $a \trianglelefteq b$ and $b \trianglelefteq a$.

Definition 4.5 The priority update of a *unified* plausibility model $\mathcal{M} = \langle W, (\trianglelefteq_i)_{i \in N}, V \rangle$ and a \trianglelefteq -event model $\epsilon = \langle E, (\trianglelefteq_i)_{i \in N}, \text{pre} \rangle$ is the unified plausibility model $\mathcal{M} \otimes \epsilon = \langle W', (\trianglelefteq'_i)_{i \in N}, V' \rangle$ constructed as follows:

- $W' = \{(w, e) \in W \times E \mid \mathcal{M}, w \Vdash \text{pre}(e)\}$,
- $(w, a) \trianglelefteq'_i (w', b)$ iff either 1. $a \trianglelefteq_i b$, $b \not\trianglelefteq_i a$ and $w \trianglelefteq w' \vee w' \trianglelefteq w$ or 2. $a \trianglelefteq_i b$, $b \trianglelefteq_i a$ and $w \trianglelefteq w'$,
- $V'((s, e)) = V(s)$.

◁

Here are our familiar key properties in this setting:

Agent revision properties in terms of \trianglelefteq_i

- \trianglelefteq -Perfect Recall if, whenever $ha \trianglelefteq h'b$ we have $h \trianglelefteq h' \vee h' \trianglelefteq h$.
- \trianglelefteq -Preference Propagation if, whenever $h \trianglelefteq h'$ and $ja \trianglelefteq j'b$ then $ha \trianglelefteq h'b$.
- \trianglelefteq -Preference Revelation if, whenever $ha \trianglelefteq h'b$ and $jb \trianglelefteq j'a$, also $h \trianglelefteq h'$.
- \trianglelefteq -Accommodation if, whenever $(ja \trianglelefteq j'b, h' \trianglelefteq h$ and $ha \not\trianglelefteq h'b)$, for all $ga, g'b \in H$ ($g \trianglelefteq g' \leftrightarrow ga \trianglelefteq g'b$), and for all $g'a, gb \in H$ ($g \trianglelefteq g' \leftrightarrow gb \trianglelefteq g'a$).

The last axiom is slightly weaker than Accommodation. The following result is proved in the extended version of this paper.

Theorem 4.6 *Let \mathcal{H} be a unified doxastic-temporal model, \mathcal{M} a unified plausibility model, $\vec{\epsilon}$ be a sequence of unified event models, while \otimes is priority update. The following assertions are equivalent:*

- \mathcal{H} is isomorphic to the forest generated by $\mathcal{M} \otimes \vec{\epsilon}$,
- \mathcal{H} satisfies bisimulation invariance, propositional stability, synchronicity, \trianglelefteq -Perfect Recall, \trianglelefteq -Preference Propagation, \trianglelefteq -Preference Revelation and \trianglelefteq -Accommodation.

Our next source of variation is an issue that we have left open throughout our analysis so far, which may have bothered some readers.

4.3.2 Bisimulations and pre-condition languages

Our definition of event models presupposed a language for the preconditions, and correspondingly, the right notion of bisimulation in our representation results should matching (at least, on finite models) the precondition language used. For instance, if the precondition language contains a belief operator scanning the *intersection* of a plausibility \leq_i relation and an epistemic indistinguishability relation \sim , then the *zig* and *zag* clauses should not only apply to \leq_i and \sim_i separately, but also to $\leq_i \cap \sim_i$. And things get even more complicated if we allow temporal operators in our languages (cf. [10]). We do not commit to any specific choice here, since the choice of a language seems orthogonal to our main concerns. But we will discuss formal languages in the next section, taking definability of our major structural constraints as a guide.

Finally, our results can be generalized by including one more major parameter in describing processes:

4.3.3 Protocols

So far we have assumed that the same sequences of events were executable uniformly anywhere in the initial doxastic model, provided the worlds fulfilled the preconditions. This strong assumption is lifted in [10, 9], who allow the *protocol*, i.e., the set of executable sequences of events forming our current informational process, to vary from state to state. Initially, they still take the protocol to be *common knowledge*, but eventually, they allow for scenarios where agents need not know which protocol is running. These variations change the complete dynamic-epistemic logic of the system. It would be of interest to extend this work to our extended doxastic setting.

