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

Johan van Benthem[1,2]*
Cédric Dégremont[1]†

[1] ILLC Amsterdam

[2] Stanford University

johan.vanbenthem@uva.nl, cedric.uva@gmail.com

Analyzing the behavior of agents in a dynamic environment requires describing the evolution of their knowledge as they receive new information. But equally crucial are agents' beliefs over time, since most of our decisions and actions involve uncertainty, from going to work to selling shares. Beliefs, too, are information-based, and when refuted, they have to be revised in systematic ways.

These phenomena have been studied in many different formal frameworks, including, game theory [14, 5], belief revision theory [1], and formal learning theory [24, 22]. In this paper, however, we are concerned with two logic-based approaches. One are *dynamic logics for changing beliefs* that have been developed recently (van Benthem [7], Baltag and Smets [3]) using plausibility relations between worlds to represent agents' beliefs and conditional beliefs. An act of revision is then a single step of change in such a relation, triggered by some new incoming, hard or soft, information. Of course, such single steps can be iterated, leading to longer sequences. The other approach that we consider are *doxastic temporal logics* (cf. Halpern and Friedman [21], Bonanno [15]), representing time as a Grand Stage of possible histories where informational processes unfold.

Dynamic and temporal logic seem the two major logical paradigms for agency, and this paper is a contribution to clarifying their connections. In doing so, we do not operate in a void. Similar questions have been solved for knowledge in van Benthem and Pacuit [11], van Benthem, Gerbrandy, Hoshi and Pacuit [10], in the form of representation theorems showing how sequences of models produced by 'product update' in dynamic-epistemic logic form a special subclass of epistemic temporal models in the sense of

* The first author acknowledges a Weilun Visiting Professorship of Humanities, Philosophical Institute, Tsinghua University, Beijing, 2008–2009.

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

Fagin, Halpern, Moses and Vardi [20] and Parikh and Ramanujam [25]. In particular, these are the temporal models for agents endowed with Perfect Recall and 'No Miracles', learning by new observations only, possibly constrained by epistemic protocols. Our aim is to do the same for the dynamic doxastic logic of plausibility change by 'priority update', relating them to models of doxastic temporal logic. We will identify the crucial agent features behind dynamic doxastic belief revision, and position them inside the broader temporal setting. This is not just a simple generalization of the epistemic case, but the benefits are similar: comparability of frameworks, and interesting new research questions once they are merged. In this paper, we concentrate on the representation aspect. Further development of the merged theory of dynamic agents in a doxastic temporal language and logic is found in the follow-up paper van Benthem and Dégremont [9].

We start in the next section with basic terminology and background on earlier results for the epistemic setting. In Section 2 we introduce plausibility models that model static multi-agent doxastic situations. We then present the *dynamic* step by step approach to belief change (Section 3), in particular, defining priority update. Next, the global *temporal* approach to beliefs over time is presented in Section 4. In Section 5 we show how step by step priority updates of a doxastic model, perhaps constrained by a protocol, generate a doxastic temporal model. The key temporal doxastic properties that characterize priority updaters are then identified and motivated in Section 6. In section 7 we prove our main result linking the temporal and dynamic frameworks, for the special case of *total* pre-orders, and then in general in Section 8. We discuss some variations and extensions in Section 9. Finally, in a last section, we state our conclusions, mention follow-up questions involving formal languages and complete logics, and discuss possible applications to belief revision theory and learning theory.

1 Introduction: background results

Epistemic temporal trees and dynamic 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 [23, ch5]. We briefly state a recent version from [10], referring the reader to that paper for a proof, as well as generalizations and variations.

Definition 1.1 (Epistemic and Event Models, Product Update)**.**

- An *epistemic model* $\mathcal{M}$ is of the form $\langle W, (\sim_i)_{i \in N}, V \rangle$ where $W \neq \varnothing$, for each $i \in N$, $\sim_i$ is a binary relation[1] on W, and $V : Prop \to \wp(W)$.

[1] The $\sim_i$ are often taken to be equivalence relations, if only for convenience, but such options are orthogonal to our main results.

- An *event model* $\varepsilon = \langle E, (\sim_i)_{i\in N}, \texttt{pre}\rangle$ has $E \neq \varnothing$, and for each $i \in N$, $\sim_i$ is a relation on E. Finally, there is a precondition map $\texttt{pre} : E \to \mathcal{L}_{EL}$, where $\mathcal{L}_{EL}$ is the usual language of epistemic logic. We will consider some generalizations of this precondition language later.

- The *product update* of an epistemic model $\mathcal{M} = \langle W, (\sim_i)_{i\in N}, V\rangle$ with an event model $\varepsilon = \langle E, (\sim_i^\varepsilon)_{i\in N}, \texttt{pre}\rangle$ is the model $\mathcal{M}\otimes\varepsilon$ whose states are the pairs (w, e) such that w satisfies the precondition of the event e and whose epistemic relations are defined as:

$$(w, e) \sim_i' (w', e') \text{ iff } e \sim_i^\varepsilon e', w \sim_i w'$$

and whose valuation is defined by

$$(w, e) \in V(p) \text{ iff } w \in V(p)$$

An epistemic model describes what agents currently know, while product update creates the new epistemic situation after some informational event has taken place. Telling illustrations of the strength of this simple mechanism can be found in [2].

Next we turn to epistemic temporal models, introduced by [25] as a Grand Stage of unfolding informational events. In what follows, Σ^* is the set of finite sequences on any set Σ, which naturally forms a branching 'tree'.

Definition 1.2 (Epistemic Temporal Models)**.** An *epistemic temporal model* ('ETL model') $\mathcal{H}$ is a tuple $\langle \Sigma, H, (\sim_i)_{i\in N}, V\rangle$ with Σ a finite set of events, and $H \subseteq \Sigma^*$ closed under non-empty prefixes. For each $i \in N$, $\sim_i$ is a binary relation on H, and there is a valuation $V : Prop \to \wp(H)$.

Here the set of histories H functions as a *protocol* defining all admissible trajectories of an informational multi-agent process. While such ETL models are very general, many special constraints are possible. Some are the usual assumptions in epistemic logic, like having accessibility be an equivalence relation for $S5$-agents. But more important here are properties connecting epistemic accessibility with flow of time, defining general properties of an informational process and the agents participating in it. Such agents can have more idealized or more bounded powers of observation, memory, and other cognitive features. In particular, the following epistemic temporal properties drive the main representation theorem in [10]:

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

- **Propositional stability** Whenever h is a finite prefix of h', then h and h' satisfies the same proposition letters.

