INFERENCE IN ACTION

Johan van Benthem, Amsterdam & Stanford
http://staff.science.uva.nl/~johan,

January 2006, *second draft*

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) has 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, I wondered how to talk to such an audience? I came up with an example that goes back 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). 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 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:

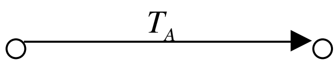

More generally, *transition models* $\boldsymbol{M} = (S, \{T_A \mid A \in Prop\})$ consist of information states S with a family of transition relations T_A over these, one for each proposition A. These suggest the following notion of inference. A sequence of propositions $P_1, \ldots, P_k$ *dynamically 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:

whenever $s_1\, T_{p1}\, s_2 \ldots T_{pk}\, s_{k+1}$, then $s_{k+1}\, C\, s_{k+1}$

We then say *the sequent* $P_1, \ldots, P_k \Rightarrow C$ *is true* in the model – $\boldsymbol{M} \models P_1, \ldots, P_k \Rightarrow C$. 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). Simple counter-examples refute Monotonicity, Contraction, Permutation, or Reflexivity – and their idea is 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:

if $\boldsymbol{P} \Rightarrow A$ and $\boldsymbol{R}, A, \boldsymbol{Q} \Rightarrow C$, 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 rules for dynamic inference:

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

Indeed, these structural rules are characteristic for dynamic inference with partial update functions. Take any set of propositions ***Prop*** as abstract objects – and a 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 (***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 $(S, \{T_A | A \in Prop\})$ 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. From (a) to (b), any abstract structure *(**Prop**, $\Rightarrow$)* induces a transition model ***M*** with states are finite sequences X, Y of propositions. Each proposition A then defines a partial function over these states:

$$T_A \;=\; \{(X, X) \mid X \Rightarrow A\} \cup \{ (X, <X, A>) \mid not\ X \Rightarrow A\}$$

We must check that the following equivalence holds:

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

'If'. Suppose that $s_1\ T_{p1}\ s_2\ \ldots\ T_{pk}\ s_k$. By the definition of the T_A, each step in this sequence either adds a proposition at the end, or 'pauses'. Here is a typical illustration:

$X \;\; T_{p1} \; <X, P_1>$	(not $X \Rightarrow P_1$)
$<X, P_1> \; T_{p2} \; <X, P_1>$	($<X, P_1> \Rightarrow P_2$)
$<X, P_1> \; T_{p3} \; <X, P_1, P_3>$	(not $<X, P_1> \Rightarrow P_3$)

We show that the end state $<X, P_1, P_3>$ is a fixed point for T_C: i.e., $<X, P_1, P_3> \Rightarrow C$. First we have $<P_1, P_2, P_3> \Rightarrow C$, and so by Left-Monotonicity $<X, P_1, P_2, P_3> \Rightarrow C$. Following the transition steps, we suppress one proposition thanks to $<X, P_1> \Rightarrow P_2$, using Left-Cut to get $<X, P_1, P_3> \Rightarrow C$. 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 $<P_1, P_2, P_3>$ dynamically imply C in our transition structure ***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 $<P_1>$; etc. Suppose this yields the following sequence of transformations:

–, $<P_1>$ $\quad$ $<P_1>, <P_1>$ (where P_1 _ P_2!) $\quad$ $<P_1>, <P_1, P_3>$

By assumption, the final state is a fixed point for T_C: P_1 , $P_3 \Rightarrow C$ is true in ***Prop***. But then since $P_1 \Rightarrow P_2$ 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 wherever 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.

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

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

Inference and links across different models Abstract consequence relations often involve just a relation between 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 wide-spread 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 & 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 & 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*** $|=$***P*** and ***M R N***, then ***N***$|=$*C*.

We will discuss properties of this generalized form of inference in Section 4 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. 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's 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 is an obvious notation for the preceding notion, viz. the *modal formula*

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

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 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 as a calculus of *plans*. Van Benthem 1998 presents natural operations on plans with Horn-type rules for them, 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 $A \rightarrow [R]B$, how to derive new ones? Here are a few examples:

from	$A \rightarrow [R]B,\ B \rightarrow [S]C$	*infer*	$A \rightarrow [R\,;\,S]C$	*composition*
from	$A \rightarrow [R]B$	*infer*	$\neg B \rightarrow [R^{\smile}]\neg A$	*converse*
from	$A \rightarrow [R]B,\ A \rightarrow [S]B$	*infer*	$A \rightarrow [R \cup S]\, B$	*union*