5 Dynamic Languages and Temporal Doxastic Languages

Our emphasis so far has been on structural properties of models. To conclude, we turn to the logical languages that can express these, and hence also, the type of doxastic reasoning our agents can be involved with.

5.1 Dynamic doxastic language

We first look at a core language that matches dynamic belief update.

5.1.1 Syntax

Definition 5.1 [Dynamic Doxastic-Epistemic language] The language of dynamic doxastic language $DDEL$ is defined as follows:

$$\phi := p \mid \neg\phi \mid \phi \vee \psi \mid \langle \leq_i \rangle \phi \mid \langle i \rangle \phi \mid E\phi \mid \langle \epsilon, \mathbf{e} \rangle \phi$$

where i ranges over N , p over a countable set of proposition letters $Prop$, and (ϵ, \mathbf{e}) ranges over a suitable set of symbols for event models. \triangleleft

All our dynamic doxastic logics will be interpreted on the following models.

5.1.2 Models

Definition 5.2 [Epistemic Plausibility Models] An *epistemic plausibility model* $\mathcal{M} = \langle W, (\leq_i)_{i \in N}, (\sim_i)_{i \in N}, V \rangle$ has $W \neq \emptyset$, and for each $i \in N$, \leq_i is a pre-order on W , and \sim_i any relation, while $V : Prop \rightarrow \wp H$. \triangleleft

Definition 5.3 [\sim, \leq -event model] An *epistemic plausibility event model* (\sim, \leq -event model for short) ϵ is of the form $\langle E, (\leq_i)_{i \in N}, (\sim_i)_{i \in N}, \mathbf{pre} \rangle$ where $E \neq \emptyset$, for each $i \in N$, \leq_i is a pre-order on E and \sim_i is a relation on W . Also, there is a precondition function $\mathbf{pre} : E \rightarrow DDEL$. \triangleleft

Definition 5.4 [Priority update] The *priority update* of an epistemic plausibility model $\mathcal{M} = \langle W, (\leq_i)_{i \in N}, (\sim_i)_{i \in N}, V \rangle$ and a \sim, \leq -event model $\epsilon = \langle E, (\leq_i)_{i \in N}, (\sim_i)_{i \in N}, \mathbf{pre} \rangle$ is the plausibility model $\mathcal{M} \otimes \epsilon = \langle W', (\leq'_i)_{i \in N}, V' \rangle$ whose structure is defined as follows:

- $W' = \{(w, e) \in W \times E \mid \mathcal{M}, w \Vdash \mathbf{pre}(e)\}$
- $(w, e) \leq'_i (w', e')$ iff $e \prec_i e'$, or $e \simeq_i e'$ and $w \leq_i w'$
- $(w, e) \sim'_i (w', e')$ iff $e \sim_i e'$ and $w \sim_i w'$
- $V'((s, e)) = V(s)$.

The result of the update is an epistemic plausibility model. \triangleleft

5.1.3 Semantics

Here is how we interpret the $DDE(L)$ language. A pointed event model is an event model plus an element of its domain. To economize on notation we use event symbols in the semantic clause. We write $\mathbf{pre}(e)$ for $\mathbf{pre}_\epsilon(e)$ when it is clear from context.

Definition 5.5 [Truth definition]

$$\text{Let } K_i[w] = \{v \mid w \sim_i v\}.$$

$$\begin{aligned} \mathcal{M}, w \Vdash p & \text{ iff } w \in V(p) \\ \mathcal{M}, w \Vdash \neg\phi & \text{ iff } \mathcal{M}, w \not\Vdash \phi \\ \mathcal{M}, w \Vdash \phi \vee \psi & \text{ iff } \mathcal{M}, w \Vdash \phi \text{ or } \mathcal{M}, w \Vdash \psi \\ \mathcal{M}, w \Vdash \langle \leq_i \rangle \phi & \text{ iff } \exists v \text{ such that } w \leq_i v \text{ and } \mathcal{M}, v \Vdash \phi \\ \mathcal{M}, w \Vdash \langle i \rangle \phi & \text{ iff } \exists v \text{ such that } v \in K_i[w] \text{ and } \mathcal{M}, v \Vdash \phi \\ \mathcal{M}, w \Vdash E\phi & \text{ iff } \exists v \in W \text{ such that } \mathcal{M}, v \Vdash \phi \\ \mathcal{M}, w \Vdash \langle \epsilon, \mathbf{e} \rangle \phi & \text{ iff } \mathcal{M}, w \Vdash \mathbf{pre}(e) \text{ and } \mathcal{M} \times \epsilon, (w, e) \Vdash \phi \end{aligned}$$

