INFERENCE IN ACTION

Johan van Benthem

ABSTRACT. Substructural logics arise whenever classical logic is put to new uses, and logicians from Serbia have been in the fore-front here. In this paper, we join the substructural tradition with another recent trend, viz. dynamic logic of information update. We show how these two approaches fit together, in particular, through a number of representation theorems concerning structural rules. The proper background for these results turn out to be modal and dynamic logics of cross-model relations. We connect this finding with recent accounts of generalized inference, interpolation, and preservation results.

1. Inference, structural rules, and information-producing actions

Inference entangled with other information sources. In the 1980s, several interesting nonstandard notions of consequence $\boldsymbol{P} \Rightarrow C$ emerged, claiming to reflect features of our common sense reasoning. Circumscription in AI looked at conclusions C true only in *minimal* models of the premises $\boldsymbol{P}$, with minimality measured by some comparison order for model size or predicate interpretation. General *non-monotonic logics* followed up on this idea, high-lighting failures of classical principles much as those found earlier in conditional logic. Structural rules, i.e., abstract properties of an inference relation $\Rightarrow$, seemed a natural focus for defining 'styles of reasoning', in terms of their basic mechanics. This idea was reinforced when it turned out that very different notions of consequence, such as the *resource-conscious inferences* found in categorial grammar (van Benthem 1991) have illuminating sets of structural rules setting them apart from others. Likewise, van Benthem 1996 showed how deviant structural rules emerge in a natural fashion when analyzing so-called *dynamic semantics*, emphasizing how inference and information change are intertwined in understanding and using language. In this paper I present some further thoughts on the notion of inference emerging from all this, and its entanglement with information update and general action.

Before going to abstract structural rules and bare mechanics, however, consider an example. The Amsterdam Science Museum *NEMO* (http://www.nemo-amsterdam.nl/) organizes Kids' Lectures ON Science for 8-year olds. While preparing for this event, I wondered how to talk to such an audience? I came up with an example that goes back essentially to Antiquity:

The Restaurant. "In a restaurant, your Father has ordered Fish, your Mother ordered Vegetarian, and you have Meat. Out of the kitchen comes some new person with the three plates. What will happen?" The children got excited, many little hands were raised, and one said: "He asks who has the Meat". "Sure enough", I said: "He asks, hears the answer, and puts the plate. What happens next?" Children said "He asks who has the Fish!" Then I asked once more what happens next. And now one could see the Light of Reason start shining in those little eyes. One girl shouted: "He does not ask!" Now, *that* is logic...

In my view, the Restaurant is about the simplest realistic logical scenario (van Benthem 2007, van Benthem 2008). Several basic informational actions take place intertwined: questions, answers, and inferences, and the setting crucially involves more than one agent. Also, actions can be analyzed for their informational content after they have taken place, but they can also be planned beforehand. Thus there is no natural border-line here between inference and other actions that produce information. I would say that 'logical analysis' of even this basic scenario involves all of them – and a logical system should account explicitly for that interplay. Indeed, the same entanglement is found in Indian Logic, a tradition parallel to our western one, where various sources of obtaining information were treated on a par: including making an *observation*, drawing a *conclusion*, or *asking* someone!

Dynamic inference over abstract transition models. Having said all this, let us first go to a very general abstract way of bringing actions into logic. We can view new propositions A dynamically as *partial functions* T_A taking input states meeting the preconditions of update with A to output states:

$$\bigcirc \xrightarrow{\quad T_A \quad} \bigcirc$$

More generally, *transition models* $\boldsymbol{M} = (S, \{T_A \mid A \in \mathbf{Prop}\})$ consist of information states S with a family of transition relations T_A over these, one for each proposition A in some given set $\mathbf{Prop}$. These suggest the following notion of inference. A sequence of propositions $P_1, \ldots, P_k$ *dynamically implies* conclusion C in transition model $\boldsymbol{M}$, if any sequence of premise updates starting from any state in $\boldsymbol{M}$ ends in a fixed point for the conclusion:

$$\text{whenever } s_1 T_{p_1} s_2 \ldots T_{p_k} s_{k+1}, \text{ then } s_{k+1} C s_{k+1}$$

We say *the sequent $P_1, \ldots, P_k \Rightarrow C$ is *true* in the model* – $\boldsymbol{M} \vDash P_1, \ldots, P_k \Rightarrow C$. Here $\boldsymbol{P}, Q, R$ stand for finite sequences of propositions, and A, B, C for single ones. Dynamic inferential sequents lack the structural rules of classical consequence (van Benthem 1996, Chapter 7). Simple counter-examples refute Monotonicity, Contraction, Permutation, or Reflexivity – and their general idea amounts to this: any sequential recipe for some desired effect may be disturbed by inserting instructions, deleting repeats of an instruction, permuting instructions, etc. Even the Cut Rule fails in its general form:

$$\text{if } \boldsymbol{P} \Rightarrow A \text{ and } \boldsymbol{R}, A, \boldsymbol{Q} \Rightarrow C, \text{ then } \boldsymbol{R}, \boldsymbol{P}, \boldsymbol{Q} \Rightarrow C$$

