PDL restrictions, belief revision and abduction in DEL

Thomas Bolander
Technical University of Denmark

This paper is published in the festschrift for Andreas Herzig on the occasion of his 65th birthday. The content is heavily inspired by the many interesting and enlightening discussions I have had with him over the years. Andreas taught me many things and made me aware of many interesting open problems in (dynamic) epistemic logic. Andreas was the first to make me aware that standard dynamic epistemic logic (DEL) does not admit belief revision, something that I will make formally precise below, and then discuss to which extent can be solved using plausibility models and model restrictions (the latter to avoid state size explosions). Andreas was also the first to introduce me to the wonders of propositional dynamic logic (PDL), which will be used to define the relevant model restrictions. The final topic of the paper is abduction: I will sketch a method for only keeping track of the most plausible worlds, and rely on abduction in case of surprising action outcomes.

1 Preliminaries on modal logic

Given a binary relation R on a set S, we write xRy when $(x, y) \in R$, and define, for any $x, y \in S$, $xR := \{(x', y') \in R \mid x' = x\}$ and $Ry := \{(x', y') \in R \mid y' = y\}$. Similarly for $X, Y \subseteq S$, we define $XR := \{(x', y') \in R \mid x' \in X\}$ and $RY := \{(x', y') \in R \mid y' \in Y\}$. A *modal similarity type* τ is in this paper a finite set of *modal operators* (boxes) $\Box_1, \Box_2, \ldots$ [5]. Given a set P of *atomic propositions* (*atoms*), the

Thanks to: Andreas Herzig for continuous inspiration; Sonja Smets and Malte Velin for discussions on abduction and initial joint work on abduction in plausibility models that eventually lead to this work; Jerome Lang, who originally proposed the weakening principle introduced in Section 5.1.

modal language over τ and P, denoted $ML(\tau, P)$, is given the following grammar, where $p \in P$ and $\Box \in \tau$: $\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi$. Other propositional connectives as well as $\bot$ and $\top$ are defined by abbreviation in the usual way. The dual of $\Box$ is denoted $\Diamond$.

Definition 1.1. *A* (Kripke) model *for $ML(\tau, P)$ is $M = (W, R, V)$, where W is a non-empty set of* worlds*, $R : \tau \to 2^{W \times W}$ assigns to each modal operator $\Box \in \tau$ a binary relation $R_\Box$ on W, and $V : W \to 2^P$ is a* valuation function *assigning to each world the set of atoms true there. For a set of* designated *worlds $W' \subseteq W$, we call (M, W') a* pointed model*. When $W' = \{w\}$, we often write (M, w) for $(M, \{w\})$.*

Satisfaction is defined as follows, for pointed models $(M, W') = ((W, R, V), W')$ and $\varphi \in ML(\tau, P)$, with standard propositional clauses:

$$(M, W') \models \varphi \text{ iff } (M, w) \models \varphi \text{ for all } w \in W'$$
$$(M, w) \models p \text{ iff } p \in V(w), \text{ for } p \in P$$
$$(M, w) \models \Box\varphi \text{ iff } (M, wR_\Box) \models \varphi, \text{ for } \Box \in \tau$$

When $(M, W) \models \varphi$, φ is *universally true* in M, denoted $M \models \varphi$. When $M \models \varphi$ for all models M, φ is *valid*, denoted $\models \varphi$. According to the semantics, evaluating $\Box\varphi$ at a world w amounts to evaluating φ at the subset of worlds $wR_\Box$. The $\Box$ modality asks us to change our perspective from the current world w to an alternative set of worlds $wR_\Box$. Thus we can think of $\Box$ as a modality for picking out a subset of worlds of the original model. There is also a more drastic way of picking out a subset of worlds in a model: we restrict the model to that subset.

Definition 1.2. *Given a Kripke model $M = (W, R, V)$ and $W' \subseteq W$, we define the* restriction *of M to W' as $M|W' = (W', R', V')$ with $R'_\Box = R_\Box \cap (W')^2$ for all $\Box \in \tau$ and $V'(w) = V(w)$ for all $w \in W'$. For formulas φ, we define $M|\varphi := M|\{w \in W \mid M, w \models \varphi\}$.*

2 Program modalities for model restrictions

In *public announcement logic* (*PAL*) [12], we have for each formula φ a *public announcement* modality $[\varphi]$ with the following semantics:

$$(M, w) \models [\varphi]\psi \text{ iff } (M, w) \models \varphi \text{ implies } (M|\varphi, w) \models \psi$$

Note the difference in how $M, w \models \Box\psi$ and $M, w \models [\varphi]\psi$ are evaluated:

1. $M, w \models \Box\psi$: The $\Box$ modality picks out a subset of worlds using $R_\Box$, and then evaluates ψ in those worlds.
2. $M, w \models [\varphi]\psi$: The $[\varphi]$ modality picks out a subset of worlds using φ, prunes away all other worlds, and finally evaluates ψ in w.