\triangleleft

The knowledge operator K_i and the universal modality \mathbf{A} are defined as usual.

5.1.4 Reduction axioms

The methodology of dynamic epistemic and dynamic doxastic logics revolves around *reduction* axioms. On top of some complete static base logic, these fully describe the dynamic component. Here is well-known *Action – Knowledge* reduction axiom of [2]:

$$[\epsilon, \mathbf{e}]K_i\phi \leftrightarrow (\mathbf{pre}(e) \rightarrow \bigwedge \{K_i[\epsilon, \mathbf{f}]\phi : e \sim_i \mathbf{f}\}) \quad (1)$$

Similarly, here are the key reduction axioms for $\langle \epsilon, \mathbf{e} \rangle \langle \leq_i \rangle$ with priority update:

Proposition 5.6 *The following dynamic-doxastic principle is sound for plausibility change:*

$$\begin{aligned} \langle \epsilon, \mathbf{e} \rangle \langle \leq_i \rangle \phi \leftrightarrow \\ (\mathbf{pre}(e) \wedge (\langle \leq_i \rangle \bigvee \{\langle \mathbf{f} \rangle \phi : e \simeq_i \mathbf{f}\} \vee \\ E \bigvee \{\langle \mathbf{g} \rangle \phi : e <_i \mathbf{g}\})) \end{aligned} \quad (2)$$

The crucial feature of such a dynamic ‘recursion step’ is that the order between *action* and *belief* is reversed. This works because, conceptually, the current beliefs already *pre-encode* the beliefs after some specified event. In the epistemic setting, principles like this also reflect agent properties of Perfect Recall and No Miracles [11]. Here, they rather encode radically ‘event-oriented’ revision policies, and the same point applies to the principles we will find later in a doxastic temporal setting.

Finally for the existential modality $\langle \epsilon, \mathbf{e} \rangle E$ we note the following fact:

Proposition 5.7 *The following axiom is valid for the existential modality:*

$$\langle e, \epsilon \rangle E\phi \leftrightarrow (\text{pre}(e) \wedge (E \bigvee \{ \langle f \rangle \phi : f \in \text{Dom}(\epsilon) \})) \quad (3)$$

We do not pursue further issues of axiomatic completeness here, since we are just after the model theory of our dynamic and temporal structures.

5.2 Doxastic epistemic temporal language

Next *epistemic-doxastic temporal models* are simply our old doxastic temporal models \mathcal{H} extended with epistemic accessibility relations \sim_i .

5.2.1 Syntax

Definition 5.8 [Doxastic Epistemic Temporal Languages]

The language of *DETL* is defined by the following inductive syntax:

$$\phi := p \mid \neg\phi \mid \phi \vee \psi \mid \langle e \rangle \phi \mid \langle e^{-1} \rangle \phi \mid \langle \leq_i \rangle \phi \mid \langle i \rangle \phi \mid E\phi$$

where i ranges over N , e over Σ , and p over proposition letters *Prop*. \triangleleft

5.2.2 Semantics

The language *DETL* is interpreted over nodes h in our trees (cf. [11]):