But dynamic inference is not totally unprincipled – and some 'substitute rules' turn out to hold. Partial update functions validate the following structural rules:

$$\text{if } \boldsymbol{P} \Rightarrow C, \text{ then } A, \boldsymbol{P} \Rightarrow C, \qquad \text{Left-Monotonicity}$$
$$\text{if } \boldsymbol{P} \Rightarrow A \text{ and } \boldsymbol{P}, A, \boldsymbol{Q} \Rightarrow C, \text{ then } \boldsymbol{P}, \boldsymbol{Q} \Rightarrow C \qquad \text{Left-Cut}$$
$$\text{if } \boldsymbol{P} \Rightarrow A \text{ and } \boldsymbol{P}, \boldsymbol{Q} \Rightarrow C, \text{ then } \boldsymbol{P}, A, \boldsymbol{Q} \Rightarrow C \quad \text{Cautious Monotonicity}$$

Indeed, these structural rules are completely characteristic for dynamic inference with partial update functions. Take any set of propositions **Prop** as abstract objects – and any binary relation $\Rightarrow$ between finite sequences of propositions and propositions. We repeat a result from van Benthem 1996, as it shows the flavour of the situation rather nicely:

THEOREM 1. *The following are equivalent for any structure* $(\textbf{Prop}, \Rightarrow)$:
   (a) $\Rightarrow$ *satisfies Left-Monotonicity, Left-Cut, Cautious Monotonicity, viewed as abstract conditions on relations of type 'sequence-to-object',*
   (b) *there is a transition model* $(S, \{T_A \mid A \in \textbf{Prop}\})$ *with partial maps* T_A *whose relation of dynamic inference coincides with the given* $\Rightarrow$.

PROOF. The direction from (b) to (a) is easy to check, as we suggested above. From (a) to (b), any abstract structure $(\textbf{Prop}, \Rightarrow)$ induces a transition model $\boldsymbol{M}$ with states are finite sequences X, Y of propositions. Each proposition A then defines a partial function over these states:

$$T_A = \{(X, X) \mid X \Rightarrow A\} \cup \{(X, \langle X, A \rangle) \mid \text{not } X \Rightarrow A\}$$

We must check that the following equivalence holds:

$$\boldsymbol{M} \vDash P_1, \ldots, P_k \Rightarrow C \text{ iff } P_1, \ldots, P_k \Rightarrow C \text{ is true in } (\textbf{Prop}, \Rightarrow)$$

'If'. Suppose that $s_1 T_{p_1} s_2 \ldots T_{p_k} s_k$. By the definition of the functions T_A, each step in this sequence either adds a proposition at the end, or it just 'pauses'. Here is a typical illustration:

$$X \, T_{p_1} \, \langle X, P_1 \rangle \qquad (\text{not } X \Rightarrow P_1)$$
$$\langle X, P_1 \rangle \, T_{p_2} \, \langle X, P_1 \rangle \qquad (\langle X, P_1 \rangle \Rightarrow P_2)$$
$$\langle X, P_1 \rangle \, T_{p_3} \, \langle X, P_1, P_3 \rangle \quad (\text{not } \langle X, P_1 \rangle \Rightarrow P_3)$$

We show that the end state $\langle X, P_1, P_3 \rangle$ is a fixed point for T_C: i.e., $\langle X, P_1, P_3 \rangle \Rightarrow C$. First we have $\langle P_1, P_2, P_3 \rangle \Rightarrow C$, and so by Left-Monotonicity $\langle X, P_1, P_2, P_3 \rangle \Rightarrow C$. Following the transition steps, we suppress one proposition thanks to $\langle X, P_1 \rangle \Rightarrow P_2$, using Left-Cut to get $\langle X, P_1, P_3 \rangle \Rightarrow C$. This argument is general. 'Pauses' involve valid sequents used to cut out items in the sequence $P_1, \ldots, P_k$ at the right places.

'Only if'. This involves the remaining structural rule. Again, here is a simple example. Let $\langle P_1, P_2, P_3 \rangle$ dynamically imply C in our transition structure $\boldsymbol{M}$. Start with the empty sequence –. We choose three particular transitions for the premises. If $- \Rightarrow P_1$ in **Prop**, the first transition is $-, -$; otherwise, take an extended sequence $\langle P_1 \rangle$; etc. Suppose this yields the following sequence of transformations:

$$-, \langle P_1 \rangle \quad \langle P_1 \rangle, \langle P_1 \rangle \text{ (where } P_1 \Rightarrow! P_2) \quad \langle P_1 \rangle, \langle P_1, P_3 \rangle$$

By assumption, the final state is a fixed point for T_C: that is, $P_1, P_3 \Rightarrow C$ is true in **Prop**. But then, since $P_1 \Rightarrow P_2$, and using Cautious Monotonicity: $P_1, P_2, P_3 \Rightarrow C$ is true in **Prop**. Again the general trick is clear. We can insert propositions in our sequence just where these are required.   □

A simple extension yields a completeness theorem for sequents on transition models M (cf. van Benthem 2003B). A sequent σ is a valid consequence of a *set of sequents* Σ iff σ is derivable from sequents in Σ using the three mentioned structural rules. Other notions of dynamic inference place other requirements on the action associated with the conclusion. Their structural properties may be determined in a similar manner to the above representation.