Are there natural hybrids between the two approaches? What if we first pick out a subset of worlds using a formula φ and then evaluate ψ in those worlds? Well, this we can easily express even without modalities, as it simply corresponds to evaluating whether $\varphi \to \psi$ is universally true in the model. How about the opposite hybrid, where we pick a subset of worlds using $R_\Box$, prune away all other worlds, and finally evaluate ψ? This would give a modality $[\Box]$ with a semantics defined by:

$$(M, w) \models [\Box]\psi \text{ iff } w \in wR_\Box \text{ implies } (M|wR_\Box, w) \models \psi$$

What would be the logic of such a modality and what could it potentially be used for? This is one of the things we will explore in the following.

First, let us try to generalise things a bit. As in the *logic of communication and change* (*LCC*) [15], it seems natural to allow compositions of the relations $R_\Box$. This is simple to achieve by allowing *propositional dynamic logic* (*PDL*) programs over these relations [10]. Further, a difference between the $\Box$ and $[\varphi]$ modalities is that the $\Box$ modality also changes the points of evaluation, which is not possible with the $[\varphi]$ modality. It seems natural to consider hybrids also allowing us to use PDL programs to define the points of evaluation. This leads us to:

Definition 2.1. *The language R-PDL*(τ, P) *(the R in R-PDL is for* restricting PDL*, as programs are used to restrict models) is given by the following grammar, where* $p \in P$ *and* $\Box \in \tau$*:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid [\pi, \pi]\psi$$
$$\pi ::= \Box \mid \varphi? \mid \pi_1; \pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^*$$

The π are standard programs in PDL [10]. In the formula $[\pi_1, \pi_2]\psi$, the first program, π_1, is used to pick out the subset of worlds that the model will be restricted to, and the second program, π_2, is used to select the points of evaluation (the designated worlds). More formally, we define as follows, with mutual recursion betweeen $\models$ and R_π:

$$(M, w) \models [\pi_1, \pi_2]\psi \text{ iff } wR_{\pi_2} \subseteq wR_{\pi_1} \text{ implies } (M|wR_{\pi_1}, wR_{\pi_2}) \models \psi$$

$$R_{\varphi?} = \{(w, w) \mid M, w \models \varphi\} \qquad R_{\pi_1;\pi_2} = R_{\pi_1} \circ R_{\pi_2}$$
$$R_{\pi_1 \cup \pi_2} = R_{\pi_1} \cup R_{\pi_2} \qquad R_{\pi^*} = (R_\pi)^*$$

When Π is a set of programs, we use Π as shorthand for $\cup_{\pi \in \Pi} \pi$. So τ^* is short for $(\cup_{\Box \in \tau} \Box)^*$ (recall that τ is the set of modalities in the logic). For any worlds x, y, $(x, y) \in R_{\tau^*}$ iff y is reachable from x (by any sequence of the $R_\Box$ relations). Since truth of a formula φ in a world w only depends on the submodel reachable from w, we get $(M, w) \models \varphi$ iff $(M|R_{\tau^*}, w) \models \varphi$. These two models, $\mathcal{M}$ and $\mathcal{M}|R_{\tau^*}$, are semantically indistinguishable (modally equivalent), and will be identified. Thus $\varphi \leftrightarrow [\tau^*, \top?]\varphi$ is valid for all φ. We define $\boxdot\varphi := [\tau^*, \Box]\varphi$.

When the modalities in τ are knowledge modalities, K_a, the programs above are the same as in *epistemic PDL* (*E-PDL*) [15] (except we would write their atomic program a as K_a). The formulas of E-PDL are then also the same as ours, except their modality is of the form $[\pi]$ whereas ours is $[\pi, \pi]$. The semantics of the E-PDL modality is given by: $M, w \models [\pi]\psi$ iff $M, wR_\pi \models \psi$. Thus R-PDL extends E-PDL, since the modality $[\pi]$ of E-PDL can in R-PDL be equivalently represented as $[\tau^*, \pi]$. This also implies that if the modalities in τ are knowledge modalities, then we can express *common knowledge that* φ in R-PDL by $[\tau^*, \tau^*]\varphi$: first restrict the model to the reachable worlds using τ^*, then evaluate φ in all those worlds (using again τ^* to reach them). R-PDL extends PAL as well, as the PAL formula $[\varphi]\psi$ can in our logic be equivalently expressed as $[\tau^*;\varphi?, \top?]\psi$: first restrict the model to the reachable worlds where φ is true, then evaluate ψ in the original world.

2.1 Expressivity

Two formulas φ_1 and φ_2 are *equivalent*, denoted $\varphi_1 \equiv \varphi_2$, if they are true in the same pointed models. If two languages L_1 and L_2 are interpreted over the same class of models, then L_2 is *at least as expressive* as L_1 if for every $\varphi_1 \in L_1$ there is a $\varphi_2 \in L_2$ such that $\varphi_1 \equiv \varphi_2$ [16]. We say that L_2 is *more expressive* than L_1 if L_2 is at least as expressive as L_1, but L_1 is