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

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

would now explicitly represents actions shifting the relevant model in the passage from premises to conclusion. Richer logics beyond the basic polymodal one may use further operations from dynamic logic here in building the *R*, such as sequential 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:

from	$A \rightarrow [R]B,\ C \rightarrow [S]D$	*infer*	$(A, C) \rightarrow [R\mathbf{x}S]\,(B, D)$	*product*

Validity is easily checked in its first-order transcription:

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

The 'plan calculus' mentioned above (cf. van Benthem 1998) describes valid reasoning with labeled sequents of this sort. It uses monotonicity inferences in antecedents and consequents of such sequents. For the sake of concreteness, here is an illustration:

Excursion: dynamic inference as plan calculus For convenience, one rewrites tagged sequents $A \rightarrow [R]B$ to a format with a backward-looking temporal operator

$P_R A \rightarrow B$, or written with a converse modality: $<R^{\vee}>A \rightarrow B$.

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 even be used to synthetize plans. Consider a resource proposition *A* and a goal proposition *G*. Our available premises encode subroutines available to us:

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

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

1 G from B, C *2 B from A* *3 C from B* *4 B from A*

Composing the associated trees required for this works out to

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

These examples may have shown the interest of taking a polymodal perspective on inference. 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 for our suggestion, 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:

$A \rightarrow <R>B$: *A* makes it possible to execute *R* so that *B* is achieved.

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

$\boldsymbol{M}, s \models (a)\phi$ iff $s\, R_a\, s\ \&\ \boldsymbol{M}, s \models \phi$

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

$(a)\phi \leftrightarrow (a)\mathbf{T}\ \&\ \phi$, $(a)T \rightarrow ([a]\phi \leftrightarrow \phi)$.

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:

$[P_1]..[P_k](C)T$

Fact 3 The structural rules of dynamic inference in Theorem 1 are all valid modal 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 $[\mathbf{P}](A)T$ and $[\mathbf{P}][A][\mathbf{Q}](C)T$ to $[\mathbf{P}][\mathbf{Q}](C)T$.

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

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

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

The loop language can also express complex existential properties of consequence relations beyond mere structural rules. 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 Children scenario in Section 1 supports more concrete scenarios than abstract state transitions, with inference intertwined with information from *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 & Solecki 1998, van Benthem 2006, van Ditmarsch, van der Hoek & 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 $\mathbf{M} = (W, \sim_i, \leq_i, V)$, 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 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 $\mathbf{M}$). E.g., announcing a fact q will make you know that q –though there are more subtle phenomena in general. 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 $(\mathbf{M}, s)$ to a new model $(\mathbf{M}|P, s)$, viz. the model $\mathbf{M}$ restricted to its sub-model consisting of just the P-worlds. To reason about this informational process, we introduce a matching dynamic operator:

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

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*:

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

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 ***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 *C* is known to all agents, or even more strongly, a model where *C* has become common knowledge among them.

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 framework of Section 2. Still, 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 $P_1, \ldots, P_k \Rightarrow \phi$ 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*:

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

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* ***M*** of epistemic models and admissible announcements.

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, where we are working with 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 special formulas obtained in this way, validity in just the above-defined Supermodel, or in arbitrary families of models $\boldsymbol{M}$ as above, makes no difference.

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.

Soundness is immediate here, 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. Next, one transforms such an abstract structure into a concrete family of epistemic models for the states, and announcement actions for the labeled transitions.

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 *A* is true in the Supermodel $\boldsymbol{M}$ of all epistemic models.

Theorem 6 The update-valid modal formulas are axiomatized precisely by the general minimal modal logic of <*a*> and *(a)* for partial functions *a*.

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

Fact 7 Any unraveled modal tree model with labeled 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 $\boldsymbol{\Delta}$. Without loss of generality, assume there are unique proposition letters true at each world. Next, any node *x* generates a subtree in the usual way, for which we define an epistemic *S5*-model $\boldsymbol{M}_{\boldsymbol{\Delta}, x}$ as in the above, whose domain is *x*'s subtree plus a fixed world *s*. Moreover, it is well-known that every finite *S5*-model $\boldsymbol{M}$ has a 'descriptive formula' $\delta(\boldsymbol{M})$ true only in $\boldsymbol{M}$ and its bisimulation invariants (van Benthem 1998, 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 $\boldsymbol{\delta}(\boldsymbol{M}_{\boldsymbol{\Delta}, x})$ for all x such that $\boldsymbol{\Delta}, x \models p$
> *upd(a)* is the disjunction of all formulas $\boldsymbol{\delta}(\boldsymbol{M}_{\boldsymbol{\Delta}, x})$ & $(\bigvee\{p_z \mid z \text{ in } \boldsymbol{M}_{\boldsymbol{\Delta}, y}\})$
> for all x, y such that $R_a^{\boldsymbol{\Delta}}\ x, y$