- **Synchronicity** Whenever $h \sim_i h'$, we have $len(h) = len(h')$.
- **Bisimulation Invariance** Whenever h and h' are epistemically bisimilar, we have $h'e \in H$ iff $he \in H$, for all events e.
- **Perfect Recall** Whenever $ha \sim_i h'b$, we also have $h \sim_i h'$.
- **Uniform No Miracles** Whenever $ga \sim_i g'b$ then, for every $h'a, hb \in H$, we also have $h'a \sim_i hb$. [2]

Dynamic-epistemic logic has borrowed once crucial idea from epistemic temporal logics. An *epistemic protocol* P maps states in an epistemic model to sets of finite sequences of pointed event models closed under taking prefixes. In general, this allows branching choices in a tree-like structure. This again defines the admissible runs of some informational process: not every observation may be available, or appropriate. More formally, let $\mathfrak{E}$ be the class of all *pointed event models*, having one 'actual event' marked. Then the set of protocols is $Prot(\mathfrak{E}) = \{\mathtt{P} \subseteq \mathfrak{E}^* \mid \mathtt{P} \text{ is closed under finite prefixes } \}$. Next comes the more general notion used in the recent literature:

Definition 1.4 (Local Protocols)**.** Given an epistemic model $\mathcal{M}$, a *local protocol* for $\mathcal{M}$ is a function $P : |\mathcal{M}| \to Prot(\mathfrak{E})$. In the particular case where the P is a constant function (mapping each world to the same set of sequences), we call the protocol *uniform*. Finally when the local protocol maps worlds to just a unique linear sequence of event models, we say that it is a *line protocol*.

To avoid technicalities, in this paper we state results with uniform line protocols. But our results generalize: see [10] for the epistemic case. Indeed, under suitable renaming of events, making different event models disjoint, line protocols even have the same expressive power as general branching protocols.

Now, given an epistemic model $\mathcal{M}$ as our initial situation, plus a uniform protocol P, we can define the resulting temporal evolution as an epistemic-temporal model $Forest(M, P) = \bigcup_{\vec{\varepsilon} \in P} M \otimes \vec{\varepsilon}$, the 'epistemic forest generated by' $\mathcal{M}$ through sequential application of the pointed event models in P using product update $\otimes$.

Finally, we can state what iterated dynamic-epistemic update means in the broader setting of epistemic-temporal logic:

Theorem 1.1 (van Benthem et al. [10])**.** *Let $\mathcal{H}$ be an arbitrary epistemic-temporal ETL model. The following two assertions are equivalent:*

[2] This says essentially that agents only get new information by acts of observation.

- *$\mathcal{H}$ is isomorphic to the temporal evolution $Forest(M, P)$ of some epistemic model $\mathcal{M}$ and uniform line protocol P,*
- *$\mathcal{H}$ satisfies Propositional Stability, Synchronicity, Bisimulation Invariance, Perfect Recall, and Uniform 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. [10] and [23, ch5] have details.

As stated in the Introduction, our paper extends this analysis to the richer setting of belief revision, where plausibility orders of agents evolve as they observe possibly surprising events. But to do so, we first need appropriate belief models, plus an appealing systematic revision mechanism.

Important remark about languages. Before moving on, it is important to stress one feature of the preceding representation theorem and results in its family. The precondition languages for event models should exactly match the notion of *bisimulation*. This means that the language should be invariant under such bisimulations, and also, that it should be strong enough to characterize a pointed model up to such bisimulations. Two technical observations follow from this:

1. To get the right definability, we should either restrict attention to finitely branching ETL models (as in [10]), or alternatively, let the precondition function of product models take values in an infinitary epistemic logic.
2. These theorems can be parametrized , in the epistemic case, and even more so, the doxastic setting. We stay at a semantic level in this paper, and state our results *up to language choice.* [3]

2 Plausibility models: static doxastic situations

As with knowledge, we first introduce static models that encode current prior (conditional) beliefs of agents. These carry a pre-order $\leq$ between worlds standing for a plausibility relation. Often this relation is taken to be total, but when we think of beliefs in terms of *multi-criteria decisions*, a pre-order allowing for incomparable situations may be all we get [19]. We will thus state our results for both total and arbitrary pre-orders.

[3] The issue of language choice returns briefly in Section 9 .The reader may also consult our companion paper [9] for an extensive discussion of syntactic issues, including other desiderata on the language, such as its expressive power for specifying the relevant properties of informational processes and the agents involved in them.

We write $a \simeq b$ ('indifference') if $a \leq b$ and $b \leq a$, and $a < b$ if $a \leq b$ and $b \not\leq a$. In what follows, $N = \{1, \ldots, n\}$ is a fixed finite set of agents.

The following definition is like the models in [14, 18, 3]:

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

$w \leq_i w'$ means that w is considered at least plausible as w' by agent i. Intuitively, the plausibility pre-orders encode current beliefs of agents. Here, we have taken them to be binary for convenience, but such relations can depend on states. An appealing intermediate case arises when we combine plausibility with an epistemic relation encoding 'hard information':

Definition 2.2 (Doxastic Epistemic Model)**.** A *doxastic plausibility model* $\mathcal{M} = \langle W, (\preceq_i)_{i \in N}, (\sim_i)_{i \in N}, V \rangle$ has $W \neq \varnothing$, for each $i \in N$, $\preceq_i$ is a pre-order on W and $\sim_i$ is binary equivalence relation on W, and $V : Prop \rightarrow \wp(W)$.

We write: $K_i[w] = \{v \in W \mid w \sim_i v\}$.

A belief operator for i is really necessity with respect to the most plausible states (i.e. the $\leq$-minimal elements) of an information partition. Though this paper does not discuss syntactic issues, it may help to state how models like these support a natural epistemic-doxastic language:

$$\mathcal{M}, w \Vdash K_i\varphi \quad \text{iff} \quad \forall v \;\; \text{with } v \in K_i[w] \text{ we have } \mathcal{M}, v \Vdash \varphi$$
$$\mathcal{M}, w \Vdash B_i\varphi \quad \text{iff} \quad \forall v \;\; \text{with } v \in Min_{\leq_i}(K_i[w]) \text{ we have } \mathcal{M}, v \Vdash \varphi$$

The setting also supports new modalities. In fact, the necessity operator for $\geq \cap \sim$ is a weakly defeasible (S4)-knowledge operator of 'safe belief' ([3]):

$$\mathcal{M}, w \Vdash \Box_i\varphi \quad \text{iff} \quad \forall v \;\; \text{with } v \leq_i w \text{ and } w \sim_i v \text{ we have } \mathcal{M}, v \Vdash \varphi$$

Remark: Alternatives. Some authors use models with just primitive plausibility relations. One can then define epistemic accessibility for a single agent as the union of that relation with its converse, accessing also less plausible worlds. We return to this perspective briefly in Subsection 9.3.

