What is DEL good for?

Alexandru Baltag

Oxford University

DEL is a Method, Not a Logic!

I take Dynamic Epistemic Logic () to refer to **a general type of logical approach to information change**, approach subsuming, but not being reducible to, any of the various dynamic-epistemic logics known in the literature.

In this wider sense, DEL is **a method, rather than a "logic"**.

DEL is not only "epistemic"!

In fact, this kind of formalism can be (and has been) applied, not only to knowledge change, but also to the dynamics of *other, non-epistemic forms of "information"*: belief change, factual change, preference (or payoff) change, dynamics of intentions, counterfactual dynamics, probabilistic dynamics etc.

So I call "dynamic epistemic logic" mainly because it arose within epistemic logic (but also because I am personally interested mainly in its epistemological significance).

The Four Ingredients of DEL

As such, the DEL approach can be characterized by **three main (obligatory) ingredients**, plus a **fourth (optional) one**:

(1) "**Dynamic Syntax**":
a *PDL-type of syntax*, with **dynamic modalities**

$$[\alpha]\varphi$$

associated to "actions" or "events" α, acting on top of any given logic for "static information" (knowledge, belief, preferences, intentions etc);

(2) "**Dynamic Semantics**":

a *semantics* for events based on "**model transformers**".

Given any class $\mathcal{M}_{states}$ of "static" models representing possible "information states", the informational "events" are represented as (partial) *transformations*

$$T : \mathcal{M}_{states} \to \mathcal{M}_{states}$$

on this class;

(3) “**Dynamic Proof System**”:

a system of axioms, often in the (equational) form of “**Reduction Laws**”, *describing the behavior of dynamic modalities* $[\alpha]\varphi$, and which can thus be used to **predict** future information states in terms of the current information state and the intervening event(s).

The Fourth Ingredient

(4) "**Dynamics as Merge**":
a specific way to generate model transformers, by first directly representing each action's inherent informational features in an "**event model**" (paralleling the given models for "static information"), and then defining an "**update operator**" as a partial map

$$\otimes : \mathcal{M}_{states} \times \mathcal{M}_{events} \rightarrow \mathcal{M}_{states}$$

that "merges" prior static models M_{states} with event models M_{events}, producing "posterior" (or "updated") static models $M_{states} \otimes M_{events}$.

The idea is that *the information-changing "event" is an object of the same "type" as the static information* that is being changed, and that *the new information state is obtained by merging* these two informational objects.

Information dynamics becomes a special case of information merge (aggregation).

Example: PAL

(1) **Syntax**:

Epistemic Logic (with, or without common knowledge $C_A\varphi$, distributed knowledge $D_A\varphi$ within a group $A \subseteq \mathcal{A}$ of agents) + public announcement modalities

$$[!\varphi]\psi$$

(2) **Semantics**:

$\mathcal{M}_{states}$ is the class of *"pointed" epistemic Kripke models*

$$M_{static} = (\mathbf{M}, s_*)$$

with $\mathbf{M} = (M, \{R_a\}_{a\in\mathcal{A}}, \| \bullet \|)$. Usually (but not always),

R_a are taken to be *equivalence relations*.

K_a is the Kripke modality for R_a, D_A is the Kripke modality for the intersection $\bigcap_{a \in A} R_a$ of all epistemic relations in G, and C_A is the Kripke modality for the reflexive-transitive closure $(\bigcup_{a \in A} R_a)^*$ of the union of all epistemic relations in A.

Semantics of PAL – continued

The transformation T_φ associated to the modality $[!\varphi]$ accepts as inputs only models $M_{static} = (\mathbf{M}, s_*)$ in which $s_* \models_{\mathbf{M}} \varphi$. The output is the *relativized model* $M^\varphi_{static} = (\mathbf{M}|^\varphi, s_*)$ obtained by restricting everything (the domain, the epistemic relations and the valuation) to the set

$$\|\varphi\|_{\mathbf{M}} = \{s \in M : s \models_{\mathbf{M}} \varphi\}$$

of all states satisfying φ (in $\mathbf{M}$).

PAL's Proof System

(3) **Reduction Laws**:

$$[!\varphi]p \iff \varphi \Rightarrow p$$

$$[!\varphi]\neg\psi \iff \varphi \Rightarrow \neg[!\varphi]\psi$$

$$[!\varphi]K_a\psi \iff \varphi \Rightarrow K_a[!\varphi]\psi$$

$$[!\varphi]D_A\psi \iff \varphi \Rightarrow D_A[!\varphi]\psi$$

What about **common knowledge**?