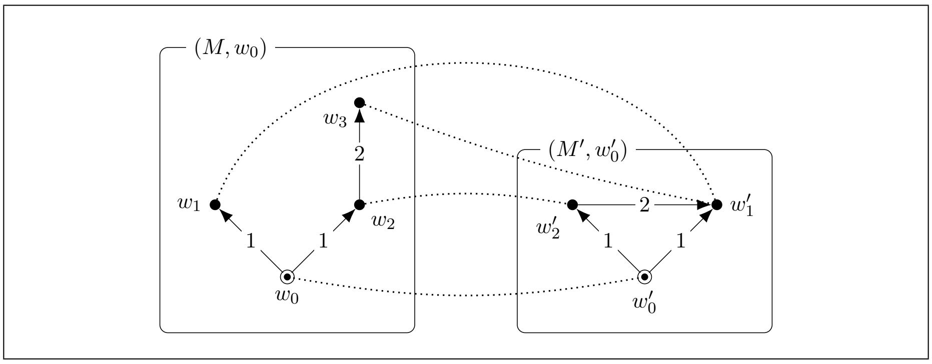

Figure 1: Two pointed models (M, w_0) (left) and (M', w'_0) (right) that are bisimilar, but distinguishable in R-PDL. Each node is a world, with designated worlds highlighted. A solid edge labelled i from w to v means that $(w, v) \in R_{\Box_i}$. The dotted edges denote a bisimulation relation between (M, w_0) and (M', w'_0).

not at least as expressive as L_2 [16]. Since we have just shown the E-PDL modality $[\pi]$ and the PAL modality $[\varphi]$ to be representable in R-PDL, R-PDL is at least as expressive as both. Is it more expressive? Yes! To prove this, it suffices to find $\varphi \in$ R-PDL and pointed models (M, w_0) and (M', w'_0) s.t. $(M, w_0) \models \varphi$ and $(M', w'_0) \not\models \varphi$, but all formulas of E-PDL and PAL have the same truth value in both models. Let $\varphi = [\Box_1^*, \top?]\Box_1\Box_2\bot$, where $\tau = \{\Box_1, \Box_2\}$. From Figure 1 we get:

1. $(M, w_0) \models \varphi$ *iff* $(M|w_0 R_{\Box_1^*}, w_0 R_{\top?}) \models \Box_1\Box_2\bot$ *iff*
   $(M|\{w_0, w_1, w_2\}, w_0) \models \Box_1\Box_2\bot$ *iff*
   true, since there is no $1, 2$-path from w_0 in $M|\{w_0, w_1, w_2\}$.

2. $(M', w'_0) \models \varphi$ *iff* $(M'|w'_0 R_{\Box_1^*}, w'_0 R_{\top?}) \models \Box_1\Box_2\bot$ *iff*
   $(M', w'_0) \models \Box_1\Box_2\bot$ iff
   false, since there is a $1, 2$-path from w'_0 to w'_1 in M'.

This shows that φ distinguishes the two models. At the same time, the figure highlights a *bisimulation relation* [5] between the two models. As truth in both E-PDL and PAL is preserved under bisimulations [15], the same formulas of those languages are true in both models. Thus:

Proposition 2.1. *R-PDL is more expressive than E-PDL and PAL.*

Note that R-PDL is more expressive than PAL despite being based on the same idea of restricting models to subsets of worlds. The essential difference is that PAL selects subsets using a formula, i.e., always selects a modally definable subset. This is not always true with PDL programs. Fig. 1 illustrates it: No formula can distinguish w_1 from w_3 in M, so it would be impossible to construct a public announcement that deletes w_3 without also deleting w_1. However, in R-PDL we *can* delete w_3 without deleting w_1, using the program $\Box_1^*$. We will here not discuss possible axiomatisations of R-PDL, but turn to some of its intended applications.

3 Adding update models

The logic LCC mentioned above is achieved by extending E-PDL with a modality $[U, e]$ for each *update model* U and *event* e in U [15].

Definition 3.1. *An* update model *for a language L extending $ML(\tau, P)$ is $U = (E, R, pre, post)$ where E is a finite, non-empty set of* events*; $R : \tau \to 2^{W\times W}$ is as in Definition 1.1; $pre : E \to L$ assigns* preconditions *to events; and $post : E \to (P \to L)$ assigns* postcondition functions *to events. For a set of* designated events *$E' \subseteq E$, we call (U, E') a* pointed update model*. When $E' = \{e\}$, we often write (U, e) for $(U, \{e\})$. An update model is* ontic *if $post(e)(p) \not\equiv p$ for some e, p;* purely epistemic *if $post(e)(p) \equiv p$ for all e, p; and* purely ontic *if $pre(e) \equiv \top$ for all e.*

Definition 3.2. *The* product update *of a Kripke model $M = (W, R, V)$ with an update model $U = (E, R, pre, post)$ is the Kripke model $M \otimes U = (W', R', V')$ where:*[1]

$$\begin{aligned} W' &= \{(w, e) \mid w \in W \text{ and } e \in E\} \\ R'_\Box &= \{((w_1, e_1), (w_2, e_2)) \mid (w_1, w_2) \in R_\Box \text{ and } (e_1, e_2) \in R_\Box\} \\ V'((w, e)) &= \{p \in P \mid M, w \models post(e)(p)\} \end{aligned}$$

The language DR-PDL is achieved by adding the following clause to the grammar of R-PDL: $\varphi ::= [U, E']\varphi$, where (U, E') is a pointed update model for DR-PDL (note the mutual recursion in allowing U to be an update model for the same language DR-PDL). The semantics is:

$$(M, w) \models [U, e]\varphi \text{ iff } (M, w) \models pre(e) \text{ implies } (M \otimes U, (w, e)) \models \varphi \quad (1)$$

[1] We are using R both for the relations on worlds and the relations on events. The context will reveal which one we refer to.