In what follows, we concentrate on pure plausibility models of our simplest sort, though our analysis will also work for more complex structures. We must now consider how such models evolve as agents observe events.

3 Dynamic Logics of Stepwise Belief Change (DDL)

Just like epistemic models, doxastic plausibility models change when appropriate triggering events are observed. It has become clear recently that a general mechanism for doing so works like the earlier product update ([3]). We start with the structures that describe complex doxastic events, crucially including the ways in which they appear to agents:

Definition 3.1 (Plausibility Event Model)**.** A *plausibility event model* ('event model', for short) ε is a tuple $\langle E, (\preceq_i)_{i \in N}, \texttt{pre}\rangle$ with $E \neq \varnothing$, each $\preceq_i$ is a pre-order on E, and $\texttt{pre} : E \to \mathcal{L}$, where $\mathcal{L}$ is the basic doxastic language.

As in the epistemic case, our analysis will work for various precondition languages for doxastic events. One specific choice is found at the end of Section 7. Combining perspectives, an 'epistemic plausibility event model' is a plausibility event model together with a collection of equivalence relations $(\sim_i)_{i \in N}$ on E.

In the following update rule, a new event itself comes with instructions as to how prior beliefs may be overridden. The principle is like that of 'Jeffrey Update' for probabilities: we follow the preferences of the plausibility event model, but if it leaves things open, we stick with prior preferences:

Definition 3.2 (Priority Update; [3])**.** *Priority update* of a plausibility model $\mathcal{M} = \langle W, (\preceq_i)_{i \in N}, V\rangle$ and an event model $\varepsilon = \langle E, (\preceq_i)_{i \in N}, \texttt{pre}\rangle$ produces the plausibility model $\mathcal{M} \otimes \varepsilon = \langle W', (\preceq'_i)_{i \in N}, V'\rangle$ defined as follows:

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

In the doxastic epistemic setting, Priority Update by epistemic plausibility event model combines the preceding mechanism with Product Update, i.e. it has one more clause:

- $(w, e) \sim'_i (w', e')$ iff $w \sim_i w'$ and $e \sim_i e'$

More motivation for this rule can be found in [3], and at the end of this section. First here is a concrete example.

As mentioned doxastic plausibility models are naturally combined with information partitions to describe scenarios involving both knowledge and beliefs. In this case Priority Update is applied to the plausibility ordering while product update is applied to information partition. We will discuss this issue in connection with the temporal models in Section 9. Let us for now present a concrete scenario that involves both knowledge and beliefs.

Reading the figures. In the following figures, the actual state (resp. event taking place) is the shaded one. Epistemic equivalence classes are represented by rectangles or ellipses. We use $<$ to display the strict plausibility ordering within such classes. Our example assumes that all agents have the same plausibility ordering. i believes φ at w is interpreted as φ holds in the i-most plausible states within i-information partition $K_i[w]$. An agent's beliefs at the actual state are thus displayed by an arrow from the actual state to the ones she considers most plausible, often just one. Thus, an arrow from x to y labelled by the agent *Enzo* means that y is the $\leq_e$-minimal state within $K_e[x]$. A similar convention applies to the event-model. Finally, we omit reflexive arrows throughout.

Example 3.3. Failed invitation. Céline and Enzo would like to invite to Denis to their Wii party. The party has been decided but none of them has informed Denis yet. Denis considers it a priori more plausible that no Wii party is taking place unless informed otherwise. This initial situation is common knowledge between Céline and Enzo. In the following figures, plain rectangles (or ellipses) will represent Denis' epistemic partition, dashed ones Enzo's and dotted ones Céline's. w and $\overline{w}$ are state names.

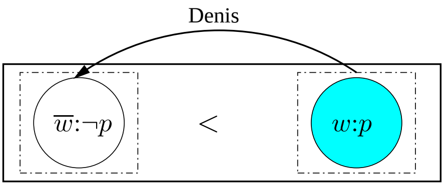

FIGURE 1. No Wii Party unless stated otherwise. Initial model.

The key event model. The telephone rings and Céline pick up the phone. Enzo hears part of the conversation and concludes that Céline is inviting Denis. In fact Céline is not on the phone with Denis. Céline think it was clear from the conversation that she was not talking to Denis.

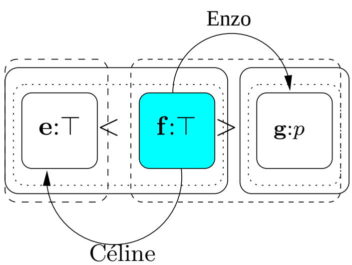

FIGURE 2. Event model of a misleading phone call.

We are now able to compute the new doxastic epistemic situation. The misunderstanding is now complete. In fact one can check that Enzo wrongly believes that it is now common knowledge between Céline and Denis that there is a Wii party while Céline wrongly believes that it is common belief between her and Enzo that Denis still does know about the Wii party and even that Denis still believes that there is no Wii party.

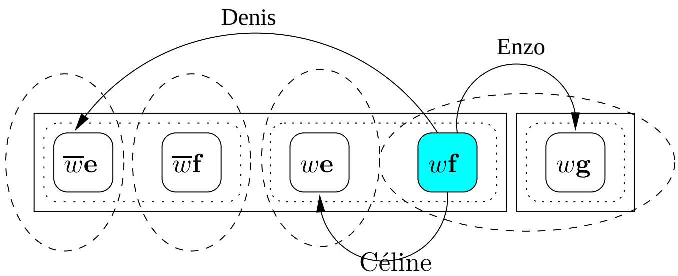

FIGURE 3. Product model of a misunderstanding.

Remark. Priority Update. In *AGM* style belief revision theory [1], new information is simply a new formula 'to be believed' by the agent. This allows for many different 'revision policies', from radical to conservative – a line also followed in a *DDL* setting by van Benthem [7]. It is important to appreciate that Priority Update is not just one such policy among many, but a *general mechanism* that can mimic many different policies depending on the richer structure of its triggers, viz. the plausibility event models [4]. If the event model has 'strong views', the update is radical, otherwise, the update remains conservative. Interestingly, this mechanisms also shifts the variety in belief revision away from fixed agent types, to case-by-case decisions: I can be radical with one input, and conservative with another.

We feel that a logic should describe a 'universal' mechanism, instead of a jungle of styles. This is why we have chosen Priority Update in this paper, leading to one representation that covers all special cases.

4 Doxastic Temporal Models: the global view

We now turn to the temporal perspective on multi-agent belief revision, as an informational process over time with global long-term features. The following models are a natural doxastic enrichment of the temporal *ETL* models of [25]. They are also close to the temporal doxastic models of [15, 21]. First the doxastic temporal models:

Definition 4.1 (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 $V : Prop \to \wp(H)$.