Well, it turns out *there is no Reduction Law* for $[!\varphi]C\psi$ in terms of **classical** static epistemic logic EL only!

Ways Out

TWO SOLUTIONS (both very fertile) have been proposed:

(a) (J. van Benthem) **Extend the language of classical EL** with *some appropriate static modality* ("conditional common knowledge" $C^{\varphi}\psi$), *that "pre-encodes" the dynamics* of common knowledge.

(b) (BMS) **Extend the proof system** of classical EL with new axioms or rules, *axiomatizing directly the dynamic logic* PAL(C).

Semantics for Conditional Common Knowledge

In a model $(\mathbf{M}, s_*)$, put

$$R_a^\varphi = R_a \cap (M \times \|\varphi\|_{\mathbf{M}}) = \{(s,t) \in R_a : t \models_{\mathbf{M}} \varphi\}.$$

Then C_A^φ is the Kripke modality for the reflexive-transitive closure $(\bigcup_{a \in A} R_a)^*$ of the union of all epistemic relations R_a^φ with $a \in A$.

Essentially, this makes $C_A^\varphi \psi$ equivalent to the infinite conjunction

$$\psi \wedge \bigwedge_{a \in A} K_a(\varphi \Rightarrow \psi) \wedge \bigwedge_{a \in A} K_a(\varphi \Rightarrow \bigwedge_{a \in A} K_a(\varphi \Rightarrow \psi)) \wedge \ldots$$

Reduction Laws for (Conditional) Common Knowle

With this, the reduction law is

$$[!\varphi]C_A\psi \iff \varphi \Rightarrow C_A^{\varphi}[!\varphi]\psi$$

and, more generally,

$$[!\varphi]C_A^{\theta}\psi \iff \varphi \Rightarrow C_A^{<!\varphi>\theta}[!\varphi]\psi$$

Another Example: "Tell Us All You Know"

Suppose we introduce a dynamic modality $[!a]\psi$, corresponding to the action by which agent a **publicly announces "all (s)he knows"**.

We interpret this in a *language-independent manner*: a announces which states (s)he considers possible (or equivalently, which states she knows to be impossible).

Semantics of $!a$

On a pointed model $M_{static} = (\mathbf{M}, \mathbf{s}_*)$, this acts as the public announcement $!s(a)$ of the set

$$s(a) := \{t \in M : s_* R_a t\},$$

representing agent a's **current information cell** (in the partition induced by a's equivalence relation).

So the semantics of $!a$ is given by **relativizing** (i.e. restricting all the components of) M_{static} to the set $s(a)$.

Proof System for $!a$

Reduction Laws:

$$[!a]p \iff p$$

$$[!a]\neg\psi \iff \neg[!a]\psi$$

$$[!a]K_b\psi \iff D_{\{a,b\}}[!a]\psi$$

$$[!a]D_A\psi \iff D_{A\cup\{a\}}[!a]\psi$$

What about **common knowledge**?

Again, we have to introduce a new modality, formalizing yet another ("static") epistemic attitude.

Common Knowledge Conditional on Others' Knowle

For $A, B \subseteq \mathcal{A}$, we read $C_A^B \psi$ as saying that: *group* A *has common knowledge of* ψ *conditional on the knowledge of (all agents of) group* B.

Formally, C_A^B is defined as the Kripke modality for

$$\left(\bigcup_{a \in A} (R_a \cap \bigcap_{b \in B} R_b) \right)^*$$

which is the same as

$$\left((\bigcup_{a \in A} R_a) \cap \bigcap_{b \in B} R_b \right)^*$$

The Static Logic of C_A^B

The static logic C_A^B is completely axiomatized by: *Modus Ponens* and *Necessitation* (for C_A^B), together with *the standard S5 axioms* for C_A^B and the *Monotonicity Axiom*:

$$C_A^B \Rightarrow C_{A'}^{B'} \ , \quad \text{for } A \supseteq A', B \subseteq B'.$$

All the standard epistemic operators are definable:

$$K_a\psi = C_{\{a\}}^{\{a\}}\psi = C_{\{a\}}^{\emptyset}$$

$$C_A\psi = C_A^{\emptyset}\psi$$

$$D_A\psi = C_A^A\psi = C_{\{a\}}^A\psi = C_{\{a\}}^{A\setminus\{a\}}, \text{ for any } a \in A.$$

Reduction Laws

$$[!a]C_A\psi \iff C_A^{\{a\}}[!a]\psi$$

and, more generally,

$$[!a]C_A^B\psi \iff C_A^{B\cup\{a\}}[!a]\psi$$

Other Examples: Updates and Upgrades on Plausib

Another example of application of this strategy is the *dynamics of belief*, induced by *hard updates* $!\varphi$ and *soft upgrades* $\Uparrow \varphi$ and $\uparrow \varphi$ on *belief-revision models*.