So much for basic connections between logical propositions and abstract actions on some state space. Let us now develop this joint perspective in more detail.

2. Inference along a relation and planning actions: a modal view

2.1. Inference and links across different models. Abstract consequence relations often involve just a relation between two propositions which are supposed to be true in some fixed situation under consideration. But as we have seen just now, inference may also take place in settings where the relevant situation changes, or at least, where we shift between situations where propositions can be true. There can be many reasons for this. One is information update, but there are many other channels for information flow. In a traditional Indian inference schema, one is at the foot of a mountain where propositions can be decided by direct observation, but one wants to know what is happening at the top of the mountain, which is not open to detailed inspection. That is the case where inferences come to the rescue, such as that from observing smoke at the top down here to the existence of a fire up there. Essentially the same example is pivotal in Barwise and Seligman 1995 with situation semantics in terms of information flow in networks. Van Benthem 1998 discusses abstract 'information links' between models and the need for a basic logic of these. This theme was taken further in Barwise and van Benthem 1999, who introduce the notion of entailment along some inter-model relation:

DEFINITION 2. *P entails C along relation R* if, whenever $M \vDash P$ and $M R N$, then $N \vDash C$.

We will discuss structural properties of this generalized form of inference below.

REMARK. Note that we have 'relocated the dynamics' here, as compared with Section 1. There, we made the propositions themselves into actions transforming states of some given transition model. Here, however, we retain classical 'static' propositions P, C denoting properties of states, whereas the dynamics shows rather in the state-shifting transition relations R.

But for now, let us make one further move. If these relations R are so important, then why not put them explicitly into our language? This makes all the more sense, since we need not assume just one relation of interest when jumping across

situations. Now, there already exists an obvious notation for the preceding notion of state shifting, viz. the *modal formula*

$$\boldsymbol{P} \to [R]C$$

Of course, to express sequent validity as before, this formula would have to be true in some universe of 'relevant models', whose nature is yet to be stipulated.

This modal language lies one step up in expressive power from the standard austere sequent format used in formulating properties of inference relations, but one can still view it as a sort of perspicuous notation for very basic properties, and their interplay with Boolean and action structure. In the remainder of this section, we take a closer look at this modal format, under various interpretations, and with further kinds of statement.

Action-tagged sequents and calculus of plans. The poly-modal format also serves further purposes, as a 'calculus of plans'. Van Benthem 1998 presents natural operations on plans with Horn-type rules, and analyzes connections with resolution in first-order logic. Just by way of illustration, let us say we want to infer, not what is true in the current situation, but what can be *made true* by performing suitable actions. So, given some of our transfer statements $A \to [R]B$, how to derive new ones with complex transitions? Here are a few examples:

$$\begin{aligned}
&\text{from } A \to [R]B, \ B \to [S]C \ \text{ infer } \ A \to [R;S]C && \text{composition} \\
&\text{from } A \to [R]B && \text{infer } \ \neg B \to [R^{\cup}]\neg A && \text{converse} \\
&\text{from } A \to [R]B, \ A \to [S]B \ \text{ infer } \ A \to [R \cup S]B && \text{union}
\end{aligned}$$

These laws tag ordinary propositional implications with actions. 'Labeled sequents'

$$\boldsymbol{P} \Rightarrow_R C$$

with possibly complex relations R would now explicitly represent actions shifting the relevant model in the passage from premises to conclusion. Richer logics beyond the basic polymodal one use sequential operations from dynamic logic here in building the R, such as composition, choice, or finite Kleene iteration. Indeed, logical inference even suggests the use of *parallel* composition of actions to obtain conjunctions of effects, as in the next rule:

$$\text{from } A \to [R]B, \ C \to [S]D \ \text{ infer } \ (A, C) \to [R \times S](B, D) \quad \text{product}$$

Validity is easily checked in its first-order transcription:

$$\forall xyzu(((Ax \ \& \ Cy) \ \& \ (Rxz \ \& \ Syu)) \to (Bz \ \& \ Du)).$$

The resulting calculus describes valid reasoning with labelled sequents of this sort. It uses monotonicity inferences in antecedents and consequents. For the sake of concreteness, here is an illustration:

2.2. Excursion: dynamic inference as plan calculus. For convenience, rewrite sequents $A \rightarrow [R]B$ to a format with a past temporal operator

$$P_R A \rightarrow B, \text{ or written with a converse modality: } \langle R^{\cup} \rangle A \rightarrow B.$$