Doxastic Epistemic Temporal models ($DETL$ models for short) are Doxastic Temporal models extended by a collection of epistemic equivalence relations $(\sim_i)_{i\in N}$ on H.

Given some history $h \in H$ and event $e \in \Sigma$, we let he stand for the concatenation of h with e. Given that plausibility links are not themselves events, the model H may again be viewed as a 'forest', a disjoint union of event trees. We sometimes refer to $DoTL$ models as doxastic temporal forests. Figure 4 gives a concrete illustration of a practical setting with this abstract format. It display the evolution of a doctor's knowledge (dashed rectangles) and belief (diagnosis) - about what is wrong with her patient - as she performs medical tests and observe their positive or negative results (labelled edges). An arrow towards a state labelled `Environ` means that at this stage of the diagnostic process, the doctor think the patient's symptoms have an environmental cause. We omit reflexive and symmetric arrows.

Our models also gain concretenes by considering doxastic temporal languages interpreted on them. While these are the subject of our follow-up paper [9], we display a few truth conditions:

$$\begin{array}{lll}
\mathcal{H}, h \Vdash \langle e\rangle\varphi & \text{iff} & \exists h' \in H \ \text{ with } h' = he \ \text{ and } \ \mathcal{H}, h' \Vdash \varphi \\
\mathcal{H}, h \Vdash \Box_i\varphi & \text{iff} & \forall h\,' \text{ with } h' \leq_i h \text{ and } h \sim_i h' \text{ we have } \mathcal{H}, h' \Vdash \varphi \\
\mathcal{H}, h \Vdash K_i\varphi & \text{iff} & \forall h\,' \text{ with } h \sim_i h' \text{ we have } \mathcal{H}, h' \Vdash \varphi \\
\mathcal{H}, h \Vdash B_i\varphi & \text{iff} & \forall h\,' \text{ with } h' \in Min_{\leq_i} K_i[h] \text{ we have } \mathcal{H}, h' \Vdash \varphi
\end{array}$$

Dégremont [16] has comparisons of this framework with others, such as 'belief functions', or the models in [15].

5 From DDL Models To Doxastic Temporal Models

Now we come to the main question of this paper. Like AGM-style belief revision theory, Dynamic Doxastic Logic analyses one-step update scenarios. But unlike, AGM heory, it has no problem with iterating these updates to form longer sequences. Indeed let us put Example 3.3 together: Figure 5 looks like a doxastic epistemic forest model already. We will make this precies now, but as in the epistemic case, we need one more ingredient.

In many informational processes, such as learning, or belief revision in games, the information that agents receive may be highly constrained. Thus, there is crucial information in the set of admissible histories of the process, its 'protocol'. This notion can be defined formally just as before in Definition 1.4. Let $\mathfrak{E}$ be the class of all pointed plausibility event models. The set of *protocols* $Prot(\mathfrak{E}) = \{\mathtt{P} \subseteq (\mathfrak{E}^* \mid \mathtt{P} \text{ is closed under finite prefixes}\}$. What we need is again a slightly more flexible version:

If we take a look at the figure describing Example 3.3 we see that it really looks like an doxastic (epistemic) forest already. Actually we could continue

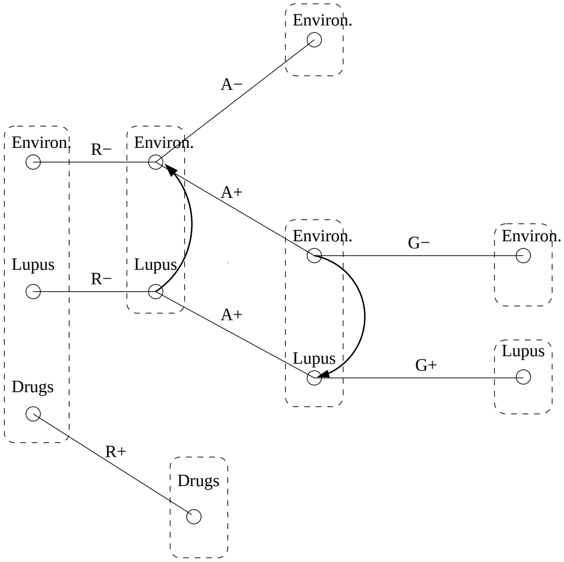

FIGURE 4. A medical investigation over time.

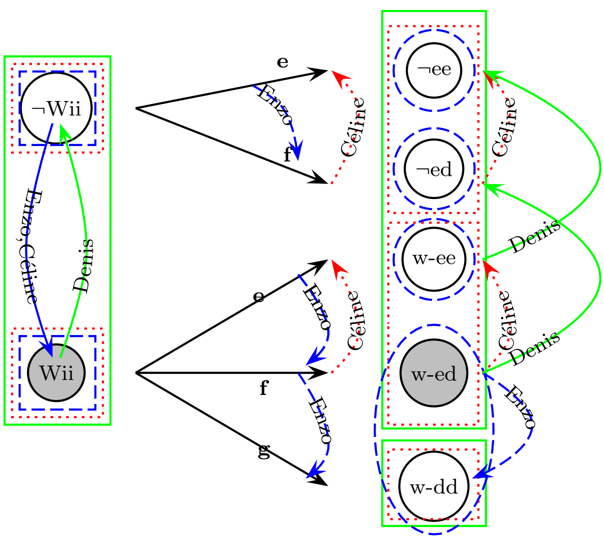

FIGURE 5. The Wii-party misunderstanding in temporal perspective.

the story, and the further updates would generate a larger forest. More generally, priority update of a plausibility model according to a protocol generates a doxastic temporal forest.

Definition 5.1 (Doxastic Protocols)**.** Given a doxastic plausibility model $\mathcal{M}$, a *local protocol* for $\mathcal{M}$ is a function $P : |\mathcal{M}| \to Prot(\mathfrak{E})$. If P is a constant function, the protocol is called *uniform*. When P maps states to a linear nested sequence of event models, we call it a *line protocol*.

In line with Section 1, we state our main theorems in terms of uniform line protocols, leaving variations and extensions to [16].

Iterated Priority Update of a doxastic plausibility model according to a uniform line protocol P generates a doxastic temporal forest model. We construct the forest by induction, starting with the doxastic plausibility model and and then checking which events can be executed according to the precondition and to the protocol. Finally the new plausibility order is updated at each stage according to Priority Update. Since Priority Update describes purely doxastic, non-ontic change, the valuation stays the same as in the initial model. (For ways of adding real factual change, cf. [12].) For simplicity, we write $P(w) = \vec{\varepsilon}$ where $\vec{\varepsilon}$ is a finite sequence of event models.