Definition 5.9 [Truth definition]

$$\text{Let } K_i[h] = \{h' \mid h \sim_i h'\}.$$

$\mathcal{H}, h \Vdash p$	iff	$h \in V(p)$
$\mathcal{H}, h \Vdash \neg\phi$	iff	$\mathcal{H}, h \not\Vdash \phi$
$\mathcal{H}, h \Vdash \phi \vee \psi$	iff	$\mathcal{H}, h \Vdash \phi$ or $\mathcal{H}, h \Vdash \psi$
$\mathcal{H}, h \Vdash \langle e \rangle \phi$	iff	$\exists h' \in H$ s.t. $h' = he$ and $\mathcal{H}, h' \Vdash \phi$
$\mathcal{H}, h \Vdash \langle e^{-1} \rangle \phi$	iff	$\exists h' \in H$ s.t. $h'e = h$ and $\mathcal{H}, h' \Vdash \phi$
$\mathcal{H}, h \Vdash \langle \leq_i \rangle \phi$	iff	$\exists h' \text{ s.t. } h \leq_i h'$ and $\mathcal{H}, h' \Vdash \phi$
$\mathcal{H}, h \Vdash \langle i \rangle \phi$	iff	$\exists h' \text{ s.t. } h' \in K_i[h]$ and $\mathcal{H}, h' \Vdash \phi$
$\mathcal{H}, h \Vdash E\phi$	iff	$\exists h' \in H$ s.t. $\mathcal{H}, h' \Vdash \phi$

\triangleleft

Now we have the right syntax to analyze our earlier structural conditions.

5.3 Defining the frame conditions

We will prove semantic *correspondence results* (cf. [13]) for our crucial properties using somewhat technical axioms that simplify the argument. Afterwards, we present some reformulations whose meaning for belief-revising agents may be more intuitive to the reader:

5.3.1 The key correspondence result

Theorem 5.10 (Definability) *Preference Propagation, Preference Revelation and Accommodation are definable in the doxastic-epistemic temporal language DETL.*

- \mathcal{H} satisfies *Preference Propagation* iff the following axiom is valid:

$$\begin{aligned} E\langle a \rangle \langle \leq_i \rangle \langle b^{-1} \rangle \top &\rightarrow \\ ((\langle \leq_i \rangle \langle b \rangle p \wedge \langle a \rangle q) &\quad (PP) \\ \rightarrow \langle a \rangle (q \wedge \langle \leq_i \rangle p) &\end{aligned}$$

- \mathcal{H} satisfies *Preference Revelation* iff the following axiom is valid:

$$\begin{aligned} E\langle b \rangle \langle \leq_i \rangle \langle a^{-1} \rangle \top &\rightarrow \\ (\langle a \rangle \langle \leq_i \rangle (p \wedge \langle b^{-1} \rangle \top) &\rightarrow \langle \leq_i \rangle \langle b \rangle p) \quad (PR) \end{aligned}$$

- \mathcal{H} satisfies *Accommodation* iff the following axiom is valid:

$$\begin{aligned} E\langle a \rangle \langle \leq_i \rangle \langle b^{-1} \rangle \top & \\ \wedge E[\langle a \rangle (p_1 \wedge E(p_2 \wedge \langle b^{-1} \rangle \top)) & \\ \wedge [a](p_1 \rightarrow [\leq_i] \neg p_2)] &\quad (AC) \\ \rightarrow ((\langle \leq_i \rangle \langle b \rangle q \rightarrow [a] \langle \leq_i \rangle q) & \\ \wedge (\langle a \rangle \langle \leq_i \rangle (r \wedge \langle b^{-1} \rangle \top) \rightarrow \langle \leq_i \rangle \langle b \rangle r) &\end{aligned}$$

Proof. We only prove the case of *Preference Propagation*, the other two are in the extended version of the paper. We drop agent labels for convenience.

(PP) characterizes Preference Propagation We first show that (PP) is valid on all models \mathcal{H} based on preference-propagating frames. Assume that $\mathcal{H}, h \Vdash E\langle a \rangle \langle \leq_i \rangle \langle b^{-1} \rangle \top$ (1). Then there are $ja, j'b \in H$ such that $ja \leq j'b$ (2). Now let $\mathcal{H}, h \Vdash ((\langle \leq_i \rangle \langle b \rangle p \wedge \langle a \rangle q)$ (3). Then there is $h' \in H$ such that $h \leq h'$ (4) and $\mathcal{H}, h' \Vdash \langle b \rangle p$ (5), while also $\mathcal{H}, ha \Vdash q$ (6). We must show that $\mathcal{H}, h \Vdash \langle a \rangle (q \wedge \langle \leq_i \rangle p)$ (7). But, from (2),(4),(6) and *Preference Propagation*, we get $ha \leq h'b$, and the conclusion follows by the truth definition.