These can be viewed as implications $\Pi \rightarrow B$ where the plan Π describes a preceding successful execution of some actions from given resources. A calculus with action-tagged sequents can synthetize plans. Consider a resource proposition A and a goal proposition G. Our available premises encode available subroutines:

$$P_S B \wedge C \rightarrow G, \quad P_T B \rightarrow C, \quad P_U A \rightarrow B$$

We now 'derive' G from A by the following heuristics:

$$1 \ G \text{ from } B,C \quad 2 \ B \text{ from } A \quad 3 \ C \text{ from } B \quad 4 \ B \text{ from } A$$

Composing the associated trees required for this works out to

$$1 \ P_S \wedge C \quad 2 \ P_S P_U A \wedge C \quad 3 \ P_S P_U A \wedge P_T B \quad 4 \ P_S P_U A \wedge P_T P_U A$$

Our discussion so far may have shown the interest of taking a polymodal perspective on abstract inference relations. On the basis of this evidence, let us now state our general recommendation in this section:

The minimal modal logic is the basic structural logic for 'inference in action'!

Further uses of polymodal logic as abstract sequent calculus. To add yet more evidence, a full poly-modal language can express many facts beyond the above tagged sequents for entailment along a relation. Thus, existential modalities can state 'enabling principles' from inferential and computational practice:

$$A \rightarrow \langle R \rangle B: \quad A \text{ makes it possible to execute } R \text{ so that } B \text{ is achieved.}$$

As another example, we show how one can use a *loop modality* to analyze earlier substructural rules for dynamic inference in a standard modal setting (van Benthem 1996, 2003B). First, we add a modality' (a) defining the fixed-points of Section 1:

$$\boldsymbol{M}, s \vDash (a)\phi \quad \text{iff} \quad sR_a s \ \& \ \boldsymbol{M}, s \vDash \phi$$

The loop language is decidable, and it has a complete axiomatization with axioms

$$(a)\phi \ \Leftrightarrow \ (a)T \ \& \ \phi, \quad (a)T \rightarrow ([a]\phi \ \Leftrightarrow \ \phi)$$

This language reads our earlier dynamic sequents $P_1, \ldots, P_k \Rightarrow C$ as modal formulas – with letters inside boxes taken as action labels:

$$[P_1] \ldots [P_k](C)T$$

FACT 3. *The structural rules of dynamic inference listed in Theorem* 1 *are all valid principles of the modal loop language.*

PROOF. One reads these modal consequences as running from premises true in a whole transition model to their conclusions. E.g., Left Cut went

from $[\boldsymbol{P}](A)T$ and $[\boldsymbol{P}][A][\boldsymbol{Q}](C)T$ to $[\boldsymbol{P}][\boldsymbol{Q}](C)T$.

This follows from the loop law $((A)T \ \& \ [A]\phi) \rightarrow \phi$. Cautious Monotonicity went

from $[\boldsymbol{P}](A)T$ and $[\boldsymbol{P}][\boldsymbol{Q}](C)T$ to $[\boldsymbol{P}][A][\boldsymbol{Q}](C)T$,

and this is a consequence of $((A)T \,\&\, \phi) \to [A]\phi$. □

The loop language can also express complex existential properties of consequence relations beyond the mere structural rules that we started with. All this reinforces our conclusion that a poly-modal logic seems a natural stage for a richer abstract theory of dynamic inference.

3. Inference and information update

Dynamic-epistemic logic. The Restaurant story in Section 1 supports more concrete scenarios than abstract state transitions, with inference intertwined with information from concrete events, viz. *public announcements* of true propositions. These represent incoming 'hard information' of a public nature. This is the realm of modern *dynamic-epistemic logic* (Baltag, Moss and Solecki 1998, van Benthem 2006, van Ditmarsch, van der Hoek and Kooi 2007). To make our point here, just assume some standard epistemic language with operators $K_i\phi$ for knowledge: 'agent i knows that ϕ'. These modal operators are interpreted in semantic models $\boldsymbol{M} = (W, \sim_i, \leqslant_i, V)$, where the $\sim_i$ are epistemic accessibility relations giving an agent's current range of uncertainty. Then knowledge at a world w means truth at all worlds accessible from w via $\sim_i$. Complete epistemic logics are well-known, but we will formulate some less-known dynamic variants.