Definition 5.2 (*DoTL* model generated by a sequence of updates)**.** Each initial plausibility model $\mathcal{M} = \langle W, (\preceq_i)_{i\in N}, V\rangle$ and each sequence of plausibility sevent models $\vec{\varepsilon} = (\varepsilon_j)_{j\in\omega}$ where $\varepsilon_j = \langle E_j, (\preceq_i^j)_{i\in N}, \mathtt{pre}_j\rangle$ yields a *generated DoTL plausibility model* $\langle \Sigma, H, (\leq_i)_{i\in N}, \mathbf{V}\rangle$ as follows:

- Let $\Sigma := \bigcup_{i=1}^{m} E_i$.
- Let $H_1 := W$, and for each $1 < n \leq m$, let $H_{n+1} := \{(we_1 \dots e_n) \mid (we_1 \dots e_{n-1}) \in H_n$ and $\mathcal{M} \otimes \varepsilon_1 \otimes \dots \otimes \varepsilon_{n-1} \Vdash \mathtt{pre}_n(e_n)\}$.

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 $h \leq_i h'$.
- Finally, set $wh \in \mathbf{V}(p)$ iff $w \in V(p)$.

Our task is to identify just when a doxastic temporal model is isomorphic to the 'forest' thus generated by a sequence of priority updates. In particular, this will uncover the key doxastic properties of agents assumed in this belief revision mechanism.

6 Crucial Frame Properties for Priority Update

We first get a few more general properties of our information process out of the way. The first of these merely says that in that process, the facts of the world do not change, only agents' beliefs about it:

Definition 6.1. 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', h and h' satisfy the same proposition letters.

Note that this can be generalized to include real world change. Next comes a basic property of the events that we allowed as revision triggers:

6.1 Bisimulation Invariance

Since the aim of this notion is to guarantee the existence of pre-conditions behind events in some modal language. Depending on the language parameter we choose, one has to choose the corresponding Bisimulation notion. As mentioned in Section 1 we will state our results up to language choice, therefore we give an abstract definition of bisimulation below. We will however give a concrete example of language instantation when stating a corollary of our result for doxastic epistemic models.

Let τ be a finite collection of binary relations $\langle R_1, \dots, R_n\rangle$ on $H \times H$.

Definition 6.2 (τ-Bisimulation)**.** Let $\mathcal{H}$ and $\mathcal{H}'$ be two $DoTL$-models based on the same alphabet Σ. A relation $Z \subseteq H \times H'$ is a τ-Bisimulation if, for all $h \in H$, $h' \in H'$ and all $R_i \in \tau$

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

(back) If hZh' and hR_ij, then there is a $j' \in H'$ with jZj' and $h'R'_ij'$,

(forth) If hZh' and $h'R'_ij'$, then there is a $j \in H$ with jZj' and hR_ij.

If Z is a τ-bisimulation and hZh', we say h and h' are τ-bisimilar.

Definition 6.3 (τ-Bisimulation Invariance)**.** A $DoTL$ model $\mathcal{H}$ satisfies *τ-bisimulation invariance* if, for all τ-bisimilar histories $h, h' \in H$, and all events e, $h'e \in H$ iff $he \in H$.

Note that these definitions apply also to $DETL$ models. Here is an example. $(\sim_i \cap \leq_i)_{i\in N}$-*Bisimulation Invariance* will leave all formulas of the basic doxastic language with safe belief invariant, and hence our earlier preconditions for events. If we want these preconditions to be richer, then we need more clauses in the bisimulation – and the same is true if we want the bisimulation to preserve explicit temporal formulas involving events.

6.2 Agent-Oriented Properties

Now we come to the relevant agent properties. These depend on single agents i only, and hence we will drop agent labels and prefixes "for each $i \in N$" for the sake of clarity. Also, in what follows, when we write ha for events a, we assume that $ha \in H$.

Definition 6.4. Let $\mathcal{H} = \langle \Sigma, H, (\leq_i)_{i \in N}, V \rangle$ be a $DoTL$ model. $\mathcal{H}$ satisfies:

- **Synchronicity** Whenever $h \leq h'$, we have $len(h) = len(h')$.

This says intuitively that agents have a correct belief about the exact stage the process is in. The following two properties trace the belief revising behavior of priority-updating agents more precisely:

- **Preference Propagation** if $ja \leq j'b$, then $h \leq h'$ implies $ha \leq h'b$.
- **Preference Revelation** If $jb \leq j'a$, then $ha \leq h'b$ implies $h \leq h'$.

What do the latter properties say? In the earlier epistemic representation theorems, the corresponding properties of Perfect Recall and No Miracles described observational agents with ideal memory, the two basic features behind the Product Update rule. Likewise, our new properties express the two basic features 'hard-wired into' the Priority Update rule, its 'radicalism' and its 'conservatism'. Preference Propagation says that, if the last-observed events ever allowed a plausibility preference, then they always do – or stated contrapositively, if they ever 'over-rule' an existing plausibility, then they always do. This reflects the first radical clause in the definition of Priority Update. Next, Preference Revelation says that when an agent has no strict plausibility preference induced by two observed events, then she will go with her prior plausibility. This reflects the second, conservative clause in Priority Update. As we have said before, this is a qualitative description of a 'Jeffrey-style' updating agent in a probabilistic setting.

7 The Main Representation Theorem

Now we prove our main result relating DDL and DTL models, both with total orders.

Theorem 7.1. *Let $\mathcal{H}$ be any doxastic-temporal model with a total plausibility order. Then the following two assertions are equivalent:*

1. *There exists a total plausibility model $\mathcal{M}$ and a sequence of total plausibility event models $\vec{\mathcal{E}}$ such that $\mathcal{H}$ is isomorphic to the forest generated by the Priority Update of $\mathcal{M}$ by the sequence $\vec{\mathcal{E}}$.*
2. *$\mathcal{H}$ satisfies Propositional Stability, Synchronicity, Bisimulation Invariance, Preference Propagation, and Preference Revelation.*

Proof. **Necessity** ($1 \implies 2$)**.** We show that the given conditions are 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 either clause in the definition of priority update that $a \leq b$ (2). Now assume that $h \leq h'$ (3). It follows from (2), (3) and again priority update that $ha \leq h'b$.

Preference Revelation. Assume that $jb \leq j'a$ (1). It follows from 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 so $h \leq h'$ (7).

Note that we did not make use of totality in this direction of the proof.

Sufficiency ($2 \implies 1$)**.** Given a *DoTL* model $\mathcal{M}$ satisfying the stated conditions, we show how to construct a matching doxastic plausibility model and a sequence of event models.