Define $[U, E']\varphi := \bigwedge_{e \in E'} [U, e]\varphi$. In R-PDL, we added PDL programs to modify Kripke models by restricting them to a certain subset of worlds and/or point out a certain set of designated worlds. It seems natural to consider similar "model modifiers" on update models. Define the *restriction* of an update model U to a set of events $E' \subseteq E$, denoted $U|E'$, similarly to Definition 1.2: the submodel generated by the events in E'. Given a pointed update model $(U, E') = ((E, R, pre, post), E')$ and two programs π_1 and π_2, we then define: $(U, E', \pi_1, \pi_2) := (U|R_{\pi_1}, E'R_{\pi_2})$. The two programs π_1 and π_2 are applied to update models the same way they applied to Kripke models: π_1 restricts the update model to a subset of events, and π_2 picks out the designated events. Since (U, E', π_1, π_2) is just another pointed update model, we already have a modality for it in the language, so adding π_1 and π_2 here doesn't add to the expressivity of DR-PDL, but can still provide convenient notation. For instance, consider the notion of an *associated local action* of an agent i for a given update model (action) (U, E') [6]. This was defined by closing the designated events under $\sim_i$, but can now be expressed as (U, E', τ^*, K_i).

We use *standard dynamic epistemic logic* (*standard DEL*) to refer to any language achieved by expanding a modal language $ML(\tau, P)$ with a modality $[U, e]$ having a product update semantics as defined above.

4 Belief revision

Call a modality $\Box$ *dynamic* if there is a pointed model (M, w) s.t. the semantic clause for $(M, w) \models \Box\varphi$ evaluates φ in a model distinct from M (or, rather, distinct from the submodel of M generated by w). Our $[\pi_1, \pi_2]$ modality is dynamic whenever π_1 is not equivalent to τ^*. Specifically, public announcements $[\varphi]$ are dynamic whenever $\varphi \not\equiv \top$ (corresponding to $[\tau^*;\varphi?, \top?]$ with $\varphi \not\equiv \top$).

Definition 4.1. *Let L be a language containing a belief modality B. We say that a dynamic modality $\Box$ of L* admits (propositional) belief revision *if there exists a propositional formula φ and a pointed model (M, w) for L such that $(M, w) \models B\varphi \wedge \neg\Box B\varphi$.*

The condition expresses that initially, in w, φ is believed true, but this belief is not preserved by the dynamic update $\Box$.

Proposition 4.1. *If U is a purely epistemic update model in standard DEL, then $[U, e]$ does not admit belief revision.*

Proof. Let φ be any propositional formula and (M, w) any pointed model with $(M, w) \models B\varphi$. We need to show that $(M, w) \models [U, e]B\varphi$. Suppose to achieve a contradiction that $(M, w) \not\models [U, e]B\varphi$. This means $(M, w) \models pre(e)$ and $(M \otimes U, (w, e)) \not\models B\varphi$. From this we get that $M \otimes U$ contains a world (w', e') such that $(w, e)R_B(w', e')$ and $(M \otimes U, (w', e')) \models \neg\varphi$. By definition of product update, we have wR_Bw', and since U is purely epistemic and φ is propositional, then also $(M, w') \models \neg\varphi$. This implies that $(M, w) \not\models B\varphi$, which is a contradiction. □

One can resort to ontic update models to achieve belief revision in the sense of Definition 4.1. However, conceptually, postconditions model ontic change, e.g. flipping a switch, not belief revision concerning a static world. Thus standard DEL doesn't admit belief revision in any natural way. This is a well-known problem and criticism of standard DEL for modelling beliefs. Specifically, if an agent believes p and a (truthful) public announcement of $\neg p$ is made, she will afterwards not consider any worlds possible, and will hence believe $\bot$ (if initially Bp, then p is true in all the R_B-accessible worlds, and all of these are deleted by the announcement). This of course even holds if $\neg p$ is directly sensed, e.g. if $p =$ *'there is milk in the fridge'*, and initially she believes p, but then opens the fridge to discover $\neg p$. In standard DEL, it is impossible for her to then adopt her beliefs and start believing $\neg p$ instead. This issue is generalised by Prop. 4.1 above: Not only does public announcements (or direct sensing) not admit belief revision, *no* purely epistemic updates admit it in standard DEL. The problem is sometimes referred to as the problem that agents "can't recover from false beliefs". The standard reply in DEL is to move to *plausibility models* [3], considered next.

5 Plausibility models

Given a set X and a relation $\leq$ on X, the set of *least* elements of X is $\text{Min}_{\leq} X := \{x \in X \mid x \leq x' \text{ for all } x' \in X\}$. A *well-preorder* on X is a reflexive, transitive relation $\leq$ s.t. every non-empty subset has least elements, i.e., for all non-empty $Y \subseteq X$, $\text{Min}_{\leq} Y \neq \emptyset$. We write $x < y$

when $x \leq y$ and $y \not\leq x$, and $x \simeq y$ when both $x \leq y$ and $y \leq x$. Our well-preorders will encode *plausibility orders* on sets of worlds W, where $w \leq v$ expresses that w is *at least as plausible as* v, $w < v$ that w *is more plausible than* v, and $w \simeq v$ that they are *equiplausible* [3]. The most plausible worlds in W are the elements of $\mathrm{Min}_{\leq} W$. Due to space limitations, and to keep the exposition simple, we restrict attention to single-agent plausibility models over modalities $\tau = \{B^n \mid n \in \mathbb{N} \cup \{\infty\}\}$, where $B^n\varphi$ reads: "φ is believed to degree n". We introduce the standard belief operator B by abbreviation: $B\varphi := B^0\varphi$.