The simplest event producing information is a public announcement $!P$ of some true proposition P (i.e., true at the actual world s in $\boldsymbol{M}$). E.g., announcing a fact q will make you know that q – though there are more subtle phenomena with complex P. The widespread intuitive idea of new information as elimination of current possibilities arises here as an action of *model change*. The event $!P$ takes the current model $(\boldsymbol{M}, s)$ to a new structure $(\boldsymbol{M}|P, s)$, viz. the model $\boldsymbol{M}$ restricted to its sub-model consisting of just the P-worlds. To reason about this informational process, we introduce a matching dynamic operator:

$$\boldsymbol{M}, s \vDash [!P]\phi \quad \text{iff} \quad \boldsymbol{M}|P, s \vDash \phi.$$

The principles which analyze the effects of public announcements on what agents know yield a logical system PAL which is axiomatized completely by the usual laws of epistemic logic plus the following *reduction axioms*:

$$[!P]q \,\Leftrightarrow\, P \to q \qquad \text{for atomic facts } q$$
$$[!P]\neg\phi \,\Leftrightarrow\, P \to \neg[!P]\phi$$
$$[!P](\phi \wedge \psi) \,\Leftrightarrow\, [!P]\phi \wedge [!P]\psi$$
$$[!P]K_i\phi \,\Leftrightarrow\, P \to K_i(P \to [!P]\phi)$$

The last axiom here is crucial, in that it reduces knowledge after an announcement to *conditional knowledge* which agents had before the announcement was made. This is called 'pre-encoding'. In this dynamic perspective, classical consequence from premises $\boldsymbol{P}$ to a conclusion C works as follows. Updating the current model

with successive announcements $!P_1, \ldots, !P_n$ leads to a new model where the conclusion C is known to all agents, or even more strongly, a model where C has become common knowledge among them. We will make this precise in a moment.

Dynamic epistemic logic, in this and more sophisticated update scenarios, provides an appropriate setting for analyzing inferences that agents make together with information which they receive from communication, observation, or other sources. This framework is more concrete than the general transition-based paradigm of Section 2. Still, as we shall see, its general properties lie close to the structural rules that we gave before in Theorem 1.

Structural rules revisited: dynamic inference in communication. Dynamic epistemic logic supports our earlier dynamic inference. Dynamic propositions are announcements $!A$ of epistemic formulas A. Dynamic validity of a sequent $P_1, \ldots, P_k \Rightarrow \phi$ in our earlier sense now says that,

> Starting with any epistemic model whatsoever, successive announcements of the premises result in a model where announcement of ϕ effects no further change: i.e., ϕ was already true everywhere even before it was announced.

This amounts to validity of the following dynamic-epistemic formula, which says that the conclusion becomes *common knowledge*:

$$[!P_1] \ldots [!P_k]C_G\phi \qquad (\#)$$

We can read this validity as referring to the '*Supermodel*' of all epistemic models related by arbitrary announcement steps. But when modeling more realistic scenarios of conversation or enquiry, we can also relativize the preceding notions to smaller *restricted families* $\mathfrak{M}$ of epistemic models and admissible announcements. Such families were studied recently in van Benthem, Gerbrandy and Pacuit 2007.

It is easy to see that the classical structural rules all fail for this new notion of dynamic validity under premise announcements. A result from van Benthem 2003B makes the connection more precise – but we first need to define a suitable notion of validity, which we will state in terms of abstract propositions as before:

DEFINITION 4. Consider a meta-sequent $\Sigma \rightarrow \sigma$ going from a set of sequents Σ to one sequent σ. We call such a meta-sequent *update-valid* if all its substitution instances with actual epistemic formulas, reading sequents as dynamic-epistemic formulas as before, leads to a valid implication between DEL-formulas of type $(\#)$.

For the sequents obtained in this way, validity in just the above-defined Supermodel, or in arbitrary families of models $\mathfrak{M}$ as above, makes no difference.

Next, we show that the special DEL setting adds nothing beyond our earlier abstract analysis of dynamic inference. Or, reading the theorem to follow in another way, our earlier abstract setting from the end of Section 1 can be represented without loss of generality in concrete DEL update models.

THEOREM 5. *The update-valid structural inferences $\Sigma \rightarrow \sigma$ are precisely those whose conclusions σ are derivable from their premise sets Σ by the rules of Left-Monotonicity, Left-Cut, and Cautious Monotonicity.*

We merely sketch the idea of the proof: for details, cf. van Benthem 2003B. Soundness is immediate, as our structural rules are valid in the special DEL transition models. Completeness uses a two-step representation argument. One first finds an counterexample on some abstract transition model via the earlier representation method (Theorem 1). Next, one transforms such an abstract structure into a concrete family of epistemic models for the states, and announcement actions for the labeled transitions. A bit more detail will be found below.

Modal logic as structural sequent logic again. The preceding style of analysis of structural rules for sequents can be extended to our complete polymodal language. We call a polymodal formula ϕ *update-valid* if every formula of dynamic-epistemic logic resulting from ϕ by uniformly replacing all proposition letters p with standard epistemic formulas, and all atomic actions a with concrete public update actions $!A$ for epistemic logic formulas P, is true in the Supermodel $\mathfrak{M}$ of all epistemic models.