Construction Here is the initial plausibility model $\mathcal{M}_0 = \langle W, (\preceq_i)_{i\in N}, \hat{V}\rangle$:

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

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

- $E_j := \{e \in \Sigma \mid \text{there is a history } he \in H \text{ with } len(h) = j\}$
- Set $a \preceq_i^j b$ iff there are $ha, h'b \in H$ such that $len(h) = len(h) = j$ and $ha \leq_i h'b$.
- For each $e \in E_j$, let $\mathtt{pre}_j(e)$ be the formula that characterizes the set $\{h \mid he \in H \text{ and } len(h) = j\}$. By general modal logic, our condition of Bisimulation Invariance guarantees that there is such a formula. Again as mentioned at the end of Section 1 this sentence may be an infinitary one in general (if we don't assume the doxastic temporal models to be finitely branching). We give a concrete instantiation when we discuss the epistemic doxastic corollary of our result.

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

Claim 7.1 (Correctness)**.** Let $\leq$ be the plausibility relation in the given doxastic temporal model. Let $\preccurlyeq^F_{DDL}$ be the plausibility relation in the forest model induced by priority update over the just constructed plausibility model $\mathcal{F}$ and the constructed sequence of event models. We have:

$$h \leq h' \text{ iff } h \preccurlyeq^F_{DDL} h'.$$

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 $\mathcal{M}_0$. Now comes the induction step:

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 $h_1a \preccurlyeq^F_{DDL} h_2b$, whatever relationship held between h_1 and H_2 in $\mathcal{F}$.

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 original doxastic temporal model $\mathcal{M}$. It follows by the inductive hypothesis that $h_1 \preccurlyeq^F_{DDL} h_2$. But then, since a and b are indifferent by (2) and (3), priority update gives us $h_1a \preccurlyeq^F_{DDL} h_2b$.

From $Forest(DDL)$ **to** $DoTL$**.** Now let $h_1a \preccurlyeq^F_{DDL} h_2b$. Again we follow the two clauses in the definition of priority update:

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*(this is the only place where we use this property), $h_1a \leq h_2b$.

Case 2: $a \simeq b$ (4) and $h_1 \preccurlyeq^F_{DDL} h_2$. 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 $h_1a \leq h_2b$. Q.E.D.

Remark. Corollary for the Doxastic Epistemic case. We get a representation result for the doxastic epistemic case as an immediate corollary from Theorem 7.1 and Theorem 1.1. Moreover we give a concrete instantiation of this corollary by choosing the language of Safe Belief. In the result below we refer to Priority Update as the results of applying Product update to the epistemic relations and Priority Update to the plausibility orderings.

Corollary 7.2. *Let $\mathcal{H}$ be any doxastic epistemic temporal model with a total plausibility order. Then the following two assertions are equivalent:*

1. *There exists a total an epistemic plausibility model $\mathcal{M}$ and a sequence of total epistemic plausibility event models $\vec{\varepsilon}$ taking preconditions in the modal language of Safe Belief such that $\mathcal{H}$ is isomorphic to the forest generated by the Priority Update of $\mathcal{M}$ by the sequence $\vec{\varepsilon}$.*

2. *$\mathcal{H}$ satisfies Propositional Stability, Synchronicity, Perfect Recall, Uniform No Miracles, $(\sim_i \cap \leq_i)_{i \in N}$ - Bisimulation Invariance, Preference Propagation, and Preference Revelation.*

This result shows how to find, inside the much broader class of all doxastic temporal models, those whose plausibility pattern was produced by a systematic priority update process.

8 Extension to arbitrary pre-orders

The preceding result generalizes to the general case of pre-orders, allowing incomparability. Here we need a new notion that was hidden so far:

Definition 8.1 (Accommodating Events)**.** Two events $a, b \in \Sigma$ are pairwise *accommodating* if, for all $ga, g'b$: $(g \leq g' \leftrightarrow ga \leq g'b)$, i.e. a, b preserve and anti-preserve plausibility.

We can now define our new condition on doxastic-temporal models:

- **Accommodation** Events a and b are accommodating in the sense of Def. 8.1 if both $ja \leq j'b$ and $ha \not\leq h'b$ for some j, j', h, h'.

Accommodation is a uniformity property saying that, if two events allow both plausibility orders for histories, then they are always 'neutral' for determining plausibility order.This property only comes into its own with pre-orders allowing incomparable situations:

Fact 8.2. If $\leq$ is a total pre-order and $\mathcal{H}$ satisfies Preference Propagation and Preference Revelation, then $\mathcal{H}$ satisfies Accommodation.

Proof. Assume that $ja \leq j'b$ (i) and $ha \not\leq h'b$. By totality, the latter implies $hb \leq h'a$ (ii). Now let $g \leq g'$. By Preference Propagation and (i), $ga \leq g'b$. Conversely, assume that $ga \leq g'b$. By Preference Revelation, (i) and (ii), we have $g' \leq g$. Q.E.D.

We can also prove a partial converse without assuming totality:

Fact 8.3. 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. Q.E.D.

Finally, an easy counter-example shows that, even with $\leq$ total:

Fact 8.4. Accommodation does not imply Preference Revelation.

Proof. Take a simplest model where the following holds: $h'b \simeq ha \simeq j'a \simeq jb$ and $h' < h \simeq j' \simeq j$. Q.E.D.

With arbitrary pre-orders we need to impose Accommodation:

Theorem 8.1. *Let $\mathcal{H}$ be any doxastic-temporal model with a plausibility pre-order. Then the following two assertions are equivalent:*

1. *There exists a plausibility model $\mathcal{M}$, and a sequence of plausibility event models $\vec{\varepsilon}$ such that $\mathcal{H}$ is isomorphic to the forest generated by the Priority Update of $\mathcal{M}$ by the sequence $\vec{\varepsilon}$.*
2. *$\mathcal{H}$ satisfies Bisimulation Invariance, Propositional Stability, Synchronicity, Preference Revelation, and Accommodation.*

By Fact 8.3, Accommodation also gives us Preference Propagation.

Proof. **Necessity of the conditions.** ($1 \implies 2$) Checking the conditions in the Section 7 did not use totality. So we focus on the new condition:
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 $a \not< b$ (4). By definition, (2) with (4) imply that $a \simeq b$ (5). Now assume that $g \leq g'$ (6). It follows from (5), (6) and priority update that $ga \leq g'b$. The other direction is similar.
Sufficiency of the conditions. ($2 \implies 1$) Given a $DoTL$ model, we again construct a DDL plausibility model plus a sequence of event models:
Construction The plausibility model $\mathcal{M}_0 = \langle W, (\preceq_i)_{i \in N}, \hat{V}\rangle$ is as follows:

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

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

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

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

Claim 8.5 (Correctness)**.** Let $\leq$ be the plausibility relation in the doxastic temporal model $\mathcal{M}$. Let $\preccurlyeq^F_{DDL}$ be the plausibility relation in the forest $\mathcal{F}$ induced by successive priority updates of the plausibility model by the sequence of event models we just constructed. We have:

$$h \leq h' \text{ iff } h \preccurlyeq^F_{DDL} h'.$$

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 $\mathcal{M}_0$. Now for the induction step, with the same simplified notation as earlier:

From $DoTL$ **to** $Forest(DDL)$**.** We distinguish two cases.