Definition 5.1. *A Kripke model $M = (W, R, V)$ is called a* plausibility model *wrt. a relation $\leq$ on W if $\leq$ is a well-preorder and, for all $w \in W$, $wR_{B^0} = \mathrm{Min}_{\leq} W$, $wR_{B^{n+1}} = wR_{B^n} \cup \mathrm{Min}_{\leq}(W \setminus wR_{B^n})$, and $B^\infty = W$. A* plausibility update model *(also simply called an* action*) wrt. $\leq$ is an update model $U = (E, R, post, pre)$ with the same conditions on R_{B^n}.*

Since plausibility update models are standard update models, they would also be subject to the no-belief-revision result of Prop. 4.1 if it hadn't been for their non-standard product update, defined next. We use $\leq$ both for relating worlds and events, letting context disambiguate.

Definition 5.2. *Let $M = (W, R, V)$ be a plausibility model wrt. $\leq$ and $U = (E, R, pre, post)$ a plausibility update model wrt. $\leq$. The* action-priority update *of M with U is the plausibility model $M \otimes_{ap} U = (W', R', V')$ where W' and V' are as in Definition 3.2 and R' is defined as in Definition 5.1 wrt. $\leq'$ defined by:*

$$(w, e) \leq' (v, f) \text{ iff } e < f \text{ or } (e \simeq f \text{ and } w \leq v) \tag{2}$$

The semantics of the dynamic modality $[U, e]$ can now be defined as in (1), except we replace $\otimes$ by $\otimes_{ap}$. This logic *does* admit belief revision:

Example 5.1. *Let (M, w_2) be a plausibility model with $W = \{w_1, w_2\}$, $V(w_1) = p$, $V(w_2) = \emptyset$ and $w_1 < w_2$. Then $(M, w_2) \models Bp$. Consider the action (U, e) with $E = \{e\}$, $pre(e) = \neg p$ and $post(e)(p) = p$. Then $M \otimes_{ap} U$ only contains the world (w_2, e) satisfying $\neg p$. Hence $(M, w_2) \not\models [U, e]Bp$, proving that $[U, e]$ admits belief revision (cf. Definition 4.1).*

Equation (2) is the *action-priority update rule*: when deciding which of the updated worlds (w, e) or (v, f) is more plausible, the ordering on the events take precedence, with the intuition that the "incoming changes of beliefs" (the action) take precendence over the past beliefs [3]. We will refer to any alternative definition of $\leq'$ in Def. 5.2 as an *update rule*, as long as $\leq'$ is a well-preorder (which is true for action-priority update [8]). Each update rule leads to a specific type of product update.

Using our PDL operators, we get convenient and compact notation for several existing operators and notions within plausibility models. The X operator in single-agent plausibility planning [1] is simply $[K, \top?]$, and the *appearance* of (U, e) to agent a [3] is (U, e, τ^*, B_a).

5.1 Alternatives to action-priority update

We now explore whether the action-priority update rule (2) is the only natural update rule. Example 5.1 showed that plausibility updates admit belief revision, and the example didn't rely on the update rule. Thus, any other logic we might achieve by changing the update rule will also admit belief revision. We define a minimal condition on update rules:

$$\text{PRES(ERVATION): If } w \star v, e \star f \text{ then } (w,e) \star' (v,f), \text{ for } \star \in \{<, \leq, \simeq\}$$

It states that the direction of the plausibility order is directly inherited from the order on the worlds and events when these agree. A critical aspect of defining an update rule is how to relate (w, e) and (v, f) when the order on worlds and events disagree. There are 3 possibilities:

$$\begin{aligned} \text{STATE-PRIORITY}: &\quad \text{If } w < v \text{ and } f < e, \text{ then } (w,e) <' (v,f) \\ \text{ACTION-PRIORITY}: &\quad \text{If } w < v \text{ and } f < e, \text{ then } (w,e) >' (v,f) \\ \text{WEAK(ENING)}: &\quad \text{If } w < v \text{ and } f < e, \text{ then } (w,e) \simeq' (v,f) \end{aligned}$$

Proposition 5.1. *No update rule satisfies* PRES *and* WEAK.