THEOREM 6. *The update-valid modal formulas are axiomatized precisely by the general minimal modal logic of the operators $\langle a \rangle$ and (a) for partial functions a.*

PROOF. We only sketch the heart of the matter. Our crucial observation is

FACT 7. *Any unraveled modal tree model with labelled actions has a bisimilar model consisting of a family of epistemic models, with proposition letters encoded by epistemic S5 formulas, and basic actions a encoded by announcements $!A$.*

More precisely, consider any abstract tree model Δ. Without loss of generality, assume there are unique proposition letters p_x true at each node x. Next, any node x generates a subtree in the usual way, for which we define an epistemic S5-model $\boldsymbol{M}_{\Delta,x}$, whose domain is x's subtree plus a fixed world s. Moreover, every finite S5-model $\boldsymbol{M}$ has a 'descriptive formula' $\delta(\boldsymbol{M})$ true only in $\boldsymbol{M}$ and its bisimulation invariants (cf. van Benthem 2006). Now we are in a position to define the required translations for proposition letters and atomic actions:

> upd(p) is the disjunction of all formulas $\delta(\boldsymbol{M}_{\Delta,x})$ for all x such that $\Delta, x \vDash p$
> upd(a) is the disjunction of all formulas $\delta(\boldsymbol{M}_{\Delta,x})$ & $\left(\bigvee \{p_z \mid z \in \boldsymbol{M}_{\Delta,y}\} \right)$
> for all x, y such that $R_a^\Delta x, y$

These translations lift to arbitrary modal formulas ϕ, taking them to DEL counterparts upd(ϕ). Here is our claim, with $\mathfrak{M}$ again the Supermodel consisting of all epistemic models:

FACT 8. *For all modal formulas ϕ, $\Delta, x \vDash \phi$ iff $\mathfrak{M}, (\boldsymbol{M}_{\Delta,x}, s) \vDash$ upd(ϕ).*

This shows that satisfiable modal formulas have true substitution instances with epistemic update in the Supermodel $\mathfrak{M}$. The converse is simpler. $\mathfrak{M}$ may itself be seen as a modal model. To go from this *class* to a *set*, observe that any satisfiable modal formula at some 'world' $(\boldsymbol{M}, s)$ can also be satisfied in the set of $(\boldsymbol{M}, s)$ and all its submodels, as only these can be reached via update actions. $\square$

There are many further interesting questions about complete logics of these update universes. But our main finding here is this:

The structural rules of abstract dynamic inference are just those for concrete information update and the resulting knowledge resulting from it.

4. Inference, interpolation and preservation

4.1. Entailment along a relation and syntactic interpolants. Moving from a model (M, s) to $(M|A, s)$ via a true public announcement $!A$ is just one case of an important inter-model relation which is relevant to information change: Barwise and van Benthem 1999 consider various others, moving from the above abstract modal framework to specific relations. In this context, with suitable logical languages, such forms of generalized consequence have a special form located in *interpolants* of certain syntactic forms. Here is a characteristic example.

THEOREM 9. *The following are equivalent for all first-order formulas A, B:*
  (a) *A entails B along submodels*
  (b) *there is a* universal formula C *such that $A \vDash C \vDash B$.*

PROOF. The direction from (b) to (a) is immediate since universal formulas are preserved under submodels. Conversely, suppose that A entails B along submodels. Then one proves that $\mathrm{univ}(A) \vDash B$, with $\mathrm{univ}(A)$ the set of universal logical consequences of A. The argument is just like that for the usual Los–Tarski Theorem. Suppose there is no universal interpolant. For any model M of $\mathrm{Univ}(A)$, consider Σ consisting of the atomic diagram of M together with the formula A. This set must be finitely satisfiable – since otherwise, $A \vDash C$ for some universal formula C denying the existence of some finite submodel, but this contradicts the truth of $\mathrm{univ}(A)$ in M. Therefore, the whole set Σ is satisfiable, and there is a model N extending M where A holds. But then, by entailment along submodels, B must hold in M itself. Finally, a simple application of Compactness to $\mathrm{univ}(A) \vDash B$ produces one universal consequence of A which implies B. $\square$

In particular, then, entailment along submodels is recursively enumerable, and hence axiomatizable in principle, for first-order formulas. Moreover, as a special case, Theorem 9 implies the Los–Tarski Preservation Theorem. A first-order formula A is preserved under submodels iff A entails A along submodels, and so A has a universal interpolant C with itself, which makes C equivalent with A.