Case 1. $ha \leq h'b, h \leq h'$**.** By the inductive hypothesis, $h \leq h'$ implies $h \preccurlyeq^F_{DDL} h'$ (1). Since $ha \leq h'b$, it follows by the construction that $a \leq b$ (2). Then, by (1), (2) and priority update, we get $ha \preccurlyeq^F_{DDL} h'b$.

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 $ha \preccurlyeq^F_{DDL} h'b$.

From $Forest(DDL)$ **to** $DoTL$**.** We again distinguish two cases.

Case 1. $ha \preccurlyeq^F_{DDL} h'b, h \preccurlyeq^F_{DDL} h'$**.** By the definition of priority update, $ha \preccurlyeq^F_{DDL} h'b$ implies that $a \leq b$ (1). There are two possibilities.

Case 1.1: The special clause of the construction has been used, and a, b are accommodating (2). By the inductive hypothesis, $h \preccurlyeq^F_{DDL} h'$ implies $h \leq h'$ (3). But (2) and (3) imply that $ha \leq h'b$.

Case 1.2: Clause (1) holds because for some $ja, j'b \in H$ in the $DoTL$ model, $ja \leq j'b$ (4). By the inductive hypothesis, $h \preccurlyeq^F_{DDL} h'$ implies $h \leq h'$ (5). Now it follows from (4), (5) and Preference Propagation that $ha \leq h'b$.

Case 2. $ha \preccurlyeq^F_{DDL} h'b, h \not\preccurlyeq^F_{DDL} h'$**.** Here is where we put our new accommodation clause to work. Let us label our assertions: $h \not\preccurlyeq^F_{DDL} h'$ (1) and $ha \preccurlyeq^F_{DDL} h'b$ (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 must have $ha \leq h'b$. Q.E.D.

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 an initial plausibility model.

9 More extensions and variations of the theorem

Several further scenarios can be treated in the same manner. In particular, it is easy to combine the epistemic analysis in Section 1 with ours to include agents having both *knowledge and belief*. Here are three more directions:

9.1 From uniform to local protocols

So far we have considered uniform line protocols. We have already suggested that line protocols are powerful enough to mimic branching protocols through renaming of events, and then taking a disjoint union of all branching alternatives. But uniformity is a real restriction, and it can be lifted. *Local protocols* allow the set of executable sequences of pointed events models forming our current informational process to vary from state to state. Indeed, agents need not even know which protocol is running. As was done in [10] for the epistemic case, we can still get our representation theorems, by merely dropping the condition of Bisimulation Invariance. While this seems a simple move, local protocols drastically change the complete dynamic-doxastic logic of the system (cf. [9] and [16] for details).

9.2 Languages and bisimulations

As we have noted in Section 4, our doxastic-temporal models support various languages and logics. These will be pursued in [9], but we do make a few points here. One is that complete doxastic-temporal logics for the above special model classes will have valid principles reflecting the reduction axioms of dynamic-doxastic logic. In fact, these doxastic-temporal *correspond* to Preference Propagation and Preference Revelation in the sense of modal correspondence theory. Thus, our structural analysis of priority-updating agents extends to the level of valid reasoning.

Proposition 9.1. *The following law is sound for plausibility change:*

$$\langle \varepsilon, \mathbf{e}\rangle\langle\leq_i\rangle\varphi \leftrightarrow (\mathtt{pre}(e) \wedge (\langle\leq_i\rangle \bigvee\{\langle\mathbf{f}\rangle\varphi \ : \ e\simeq_i f\} \vee \mathtt{E}\bigvee\{\langle\mathbf{g}\rangle\varphi \ : \ e <_i g\}))$$

Analogies with recursion axioms One can understand the following formal axioms in their original format with existential modalities by analogy with the dynamic-doxastic recursion axiom for priority update just given:

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

$$\mathtt{E}\langle a\rangle\langle\leq_i\rangle\langle b^{-1}\rangle\top \rightarrow (((\langle\leq_i\rangle\langle b\rangle p \wedge \langle a\rangle q) \rightarrow \langle a\rangle(q \wedge \langle\leq_i\rangle p)$$

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

$$\mathtt{E}\langle b\rangle\langle\leq_i\rangle\langle a^{-1}\rangle\top \rightarrow (\langle a\rangle\langle\leq_i\rangle(p \wedge \langle b^{-1}\rangle\top) \rightarrow \langle\leq_i\rangle\langle b\rangle p)$$

In fact, a doxastic-temporal language has two main purposes in our setting: (a) stating 'local' preconditions for events, (b) specifying 'global' properties of the temporal evolution of the current process. As is well-known [13] a choice of language here corresponds with a choice of a semantic invariance relation, usually some weaker or stronger variant of *bisimulation.* In stating our results, we kept to the simplest version, of a modal bisimulation adequate for the static language of safe belief. But this can be varied, and one can also have stronger notions of bisimulation, respecting more structure, that work for more expressive doxastic languages, or for combined doxastic-epistemic languages. If we add epistemic structure, the bisimulation also needs back-and-forth clauses for the intersection of the epistemic accessibility and doxastic plausibility relations.

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 *back* and *forth* clauses should not only apply to $\leq_i$ and $\sim_i$ separately, but also to $\leq_i \cap \sim_i$. (Indeed $\cap$ is not safe for bisimulation.) And things get even more complicated if we allow temporal operators in our languages (cf. [10]). We do not want to commit to any specific choice here, since the choice of a language seems orthogonal to our main concerns. Even so, we will discuss formal languages in the next section, taking definability of our major structural constraints as a guide.

9.3 Alternative model classes

Finally, we mentioned in Section 4 that one can also work with a primitive plausibility relation that merges epistemic indistinguishability and doxastic plausibility. This is done for Priority Update in [3], and we indicate briefly the notions involved in the corresponding results:

Definition 9.1. The priority update of a *unified* plausibility model $\mathcal{M} = \langle W, (\trianglelefteq_i)_{i\in N}, V\rangle$ and a $\trianglelefteq$-event model $\varepsilon = \langle E, (\trianglelefteq_i)_{i\in N}, \mathtt{pre}\rangle$ is the unified plausibility model $\mathcal{M} \otimes \varepsilon = \langle W', (\trianglelefteq'_i)_{i\in N}, V'\rangle$ constructed as follows:

- $W' = \{(w, e) \in W \times E \mid \mathcal{M}, w \Vdash \mathtt{pre}(e)\}$,
- $(w, a) \trianglelefteq'_i (w', b)$ iff either 1. $a \trianglelefteq_i b$, $b \ntrianglelefteq a$ and $w \trianglelefteq w' \vee w' \trianglelefteq w$ or 2. $a \trianglelefteq_i b$, $b \trianglelefteq a$ and $w \trianglelefteq w'$,
- $V'((s, e)) = V(s)$.