Next, we assume that axiom (PP) is valid on a doxastic temporal frame, that is, true under any interpretation of its proposition letters. So, assume that $ja \leq j'b$ (1), and also $h \leq h'$ (2). Moreover, let $ha, h'b \in H$ (3). First note that (1) automatically verifies the antecedent of (PP) in any node of the tree. Next, we make the antecedent of the second implication in (PP) true at h by interpreting the proposition letter p as just the singleton set of nodes $h'b$, and q as just ha (4). Since (PP) is valid, its consequent will also hold under this particular valuation V . Explicitly we have $\mathcal{H}, V, h \Vdash \langle a \rangle (q \wedge \langle \leq_i \rangle p)$. But spelling out what p, q mean there, we get just the desired conclusion that $ha \leq h'b$. QED

The preceding correspondence argument is really just a *Sahlqvist* substitution case (cf. [13]), and so are the other two. We do not prove a further completeness result, but will show one nice derivation, as a syntactic counterpart to our earlier Fact 3.5.

$$\begin{aligned} & \mathbf{E} [\langle a \rangle (\psi \wedge \mathbf{E} (\phi \wedge \langle b^{-1} \rangle \top)) \wedge [a] (\psi \rightarrow [\leq_i] \neg \phi)] \\ & \rightarrow (\langle a \rangle \langle \leq_i \rangle (\phi \wedge \langle b^{-1} \rangle \top) \rightarrow \langle \leq_i \rangle \langle b \rangle \phi) \end{aligned} \quad (F)$$

Here is an auxiliary correspondence observation:

Fact 5.11 *On total doxastic temporal models the following axiom is valid:*

$$\begin{aligned} & \langle a \rangle (\psi \wedge \mathbf{E} (\phi \wedge \langle b^{-1} \rangle \top)) \rightarrow \\ & (\langle a \rangle (\psi \wedge \langle \leq_i \rangle \phi) \vee \mathbf{E} \langle b \rangle (\phi \wedge \langle \leq_i \rangle (\psi \wedge \langle a^{-1} \rangle \top)) \end{aligned} \quad (Tot)$$

Now we can state an earlier semantic fact in terms of axiomatic derivability in some obvious minimal system for the language *DETL*:

Fact 5.12

- $\vdash ((PP) \wedge (F)) \rightarrow (AC)$
- $\vdash ((PR) \wedge (Tot)) \rightarrow (F)$

We leave the simple combinatorial details to the extended version of this paper. We now get an immediate counterpart to Fact 3.5:

Corollary 5.13

$$\vdash ((PP) \wedge (PR) \wedge (Tot)) \rightarrow (AC) \quad (4)$$

5.3.2 Two intuitive explanations

Here are two ways to grasp the intuitive meaning of our technical axioms.

Reformulation with safe belief. An intermediate notion of knowledge first considered by [24] has been argued for doxastically as *safe belief* by [3] as describing those beliefs we do not give up under true new information. The safe belief modality \square^{\geq} is just the universal dual of the existential modality $\langle \geq \rangle$ scanning the converse of \leq . Without going into details of its logic (e.g., safe belief is positively, but not negatively introspective), here is how we can rephrase our earlier axiom:

- \mathcal{H} satisfies Preference Propagation iff the following axiom is valid on \mathcal{H} :

$$\mathbf{E} \langle a \rangle \langle \geq \rangle \langle b^{-1} \rangle \top \rightarrow (\langle a \rangle \square^{\geq_i} p \rightarrow \square^{\geq_i} [b] p) \quad (PP')$$

A similar reformulation is easy to give for Preference Revelation. These principles reverse action modalities and safe belief much like the better-known Knowledge-Action interchange laws in the epistemic-temporal case. We invite the reader to check their intuitive meaning in terms of acquired safe beliefs as informative events happen.