Such models can be given in terms of "Grove spheres", or alternatively as "**preference (or plausibility) models**": Kripke models in which the relation is assumed to be a *total preorder*.

To have reduction laws for beliefs, one needs to extend again the language, by introducing *conditional beliefs* $B^{\varphi}\psi$.

Dynamics as Merge: Event Models

A model for static information (epistemic Kripke model, or plausibility model, or preference model, or probabilistic model, or Lewis model for counterfactuals etc) can be alternatively interpreted as an "event model":

its possible worlds represent now possible *informational events*, the "valuation" gives us now the *precondition* of a given event and the *factual changes* induced by the event, and the epistemic relations (or plausibility/preference relations, or comparative similarity relations) represent the *knowledge/beliefs/preferences (or counterfactual judgments) about the current event.*

Preference Merge in Social Choice Theory

In Social Choice Theory, the main issue is how to *merge* the agent's individual preferences in a reasonable way. In the case of *two agents*, a merge operation is a function $\bigodot$, taking preference relations R_a, R_b into a *"group preference" relation* $R_a \bigodot R_b$ (on the same state space).

As usually considered, the problem is to find a *"natural" merge operation* (subject to various *fairness conditions*), for merging the agents' preference relations. Depending on the stringency of the required conditions, one can obtain either an Impossibility Theorem or a classification of the possible types of merge operations.

Merge by Intersection

The so-called *parallel merge* (or "*merge by intersection*") simply takes the merged relation to be

$$\bigcap_a R_a.$$

In the case of two agents, it takes:

$$R_a \bigodot R_B := R_a \cap R_b$$

This could be thought of as a *"democratic" form of preference merge.*

Distributed Knowledge is 'Static' Parallel Merge

This form of merge is particularly suited for "hard information" (irrevocable knowledge) K: since this is an absolutely certain, fully reliable, unrevisable and fully introspective form of knowledge, there is no danger of inconsistency. The agents can pool their information in a *completely symmetric manner*, *accepting the other's bits without reservations*.

The concept of "distributed knowledge" DK in epistemic logic corresponds to the parallel merge of the agents' hard information:

$$DK_{a,b}P = [R_a \cap R_b]P.$$

Information Dynamics as Preference Merge

We can turn the tables around by thinking of *information dynamics as a "kind" of information (or preference) merge*.

The idea is that we can think of agent a's *new observations* (her plausibility order on epistemic/doxastic actions) as being *the beliefs/informtion of another agent* $\tilde{a}$. The way the new observations/actions change a's beliefs can then understood as a *merging* of a's beliefs with $\tilde{a}$'s beliefs.

Product Update is ‘Dynamic’ Parallel Merge

When we think in this way, we can say that the so-called Product Update of Baltag, Moss and Solecki corresponds to *parallel merge (merge be intersection)*.

This is not surprising: classical DEL deals only with “hard” information, and we’ve seen that the most natural merge operation for “hard” information is the parallel merge.

Reduction Law: The Action-Knowledge Axiom

This law embodies the essence of Product Update Rule, governing the most general **dynamics of "hard information"** K:

$$[\alpha]K_a\mathbf{P} \quad \leftrightarrow \quad pre_\alpha \to \bigwedge_{\beta R_a \alpha} K_a[\beta]\mathbf{P}$$

A proposition $\mathbf{P}$ *will be known after an epistemic event* α *iff, whenever the event can take place, it is known (before the event) that* $\mathbf{P}$ *will be true after all events that are epistemically indistinguishable from* α*.*

Lexicographic Merge

In lexicographic merge, a "priority order" is given on agents, to model the group's hierarchy. For two agents a, b, we denote by $R_{a/b}$ the lexicographic merge in which agent a has priority over b.

The strict preference of a is adopted by the group; if a is indifferent, then b's preference (or lack of preference) is adopted; finally, a-incomparability gives group incomparability. Formally:

$$R_{a/b} := R_a^{>} \cup (R_a^{\simeq} \cap R_b) = R_a^{>} \cup (R_a \cap R_b) = R_a \cap (R_a^{>} \cup R_b).$$

Lexicographic merge of soft information

This form of merge is *particularly suited for "soft information"*, given by either indefeasible knowledge $\Box$ or belief B, *in the absence of any hard information*: since soft information is not fully reliable (because of lack of negative introspection for $\Box$, and of potential falsehood for B), some "screening" must be applied (and so some hierarchy must be enforced) to ensure consistency of merge.

$$s \models \Box_{a/b} P$$

iff

$\exists P_a, P_b$ s. t. $s \models \Box_a P_a \wedge \Box_b P_b \wedge \Box_a^{weak} P_b$ and $P_a \cap P_b \subseteq P$.