Proof. Consider the product update in Fig. 2, where the plausibility order is induced by PRES and WEAK. The updated model is not transitive, as we have $(w_2, e_2) \leq' (w_3, e_3) \leq' (w_1, e_1)$, but not $(w_2, e_2) \leq' (w_1, e_1)$. Hence $\leq'$ is not a well-preorder and cannot be an update rule. □

This shows that only STATE-PRIORITY and ACTION-PRIORITY can satisfy PRES. Using similar examples as in Fig. 2, only modifying the direction of the edges, we can show that when requiring PRES and either

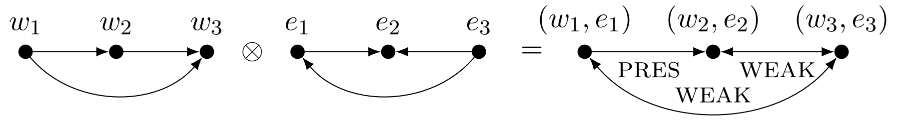

Figure 2: A product update under the assumption of PRESERVATION and WEAKENING. We have omitted the valuation, pre- and post-conditions, but assume they are chosen s.t. $pre(e_i)$ is only satisfied in w_i, $i = 1, 2, 3$. Each edge in the updated model is labelled by the principle that determined it.

STATE-PRIORITY or ACTION-PRIORITY, then there is a unique transitive order $\leq'$ satisfying the requirements. Furthermore, also for STATE-PRIORITY, this becomes a well-preorder (by a proof symmetric to the one for action-priority [8]). Thus:

Proposition 5.2. *The only two update rules that satisfy* PRES *are the action-priority rule (2) and the following dual* state-priority rule*:*

$$(w, e) \leq' (v, f) \text{ iff } w < v \text{ or } (w \simeq v \text{ and } e \leq f). \quad (3)$$

Definition 5.3. *The operator of* state-priority update*, denoted $\otimes_{sp}$, is defined as in Definition 5.2, except replacing the update rule (2) by (3).*

Could state-priority update have interesting applications? Yes, it might be relevant in epistemic planning [6], where update models are used to reason about possible futures, and there it might make sense to give precedence to your current beliefs over how your future actions might eventually affect them. Some initial explorations in this direction already exist [1, 11]. Another possible application is within abduction:

Example 5.2. *Let $M = (\{w\}, R, V)$ be a plausibility model with $V(w) = p$, where $p =$* 'my bike is parked in front of the supermarket'. *Let U_1 be the action of entering the supermarket, U_2 the action of buying groceries, and U_3 the action of returning outside. Suppose that $M \otimes_{ap} U_1 \otimes_{ap} U_2 \otimes_{ap} U_3 \models \neg p$, i.e., when I return outside, I discover that my bike was stolen. We can assume that U_1 has events $\{e_{no}, e_{ab}\}$ and U_2 has events $\{e'_{no}, e'_{ab}\}$, where e_{no} and e'_{no} are the* normal *events in which the bike is not stolen, and e_{ab} and e'_{ab} are the* abnormal *events where it is stolen. Supposing that I didn't consider the event of the bike being* *stolen most plausible at any of the previous time points, i.e., $e_{no} < e_{ab}$ and $e'_{no} < e'_{ab}$, when will I think the bike was stolen? By the principle of* chronological minimisation *[14, 13, 4], I would conclude that the bike was most likely stolen at the latest possible time, i.e., e'_{ab} happened. However, chronological minimisation is not consistent with action-priority update. In action-priority, the ordering on worlds in $M \otimes_{ap} U_1 \otimes_{ap} U_2$ is $(w, e_{no}, e'_{no}) < (w, e_{ab}, e'_{no}) < (w, e_{no}, e'_{ab}) < (w, e_{ab}, e'_{ab})$, i.e., I consider it more plaussible that it was stolen during U_1 than U_2. This corresponds to* reverse chronological minimisation*, not chronological minimisation [2]. To achieve chronological minimisation, we would use $\otimes_{sp}$.*

6 Abduction based on plausibility models

In this section, we only consider actions with propositional pre- and post-conditions, i.e., with no occurences of the B^i modalities. An issue with plausibility models (and all other existing approaches to DEL) is that to guarantee that an agent doesn't end up believing $\bot$, the agent needs to keep track of any event that *could* potentially have happened, since that event might later turn out to have been the real one. Say that you leave your office for a few moments while a colleague is typing on your keyboard, adding a definition to your joint paper. While you're away, you can still hear each keystroke, but not identify which keys are pressed. Say plausibility models and plausibility update models are used to keep track of the dynamics of your beliefs. After having heard n keystrokes, your updated plausibility model will contain all 78^n possible combinations of keystrokes. Why? Well, to ensure that you don't end up believing $\bot$ when you return to the office and see what's on the screen. Plausibility models rely on observations being modelled as restrictions of the model (as with public announcements), i.e., whatever an agent might potentially observe in a state has to be already represented as one of the worlds of that state. This makes plausibility models far less attractive in applications such as epistemic planning [6, 1] and robots tracking beliefs in human-robot interaction [9, 7], as focusing on beliefs rather than knowledge then doesn't help to tame the combinatorial explosion of keeping track of all the events that might happen but are not observed.