Analogs with reduction axioms Another way to understand the above axioms in their original format with existential modalities is their clear analogy with the reduction axiom for priority update. Here are two cases juxtaposed:

$$\begin{aligned} \langle \epsilon, e \rangle \langle \leq_i \rangle p & \leftrightarrow \\ & (\mathbf{pre}(e) \wedge (\langle \leq_i \rangle \bigvee \{ \langle f \rangle p : e \simeq_i f \}) \\ & \vee \mathbf{E} \bigvee \{ \langle g \rangle p : e <_i g \}) \end{aligned} \quad (2)$$

$$\begin{aligned} & \mathbf{E} \langle a \rangle \langle \leq_i \rangle \langle b^{-1} \rangle \top \rightarrow \\ & (\langle \leq_i \rangle \langle b \rangle p \rightarrow [a] \langle \leq_i \rangle p) \end{aligned} \quad (PP)$$

$$\begin{aligned} & \mathbf{E} \langle b \rangle \langle \leq_i \rangle \langle a^{-1} \rangle \top \rightarrow \\ & (\langle a \rangle \langle \leq_i \rangle (p \wedge \langle b^{-1} \rangle \top) \rightarrow \langle \leq_i \rangle \langle b \rangle p) \end{aligned} \quad (PR)$$

Family resemblance is obvious, and indeed, *(PP)* and *(PR)* may be viewed as the two halves of the reduction axiom, transposed to the more general setting of arbitrary doxastic-temporal models.

5.4 Variations and extensions of the doxastic temporal language

5.4.1 Weaker languages

The above doxastic-temporal language is by no means the only reasonable one. Weaker forward-looking modal fragments also make sense, dropping both converse and the existential modality. But they do not suffice for the purpose of our correspondence.

Proposition 5.14 (Undefinability)

Preference Propagation, Preference Revelation and Accommodation are not definable in the forward looking fragment of DETL.

Proof. The reason is the same in all cases: we show that these properties are not preserved under taking *bounded p-morphic images*. The Figure gives an indication how this works concretely. QED




Figure 1. Preference Propagation is not preserved under p -morphic images

5.4.2 Richer languages

But there is also a case to be made for richer languages. For instance, if we want to define the frame property of *synchronicity*, we must introduce an *equilevel relation* in our models, with a corresponding modality for it. While expressing synchronicity then becomes easy, this move is dangerous in principle. Van Benthem and Pacuit [11] point at the generally high complexity of tree logics when enriched with this expressive power.

Likewise, finer epistemic and doxastic process descriptions require further temporal modalities, such as “Since” and “Until”, beyond the basic operators we used for matching the needs of dynamic doxastic logic directly.

Finally, there may be even more urgent language extensions for doxastic temporal logic, having to do with our very notion of belief. We have emphasized the notion of *safe belief*, which scans the plausibility relation \geq as an ordinary modality. This notion can be used to define the more standard notion of belief as truth in all most plausible worlds: cf. [15]. But it has been argued recently by [3], and also by [16] that we really want a more ‘entangled’ version of the latter notion as well, referring to the most plausible worlds *inside the epistemically accessible ones*. Such a notion of ‘posterior belief’ has the following semantics:

$$\mathcal{H}, h \Vdash B_i \phi \quad \text{iff} \quad \forall h' \in \text{Min}(K_i[h], \leq_i) \text{ we have } \mathcal{H}, h' \Vdash \phi$$

Technically, expressing this requires an additional intersection modality. While this extension loses some typical modal properties, it does satisfy reduction axioms in the format discussed here: cf. [21].

6 Conclusion

Agents that update their knowledge and revise their beliefs can behave very differently over time. We have determined the special constraints that capture agents operating with the ‘local updates’ of dynamic doxastic logic. This took the form of some representation theorems that state just when a general doxastic temporal model is equivalent to the forest model generated by successive priority updates of an initial doxastic model by a protocol sequence of event models. We have also shown how these conditions can be defined in an appropriate extended modal language, making it possible to reason formally about agents engaged in such updates and revisions. Our methods are like those of existing epistemic work, but the doxastic case came with some interesting new notions.

As for open problems, the paper has indicated several technical issues along the way, e.g., concerning the expressive power of different languages over our models and their complexity effects (cf. [11] for the epistemic case). In particular, we have completely omitted issues of common knowledge and common belief, even though these are known to generate complications [12].