In other words, *all a's "indefeasible knowledge" is unconditionally accepted by the group, while b's indefeasible knowledge is "screened" by a using her "weakly indefeasible knowledge".*

Relative Priority Merge

Note that, in lexicographic merge, the first agent's priority is "absolute" in the sense that her strong preferences are adopted by the group even when they are incomparable according to the second agent. But *in the presence of hard information*, the lexicographic merge of soft information must be modified (by first pooling together all the hard information and then using it to restrict the lexicographic merge). This leads us to a "more democratic" form of merge: the *(relative) priority merge* $R_{a \otimes b}$, given by $R_{a \otimes b} := (R_a \cap R_b^{\sim}) \cup (R_a^{\simeq} \cap R_b)$
$= (R_a^{>} \cap R_b^{-1}) \cup (R_a \cap R_b) = R_a \cap R_b^{\sim} \cap (R_a^{>} \cup R_b)$.

Essentially, this means that both agents have a "veto" with respect to group incomparability: the group can only compare options that *both* agents can compare; and whenever the group can compare two options, everything goes on as in the lexicographic merge: agent a's strong preferences are adopted, while b's preferences are adopted only when a is indifferent.

Relative Priority Merge can be thought of as *a combination of Merge by Intersection and Lexicographic Merge*: the "hard" information is merged by intersection; then the "soft" information is lexicographically merged; but with the proviso that it still has to be consistent with the group's hard information.

Priority Merge of Soft Information

The corresponding notion of "indefeasible knowledge" of the group is obtained as in the lexicographic merge, *except that both agents' "strong knowledge" is unconditionally accepted.* Formally:

$$s \models \Box_{a \otimes b} P$$

iff

$$\exists P_a, P_b, \varphi'_b \text{ s. t. } s \models \Box_a P_a \wedge K_b P_b \wedge \Box_b P'_b \wedge \Box_a^{weak} P'_b$$

$$\text{and } P_a \cap P_b \cap P'_b \subseteq P.$$

In other words, relative-priority group "knowledge" is obtained by pooling together the following: agent a's "indefeasible knowledge"; agent b's "irrevocable knowledge"; and the result of screening agent b's "indefeasible knowledge" using agent a's "weakly indefeasible knowledge".

Action Priority Update is Priority Merge

The natural merge operation for "soft" information is the Priority Merge. In the context of Belief Revision Theory, the AGM paradigm asks us to *give priority to the new information.*

So the natural product update operation for soft information will be given by Priority Merge, where the "new" agent $\tilde{a}$ has priority over the "old" agent a. This is exactly what "The Action-Priority" update" introduced in my joint work with Sonja Smets on belief revision.

Another Example: Counterfactual Dynamics

If we interpret the preference relations as expressing a Lewis-type “comparative similarity” relation

$$s \leq_w s'$$

between worlds, saying that world s is at least as similar to world w as world s', then we obtain the *Lewis semantics for counterfactuals*.

Example

Alice tossed a (fair) coin once, and it fell Heads up. But it could have fallen Tails up.

Below we draw the comparative similarity relation $\leq_{\mathsf{T}_1}$ for the actual world T_1:

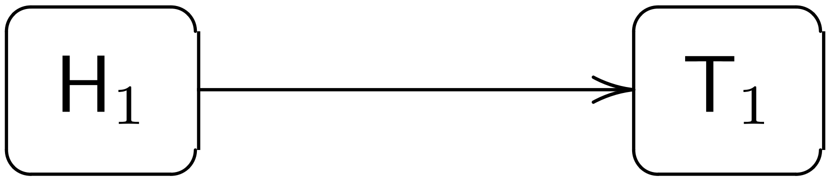

Counterfactual Conditionals

The semantics of counterfactual conditionals $\varphi \mathrel{\Box\!\!\to} \psi$ is given by:

$\varphi \mathrel{\Box\!\!\to} \psi$ **holds in a world** w **if** ψ **is true in all the worlds satisfying** φ **that are "most similar" to** w**.**

Counterfactual Event Models

A new event is now happening: Alice tosses the coin a second time (2). Let's say it'll fall Heads up this time (H_2), though this is not yet determined (or at least not yet known to Alice).