Let us here take an initial step towards addressing the issue by considering the possibility of only preserving the most plausible worlds and events when updating. Our PDL extensions now come in handy. The modality $[B, B]$ (recall Def. 2.1) will select only the most plausible worlds of a model and make them all designated. For instance, referring to the model (M, w_2) of Ex. 5.1, we have $(M, w_2) \models \neg B^\infty p \wedge [B, B] B^\infty p$: In M, the agent only believes p to degree 0, but after we prune away all not-most-plausible worlds, p is believed to arbitrary degree. Applying the $[B, B]$ modality corresponds to the agent accepting her beliefs as facts, i.e., only preserving the worlds representing her (degree 0) beliefs. Accepting your beliefs as facts is of course a potentially dangerous strategy. Consider the action (U, e) of Ex. 5.1, representing the announcement of $\neg p$. If an agent has accepted a false belief in p as a fact, she cannot incorporate this announcement, as there are no $\neg p$ worlds left in her model (this is the "no recovery from false beliefs" problem we discussed earlier). Formally, $(M, w_2) \models [B, B][U, e]\bot$: If $\neg p$ is announced after the agent has accepted her beliefs as facts, she will believe anything. We adapt an existing notion of surprise [8] to the current setting:

Definition 6.1. *An action (U, E') is called a* surprise *in a plausibility model (M, W') if $(M, W') \models [B, B][U, E', B, B]\bot$.*

Thus an action is a surprise if accepting your current beliefs as facts—including your beliefs about the action—leads to a degenerate updated model (one that has no designated worlds). As long as no surprises happen, beliefs are preserved even if we prune all not-most-plausible worlds, as we now show.

Proposition 6.1. *Let (M, W') be a plausibility model and suppose that the update model (U, E') is not a surprise in (M, W'). Then for any propositional formula φ, $(M, W') \models [U, E']B\varphi \leftrightarrow [B, B][U, E', B, B]B\varphi$.*

Proof. To prove the equivalence, it suffices to prove that $M \otimes_{ap} U$ and $M|(W'R_B) \otimes_{ap} U|(E'R_B)$ have the same most plausible worlds (since φ is propositional). Since (U, E') is not a surprise, we have $(M, W') \not\models [B, B][U, E', B, B]\bot$, and hence there exists $w \in W'$ such that $(M, w) \not\models [B, B][U, E', B, B]\bot$, and thus $(M|wR_B, wR_B) \not\models [U, E', B, B]\bot$. This implies the existence of a world $w' \in wR_B$ such that $(M|wR_B, w') \not\models [U, E', B, B]\bot$, from which we get the existence of an event $e' \in E'R_B$

with $(M|wR_B, w') \models pre(e')$. Since $w' \in wR_B$ and $e' \in E'R_B$, both w' and e' are most plausible in their respective models. This combined with $(M|wR_B, w') \models pre(e')$ gives that (w', e') is a world of both $M \otimes_{ap} U$ and $M|(W'R_B) \otimes_{ap} U|(E'R_B)$ (preconditions are propositional). It is also most plausible in both, as otherwise there would have to be either a world strictly more plausible than w' or an event strictly more plausible than e', according to action-priority update. Let now (w'', e'') be an arbitrary most plausible world in $M \otimes_{ap} U$. We will show that it is also most plausible in $M|(W'R_B) \otimes_{ap} U|(E'R_B)$. As (w', e') and (w'', e'') are both most plausible, we get $(w'', e'') \simeq (w', e')$. By action-priority update, also $w'' \simeq w'$ and $e'' \simeq e'$, so w'' is also most plausible in M, and hence exists in $M|(W'R_B)$. Since $e'' \simeq e'$, e'' is most plausible in U, and hence (w'', e'') exists in $M|(W'R_B) \otimes_{ap} U|(E'R_B)$ (preconditions are propositional). Since $w'' \simeq w'$ and $e'' \simeq e'$, also $(w'', e'') \simeq (w', e')$ in $M|(W'R_B) \otimes_{ap} U|(E'R_B)$, implying that (w'', e'') is most plausible in that model. The other direction is similar. □

This shows that as long as there are no surprises, whenever an action occurs, the agent can update her beliefs by first pruning away all non-most plausible worlds and events, and update afterwards. Say that in the keyboard typing example, you consider it most plausible that your colleague is repeatedly pressing the *a* key (so that when hearing a keystroke, you consider the event "typing *a*" most plausible). Then after having heard n keystrokes, you would still only have a model of size 1, and believe that the screen now shows "aaaaaaaaaaa..." (we assume that it does not). Only when you come back and look at the screen, you will be surprised (the action of sensing the content of the screen will be a surprise according to Def. 6.1), and need to reconstruct the less plausible events that you omitted. This is a case of abduction.

Let a plausibility model $M = (W, R, V)$, a plausibility update model $U = (E, R, pre, post)$ and an R-PDL program π be given. We introduce $M|\pi$ as an abbreviation for $M|(WR_\pi)$ and $U|\pi$ as an abbreviation for $U|(ER_\pi)$. Thus e.g. $M|B$ is an abbreviation for $M|(WR_B)$, which is the same as $M|\operatorname{Min}_{\leq} W$. We call $(M|B) \otimes_{ap} (U|B)$ a *most plausible update* of M with U. It is what we get by deleting all non-most-plausible worlds and events before performing the update. Note that when (U, E) is not a surprise in (M, W), it follows from Proposition 6.1 that the agent has the same beliefs in $(M|B) \otimes_{ap} (U|B)$ as in $M \otimes_{ap} U$. Say an action sequence $(U_1, E'_1), (U_2, E'_2), \dots, (U_n, E'_n)$ is executed. As long as no surprises occur, the agent can then simply perform most plausible updates and still preserve her beliefs, hence potentially avoiding the computational explosion of keeping track of all possible events. For instance, in case an action happens and nothing is sensed, the agent might use a principle of *epistemic inertia* to take the *skip* event as the single most plausible event. Using a most plausible update, this implies that the existing model will simply be preserved (giving us the least computationally expensive update operation possible).