Here is another result of the same type, again combining an interpolation theorem with a preservation theorem:

THEOREM 10. *The following are equivalent for all first-order formulas A, B:*
  (a) *A entails B along bisimulation in vocabulary L*
  (b) *there is a* modal L-formula C *such that $A \vDash C \vDash B$.*

PROOF. Again, the direction from (b) to (a) is immediate here, as modal formulas are invariant for bisimulation. Conversely, we prove that $\mathrm{mod}(A)$, the set of modal consequences of A, implies the formula B. Again, consider any model $M \vDash \mathrm{mod}(A)$. Using a 'modal diagram' for M this time, there must be a model N

modally equivalent to $\boldsymbol{M}$ where A holds. Now take ω-saturated elementary extensions $\boldsymbol{M}^+$, $\boldsymbol{N}^+$ respectively, and observe as usual that these have a bisimulation running between them. Thus, we have A true in $\boldsymbol{N}$, and in its elementary extension $b\boldsymbol{N}^+$, and then via entailment along bisimulation, B must be true in $\boldsymbol{M}^+$, and hence in $\boldsymbol{M}$. Again, Compactness then gives the required single modal formula which follows from A and implies B. $\qquad\square$

Barwise and van Benthem use the latter type of result, with a generalized proof to deal with potential isomorphisms, to formulate new interpolation theorems for infinitary first-order logic $L_{\infty\omega}$, a logic which lacks Craig Interpolation in the usual sense. A recent survey of new perspectives on the interpolation theorem is van Benthem 2008.

Entailment along a relation and matching interpolation properties have also been considered in the pioneering model-theoretic study Lindström 1966. We refer to that paper for more systematic background to the preceding observations.

The main point of the preceding results is yet one more perspective on general dynamic inference. One now splits a model-crossing form of logical inference into two standard consequences: one from antecedent to interpolant, and the other from interpolant to antecedent. The 'bridge' between antecedent and consequent is then provided by the invariance of the specially constructed interpolant across the relevant inter-model relation. In a slogan,

General consequence equals *standard consequence* plus *invariance*.

It would be of great interest to discover the precise range of this phenomenon.

4.2. Existential variants and higher complexity. Entailment via interpolants along first-order definable relations is itself recursively enumerable (RE), and hence we are still dealing with axiomatizable consequence relations. But simple existential variations can quickly drive up complexity. One example is the situation-theoretic inference of the type 'Smoke Means Fire': "every situation where there is smoke is part of a situation where there is fire". This is the modal 'enabling' pattern $A \to \langle R \rangle B$ of Section 2, but now for concrete model-theoretic relations R. Here is a result which shows the complexity effects of this:

FACT 11. *The general inference notion* $A \to \langle model\text{-}inclusion \rangle B$ *is not RE.*

PROOF. One easily reduces first-order satisfiability to this notion. Consider any first-order formula A, and unary predicate letter P not occurring in it. Then A is satisfiable iff the implication $\forall x\, Px \to \langle inclusion \rangle (A)^{\neg P}$ holds. $\qquad\square$

Indeed, the obvious conjecture is that this sort of extension-entailment is exactly an arithmetical Π_2^0 notion for first-order formulas.

Similar points arise for other widely used inter-model connections, such as the earlier-mentioned modal bisimulation. Even so, here is a standard logical way of expressing in different terms what this sort of entailment says:

PROPOSITION 12. *The existential notion* $A \to \langle model\text{-}inclusion \rangle B$ *is equivalent to* conservativity *of A over B w.r.t.* universal *statements.*

PROOF. (1) First, if B implies some universal sentence C, then so does A. For, let $\boldsymbol{M}$ be any model for A. It has some extension $\boldsymbol{N}$ which is a model for B. Hence C holds in $\boldsymbol{N}$, and by preservation under submodels, C also holds in $\boldsymbol{M}$.

(2) Next, let $\boldsymbol{M}$ be any model for A. Consider the atomic diagram of $\boldsymbol{M}$ together with the formula B. We show that this set is finitely satisfiable. Suppose otherwise. Then B implies some negation of a conjunction of true literals in the $\boldsymbol{M}$-diagram, and – quantifying out the new domain constants – we get a universal consequence of B which is false in $\boldsymbol{M}$, and hence does not follow from A. This refutes the given universal conservativity. $\square$

Conservativity is Π_2^0 – which explains the earlier conjecture. By similar reasoning, we can determine a counterpart for bisimulation and modal formulas:

PROPOSITION 13. *The following are equivalent for first-order formulas A, B:*
  (a) *Each model for A has a bisimilar model where B holds*
  (b) *B is conservative over A with respect to* modal *consequences.*

An independent motivation for 'existential entailment is a phenomenon in modal completeness proofs which may be called '*boosting along bisimulation*'. One first finds a Henkin model for a modal formula ϕ, and then, through techniques like unraveling, bulldozing, duplication etc., one constructs *bisimilar model* satisfying some additional pleasant property α, as well as ϕ because of its bisimulation invariance. This method really depends on a generalized inference of the form

$$\phi \rightarrow \langle \text{bisim} \rangle (\phi \,\&\, \alpha)$$

Here is a question behind many modal completeness techniques.