But from where we are standing now, we see several larger directions to pursue:

- A systematic “protocol logic” of axiomatic completeness for constrained revision processes, analogous to the purely epistemic theory of observation and conversation protocols initiated in [9],
- A comparison of our ‘constructive’ *DDL*-inspired approach to *DTL* universes with the more abstract *AGM*-style postulational approach of [14],
- A theory of variation for different sorts of agents with different abilities and tendencies, as initiated in [21],
- An analysis of knowledge and belief dynamics in games [7, 17, 4]
- Connections with formal learning theory over epistemic-doxastic temporal universes (cf. [20]).

References

- [1] A. Baltag and L. S. Moss. Logics for Epistemic Programs. *Synthese*, 139(2):165–224, 2004.
- [2] A. Baltag, L. S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, pages 43–56, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.

- [3] A. Baltag and S. Smets. Dynamic belief revision over multi-agent plausibility models. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory: Proceedings of LOFT'06*, Texts in Logic and Games, pages 11–24. Amsterdam University Press, 2006.
- [4] A. Baltag, S. Smets, and J. Zvesper. Dynamic Rationality in Extensive Games. 2008. (in this volume).
- [5] J. Benthem. Dynamic logics for belief change. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.
- [6] J. v. Benthem. Games in Dynamic Epistemic Logic. *Bulletin of Economic Research*, 53(4):219–248, 2001.
- [7] J. v. Benthem. Rational dynamics and epistemic logic in games. *International Game Theory Review*, 9(1):377–409, 2007. Erratum reprint, Volume 9(2), 2007, 377 - 409.
- [8] J. v. Benthem. Priority product update as social choice. *ILLC Prepublications Series*, PP-2008-09, 2008.
- [9] J. v. Benthem, J. Gerbrandy, T. Hoshi, and E. Pacuit. Merging frameworks for interaction: DEL and ETL. Working paper, 2008.
- [10] J. v. Benthem, J. Gerbrandy, and E. Pacuit. Merging frameworks for interaction: DEL and ETL. In D. Samet, editor, *Proceedings of Theoretical Aspects of Rationality and Knowledge (TARK 2007)*, 2007.
- [11] J. v. Benthem and E. Pacuit. The Tree of Knowledge in Action: Towards a Common Perspective. In I. H. G. Governatori and Y. Venema, editors, *Advances in Modal Logic*, volume 6. College Publications, 2006.
- [12] J. v. Benthem, J. van Eijck, and B. P. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [13] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.
- [14] G. Bonanno. Belief revision in a temporal framework. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory: Proceedings of LOFT'06*, Texts in Logic and Games, pages 43–50. Amsterdam University Press, 2006.
- [15] C. Bouillier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68(1):87–154, 1994.
- [16] D. de Jongh and F. Liu. Optimality, belief and preference. In S. Artemov and R. Parikh, editors, *Proceedings of the Workshop on Rationality and Knowledge*. ESSLLI, Malaga, 2006.
- [17] C. Dégrémont and J. Zvesper. Logique dynamique pour le raisonnement stratégique dans les jeux extensifs. In *Journées Francophones sur les Modèles Formels de l'Interaction (MFI'07)*, pages 61–73, 2007.
- [18] K. Eliaz and E. A. Ok. Indifference or indecisiveness? Choice-theoretic foundations of incomplete preferences. *Games and Economic Behavior*, 56(1):61–86, July 2006.
- [19] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, Cambridge, 1995.
- [20] K. Kelly. Ockham's razor, truth, and information. In J. van Benthem and P. Adriaans, editors, *Handbook of the Philosophy of Information*. 2008.
- [21] F. Liu. *Changing for the Better: Preference Dynamics and Agent Diversity*. Phd dissertation, ILLC Amsterdam, 2008.
- [22] S. Morris. The common prior assumption in economic theory. *Economics and Philosophy*, 11:227–253, 1995.
- [23] R. Parikh and R. Ramanujam. A knowledge based semantics of messages. *Journal of Logic, Language and Information*, 12(4):453–467, 2003.
- [24] R. Stalnaker. A theory of conditionals. In R. Stalnaker, W. L. Harper, and G. Pearce, editors, *Ifs: Conditionals, Belief, Decision, Chance and Time*. D. Reidel, Dordrecht, 1981.