This event model is

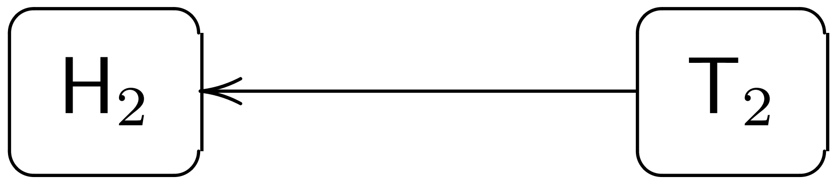

where we now represented the comparative similarity relation $\leq_{\mathsf{H}_2}$ for the actual event that will be happening.

Backtracking Update is Lexicographic Merge

The most commonly used type of "historical" or "dynamic" counterfactual is given by **backtracking**.

"Backtracking Update"

The natural update product notion for counterfactual event models, and the *only one that matches/generalizes backtracking*, is exactly the dual of the one for belief revision: **a "lexicographic merge that gives preference to the past**.

Counterfactual Semantics

The backtracking update says that: **two worlds/histories are more similar if they differ only in their last (current) event than if their differences run deeper in the past.**

$$(s, \sigma) \leq_{(w,\omega)} (s', \sigma')$$

iff:

either $s <_w s'$

or $s \simeq_w s'$, $\sigma \leq_\omega \sigma'$.

Example Continued

The backtracking update gives us the comparability relation for the actual world (T, H) after the event:

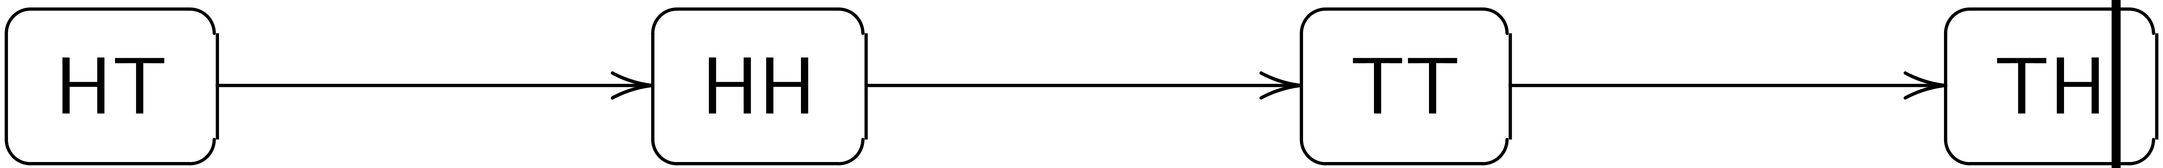

Dynamic Reduction Laws for Counterfactuals

Theorem: (Baltag and Smets 2010) There exists a **complete proof system** for DEL with hard and soft information and with counterfactual dynamics, that includes the following Reduction Law:

$$[\alpha](\varphi \,\Box\!\!\rightarrow \psi)$$

$$\Leftrightarrow$$

$$pre_\alpha \Rightarrow \left(\left(\bigvee_\beta \langle \beta \rangle \varphi \right) \Box\!\!\rightarrow \bigvee_\beta \left(\langle \beta \rangle \varphi \wedge \bigwedge_{\gamma \leq_\alpha \beta} [\gamma](\varphi \Rightarrow \psi) \right) \right)$$

What's DEL good for?

- DEL can help to **uncover various new (static) informational attitudes** (epistemic/doxastic attitudes, preference operators etc), that *pre-encode* the dynamics of other (more familiar) such attitudes.
- DEL can **classify** (static) informational attitudes in terms of **their specific dynamics** and their **sources** (the informational events that produced them), via their *corresponding Reduction Laws*.

- DEL can **characterize** (static) informational attitudes, as (a) being **preserved** by specific types of informational events, or (b) as defining the **fixed points** of such events.

- DEL can **explain** in some sense various *types of informational dynamics* in terms of various *types of preference merge*, and dually to "realize" dynamically various such types of preference merge via communication/persuation scenarios: potential connections with Social Choice Theory and Decision Theory.

- the fine distinctions made by DEL between "static" and "dynamic" revision, and thus between the corresponding static and dynamic epistemic attitudes, have applications to the understanding of *rationality and solution concepts in Game Theory*, and to solving various *epistemic puzzles* (Moore, Fitch, Voorbrak)

- the DEL investigation of fixed points of iterated informational dynamics has applications in Learning Theory and in solving epistemic paradoxes (the Surprise Exam).
- the DEL investigation of various forms of private communication, cheating, lying, interception has potential application in CS (security protocols).