OPEN PROBLEM 14. *What is the arithmetical complexity of boosting along bisimulation for given first-order formulas ϕ and α?*

Logics of model change. The preceding considerations point to something still more general than DEL as it stands, viz. a dynamic logic of various forms of *model change*. Logical operators which 'look across' models during their evaluation are becoming popular these days, not just in dynamic epistemic logics of information update. They also occur, e.g., in modal logics with so-called 'bisimulation quantifiers' which have already thrown new light on fixed-point logics such as the modal μ-calculus (cf. the chapter by Bradfield and Stirling in Blackburn, van Benthem and Wolter 2006). Thus we see the above observations on model-crossing inference as only the beginning of bringing more structure of the model-theoretic universe into our logics.

Making the vocabulary explicit. Further aspects of inference and model change might be studied in the same spirit. E.g., the preceding results also highlight the role of formal languages and explicit vocabulary in studying inference (van Benthem 2003A). Consider a *ternary* language-dependent notion of consequence $A \vDash B | L$ defined as follows: A implies every L-*consequence of B*. Ordinary valid consequence is $A \vDash B | L_B$, and conservative extension of A by B is $B \vDash A | L_A$ &

$A \vDash B|L_A$. This leads to a new calculus with ternary inferences that may also change vocabulary. E.g., $A \vDash B|L$ and $C \vDash B|L'$ imply $A \vee C \vDash B|L \cap L'$.

Calculi like this link up between logic, theories of abstract data types in computer science, and indeed, calculi of theory structure in the *philosophy of science.*

5. Conclusion

This paper is an exercise in 'logical pluralism'. We have emphasized the entanglement of standard 'inference' with other informational processes such as update through assertions or observations. We have shown that one can still use the familiar format of structural rules to determine the styles of reasoning which emerge then. Moving beyond this level, our claim was that a modal or dynamic language provides a suitable next level for studying abstract properties of general information links. And finally, we have shown how these ideas also make sense with concrete relations between models for first-order logic and other familiar systems. Dynamic inference then give rise to interesting new model-theoretic issues, such as generalized interpolation theorems, and new relations of 'boosting' along model changes. We see all this as a natural fit with 'substructural logics' as put on the map in Dosen and Schroeder-Heister 1994, Restall 2000, though much remains to be explored.

References

[1] A. Baltag, L. Moss and S. Solecki, 1998, *The Logic of Public Announcements, Common Knowledge and Private Suspicions*, Proceedings TARK 1998, 43–56, Morgan Kaufmann Publishers, Los Altos.

[2] J. Barwise and J. van Benthem, 1999, *Interpolation, preservation, and pebble games*, J. Symbolic Logic 64:2, 881–903.

[3] J. Barwise and J. Seligman, 1995, *Information Flow*, Cambridge University Press, Cambridge.

[4] J. van Benthem, 1991, *Language in Action*, North-Holland, Amsterdam.

[5] J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford.

[6] J. van Benthem, 1998, *Dynamic Odds and Ends*, ILLC Research Report ML-98-08, University of Amsterdam.

[7] J. van Benthem, 2003A, *Is there still Logic in Bolzano's Key?*, In: E. Morscher (ed.), *Bernard Bolzano's Leistungen in Logik, Mathematik und Physik*, Bd.16, Academia Verlag, Sankt Augustin, 2003, 11–34.

[8] J. van Benthem, 2003B, *Structural Properties of Dynamic Reasoning*, in: J. Peregrin (ed.), *Meaning: the Dynamic Turn*, Elsevier, Amsterdam, 15–31.

[9] J. van Benthem, 2006, *One is a Lonely Number: on the logic of communication*, In: Z. Chatzidakis, P. Koepke and W. Pohlers (eds.), *Logic Colloquium '02*, ASL and A. K. Peters, Wellesley MA, 96–129.

[10] J. van Benthem, 2007, *Computation as Conversation*, ILLC Amsterdam; In: B. Cooper et al. (eds.), *Computing in Europe 2005*, to appear.

[11] J. van Benthem, J. Gerbrandy and E. Pacuit, 2007, *Merging Frameworks for Interaction: DEL and ETL*, Proceedings TARK.

[12] J. van Benthem 2007, *Logic, Rational Agency, and Intelligent Interaction*, to appear in D. Westerstahl et al., eds., *Proceedings Beijing 2007 DLMPS Conference*.

[13] J. van Benthem 2008, *The Many Faces of Interpolation*, to appear in P. Mancosu, editor, *Special Synthese Issue in Honour of William Craig*.

[14] P. Blackburn, J. van Benthem and F. Wolter 2006, *Handbook of Modal Logic*, Elsevier, Amsterdam.

[15] H. van Ditmarsch, W. van der Hoek and B. Kooi, 2007, *Dynamic Epistemic Logic*. Synthese
      Library, Springer, Berlin.
[16] K. Dosen and P. Schroeder-Heister, eds., 1994, *Substructural Logics*, Clarendon press, Oxford.
[17] P. Lindström, 1966, *On Relations Between Structures*, Theoria 32, 172–185.
[18] G. Restall, 2000, *An Introduction to Substructural Logics*, Routledge, London.

Institute for Logic, Language and Computation (ILLC)
University of Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam
Netherlands
`johan@science.uva.nl`