Here are our basic temporal doxastic agent properties in this setting:

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

[9, 16] show how these conditions drive a general representation theorem similar to the one in Section 7 and 8.

10 Conclusion

Agents that update their knowledge and revise their beliefs leave an epistemic and doxastic 'trace' over time of epistemic and doxastic relations. We have determined the special constraints that capture agents operating with the 'local updates' of dynamic doxastic logic. This took the form of 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.

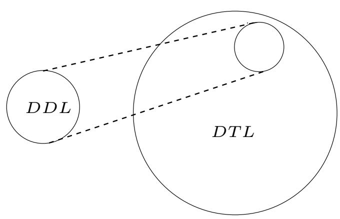

FIGURE 6. DDL inside DTL

Thus, we have determined an area where the idealized belief changers of dynamic-doxastic logic live. This identification becomes even stronger through its axiomatization in suitable doxastic-temporal languages, as explored in our follow-up paper [9].

As for open problems, this paper has already indicated several technical issues along the way. Here we list a few more:

- *Languages and logics.* There are many issues of expressive power of different languages over our models and their complexity effects (cf. [11] for the epistemic case). In particular, belief revision for individual agents has a natural companion of *belief merge* for groups of agents, and various update processes that might create common beliefs. These create complications in the epistemic case [12] and are even more challenging here. For all such languages, there are also natural questions of complete axiomatization. Indeed, we have proved axiomatic completeness for the 'protocol logic' of revision processes in [9, 16]. This is analogous to the epistemic theory of observation and conversation protocols initiated in [10].

- *Agent diversity.* Beyond the realm of priority-updating agents studied in this paper, there are also other zones of interest inside DTL models, including processes involving agents with memory bounds, that may revise beliefs in different ways, and may even see their knowledge fade into belief, and finally, oblivion.

- *Related Approaches.* A comparison of our 'constructive' DDL-inspired approach to DTL universes with the more abstract AGM-style postulational approach of [15] remains to be made.

- *Concrete scenarios: game theory and learning theory.* We intend to take our analysis to knowledge and belief dynamics in the concrete setting of *extensive games*, following up on the initial studies [8, 5]. It would also be of interest to link our doxastic-temporal models with more concrete *learning scenarios*. In the temporal universes of formal learning theory [22], learning rules are belief revision or plausibility update rules producing the new hypothesis after the next observed event. Dégremont and Gierasimczuk [17] is a first systematic attempt at connecting this with DDL and DTL models.

Acknowledgements.

We thank Tomohiro Hoshi and Eric Pacuit for useful discussions. Also, Olivier Roy, Hans van Ditmarsch, Nina Gierasimczuk, Valentin Goranko, Ram Ramanujam, Yu Junhua and Jonathan Zvesper gave helpful comments. Finally, the extensive comments of two referees helped us produce a more focused and hopefully clearer paper.

References

[1] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[2] Alexandru Baltag and Lawrence S. Moss. Logics for Epistemic Programs. *Synthese*, 139(2):165–224, 2004.

[3] Alexandru Baltag and Sonja 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*, TLG, pages 11–24. AUP, 2006.

[4] Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundation of Game and Decision Theory (LOFT7)*, volume 3 of *TLG*, pages 13–60. AUP, 2008.

[5] Alexandru Baltag, Sonja Smets, and Jonathan A. Zvesper. Dynamic Rationality in Extensive Games. In Johan van Benthem and Eric Pacuit, editors, *Proceedings of Workshop on Logic and Intelligent Interaction, ESSLLI*, 2008.

[6] Johan van Benthem. Games in Dynamic Epistemic Logic. *Bulletin of Economic Research*, 53(4):219–248, 2001.

[7] Johan van Benthem. Dynamic logics for belief change. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.

[8] Johan van Benthem. *Logical Dynamics of Information Flow and Interaction.* Cambridge University Press, Cambridge, to appear.

[9] Johan van Benthem and Cédric Dégremont. Multi-agent belief dynamics: bridges between dynamic doxastic and doxastic temporal logics. Technical report, ILLC, University of Amsterdam, 2008. `http://staff.science.uva.nl/~cdegremo/`.

[10] Johan van Benthem, Jelle Gerbrandy, Tomohiro Hoshi, and Eric Pacuit. Merging frameworks for interaction: DEL and ETL. *Journal of Philosophical Logic*, 2009.

[11] Johan van Benthem and Eric Pacuit. The Tree of Knowledge in Action: Towards a Common Perspective. In I. Hodkinson G. Governatori and Y. Venema, editors, *Advances in Modal Logic*, volume 6. College Publications, 2006.

[12] Johan van Benthem, Jan van Eijck, and Barteld P. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.

[13] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.

[14] Oliver Board. Dynamic interactive epistemology. *Games and Economic Behavior*, 49:49–80, 2004.

[15] Giacomo Bonanno. Axiomatic characterization of the agm theory of belief revision in a temporal logic. *Artif. Intell.*, 171(2-3):144–160, 2007.

[16] Cédric Dégremont. *Logics of Information and Interaction*. PhD dissertation, ILLC Amsterdam, to appear.

[17] Cédric Dégremont and Nina Gierasimczuk. Can doxastic agents learn? on the temporal structure of learning. Technical Report PP-2009-14, ILLC, 2009.

[18] Hans van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Synthese*, 147:229–275, 2005.

[19] Kfir Eliaz and Efe A. Ok. Indifference or indecisiveness? Choice-theoretic foundations of incomplete preferences. *Games and Economic Behavior*, 56(1):61–86, July 2006.

[20] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, Cambridge, 1995.

[21] Nir Friedman and Joseph Y. Halpern. Modeling belief in dynamic systems, part I: foundations. *Artificial Intelligence*, 95(2):257–316, 1997.

[22] Kevin T. Kelly. The learning power of belief revision. In *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, pages 111–124, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.

[23] Fenrong Liu. *Changing for the Better: Preference Dynamics and Agent Diversity*. PhD dissertation, ILLC Amsterdam, 2008.

[24] Eric Martin and Daniel Osherson. *Elements of Scientific Inquiry*. MIT Press, Cambridge, 1998.

[25] Rohit Parikh and Ramaswamy Ramanujam. A knowledge based semantics of messages. *Journal of Logic, Language and Information*, 12(4):453–467, 2003.