At some point a surprise might of course occur, say the ith action in the sequence is a surprise. Opposite the situation in standard DEL, the agent now actually *does* have the possibility to regain a consistent state representation. She simply has to "unprune" some of the pruned worlds or events, i.e., include also points that were originally not considered most plausible. But which ones? Chronological minimisation would ask us to include less plausible events towards the end of the action sequences, whereas reverse chronological minimisation would ask us to include less plausible events from the beginning of the action sequence. In Ex. 5.2, discovering the stolen bike in U_3 is a surprise after the action sequence U_1, U_2, since $M \otimes_{ap} U_1 \otimes_{ap} U_2$ has a single most plausible world (w, e_{no}, e'_{no}), and in this, the bike was not stolen. Correspondingly, performing most plausible updates, we would get a model $M \otimes_{ap} U_1|\{e_{no}\} \otimes_{ap} U_2|\{e'_{no}\}$ with a single world (w, e_{no}, e'_{no}). In this model, applying U_3 would lead to an empty model (since U_3 includes an announcement of $\neg p$, and p is true in (w, e_{no}, e'_{no})). This calls for abduction, which could e.g. either be to replace the computation of the most plausible update $(M|B) \otimes_{ap} (U_1|B) \otimes_{ap} (U_2|B)$ with $M \otimes_{ap} (U_1|B) \otimes_{ap} (U_2|B^1)$ (chronological minimisation on events) or $(M|B) \otimes_{ap} (U_1|B^1) \otimes_{ap} (U_2|B)$ (reverse chronological minimisation on events). More generally, one could iteratively increase the indices on the degree of belief modalities in either lexicographic, reverse lexicographic or some other monotonic order until the latest action is no longer a surprise, i.e., one has successfully performed abduction.[2]

[2]When replacing B by B^1 in a single position, state-priority and action-priority update will still give the same result, but this is not generally true when replacing B

We leave a more detailed exploration of these ideas for a future paper. The main point here was to lay the formal groundwork for starting to work with abduction in plausibility models, and to illustrate the possibility of getting the advantages of DEL in terms of expressivity, but still be able to handle computational complexity by not keeping track of *everything* that might happen at each time step. The work was directly motivated by the observed practical computational limitations of working with DEL in human-robot interaction scenarios with many unobserved actions taking place (e.g. in false-belief tasks) [9], and we plan to apply the ideas of this paper in that setting. Doing most plausible updates will of course not give computational advantages in the worst case, as surprising actions could force us to do abduction (unprune) until we end up with standard full product updates. However, the hope and expectation is that it will in many settings give a significant practical advantage. Humans clearly also don't keep track of all possible past events, but also rely on "reconstructing the past" when faced with surprising observations, and this is what the proposed approach can to some extent mimic in the rich setting of DEL.

References

[1] Mikkel Birkegaard Andersen, Thomas Bolander, and Martin Holm Jensen. Don't plan for the unexpected: Planning based on plausibility models. *Logique et Analyse*, 58(230):145–176, 2015.

[2] Andrew B Baker and Matthew L Ginsberg. Temporal projection and explanation. In *IJCAI*, pages 906–911. Citeseer, 1989.

[3] Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT7)*, volume 3 of *Texts in Logic and Games*, pages 13–60. Amsterdam University Press, 2008.

[4] John Bell. Chronological minimization and explanation. *Working papers of Common Sense*, 98, 1998.

by B^i in multiple positions. As earlier mentioned (Example 5.2), we should then apply state-priority update, $\otimes_{sp}$, to achieve chronological minimisation, and action-priority update, $\otimes_{ap}$, to achieve reverse chronological minimisation.

[5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, UK, 2001.

[6] Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21:9–34, 2011.

[7] Thomas Bolander, Lasse Dissing, and Nicolai Herrmann. DEL-based epistemic planning for human-robot collaboration: Theory and implementation. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR 2021)*, 2021.

[8] Thomas Bolander and Hermine Grosinger. Reasoning about beliefs and expectations using plausibility models for epistemic proactivity. under submission, 2025.

[9] Lasse Dissing and Thomas Bolander. Implementing Theory of Mind on a robot using Dynamic Epistemic Logic. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.

[10] David Harel, Dexter Kozen, and Jerzy Tiuryn. Dynamic logic. *ACM SIGACT News*, 32(1):66–69, 2001.

[11] Jonathan Pieper. Plausibility planning for simplified implicit coordination. Master's thesis, University of Freiburg, 2023.

[12] Jan Plaza. Logics of public announcements. In *Proceedings 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.

[13] Murray Shanahan. Explanation in the situation calculus. In *IJCAI*, pages 160–165, 1993.

[14] Yoav Shoham. *Reasoning about change: time and causation from the standpoint of artificial intelligence*. MIT Press, 1987.

[15] Johan van Benthem, Jan van Eijck, and Barteld Kooi. Logics of communication and change. *Inf. Comput.*, 204(11):1620–1662, 2006.

[16] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer Publishing Company, 2007.