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

Fact 8 For all modal formulas ϕ, $\boldsymbol{\Delta}, x \models \phi$ iff $\mathbb{M}, (\boldsymbol{M}_{\boldsymbol{\Delta}, x}, s) \models upd(\phi)$

This shows that satisfiable modal formulas have true substitution instances with epistemic update in the Supermodel $\mathbb{M}$. The converse is much simpler. $\mathbb{M}$ 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' *(**M**, s)* can also be satisfied in the set consisting of *(**M**, s)* and all its submodels, since only these can be reached via update actions. ♠

There are many further questions about complete logics of these update universes. But our main finding here is this. The structural rules of abstract dynamic inference are the same as those for concrete information update and the knowledge resulting from it.

4 Interpolation and preservation theorems

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 & van Benthem 1999 consider various others, moving from the above abstract modal framework to specific relations. In this context, they point out that, 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 $A \models C \models B$.

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 *univ(A)* |=*B*, with *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 *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* |= *C* for some universal formula *C* denying the existence of some finite submodel, but this contradicts the truth of *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 *univ(A)* |=*B* produces one universal consequence of *A* which implies *B*. ♠

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* |= *C* |= *B*.

Proof Again, the direction from (b) to (a) is immediate here, as modal formulas are invariant for bisimulation. Conversely, we prove that *mod(A),* the set of modal consequences of *A*, implies the formula *B*. Again, consider any model ***M*** |= *mod(A)*. Using a 'modal diagram' for ***M*** this time, there must be a model ***N*** modally equivalent to ***M*** where *A* holds. Now take ω–saturated elementary extensions M^+, N^+ respectively, and observe as usual that these have a bisimulation running between them. Thus, we have *A* true in ***N***, and in its elementary extension N^+, and then via entailment along bisimulation, *B* must be true in M^+, and hence in ***M***. Again, Compactness then gives the required single modal formula which follows from *A* and implies *B*. ♠

Barwise & 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.

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.

But the main point of the preceding results seems to be this. One splits a general 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.

Existential variants and higher complexity Entailment via interpolants along first-order definable relations is itself *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 $A \rightarrow \langle R \rangle B$ of Section 2, now for concrete model-theoretic relations *R*. Here is a result which shows the complexity effects of this.

Fact 11 The general inference notion $A \rightarrow \langle \textit{model-inclusion} \rangle B$ is not *RE*.

Proof The reason is that 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 $\forall x P x \rightarrow \langle \textit{inclusion} \rangle\, (A)^{\neg P}$ holds. ♠

Indeed, the obvious conjecture is that this sort of extension-entailment is exactly an arithmetical Π^0_2 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 $A \rightarrow \langle \textit{model-inclusion} \rangle B$ 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 ***M*** be any model for *A*. It has some extension ***N*** which is a model for *B*. Therefore, *C* holds in ***N***, and by preservation under submodels, *C* also holds in ***M***. (2) Next, let ***M*** be any model for *A*. Consider the atomic diagram of ***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 ***M***-diagram, and – quantifying out the new domain constants – we get a universal consequence of *B* which is false in ***M***, and hence does not follow from *A*. This refutes the given universal conservativity. ♠

Conservativity is typically Π^0_2 – which explains the earlier conjecture. By quite similar reasoning, we can determine a counterpart for bisimulation and modal formulas:

Proposition 13 The following assertions 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 found 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 shows that there exists a *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

$$\phi \rightarrow \langle bisim \rangle(\phi \;\&\; \alpha).$$

Here is an open problem behind many modal completeness techniques.

Open Problem 14 What is the arithmetical complexity of boosting along bisimulation for given first-order formulas ϕ and ϕ?

Logics of model change The preceding considerations point to something still more general, 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 throw new light on fixed-point logics such as the modal μ–calculus (cf. the chapter by Bradfield & Stirling in Blackburn, van Benthem & Wolter 2006). Thus we see the above 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 high-light the role of formal languages and explicit vocabulary in studying inference (van Benthem 2003A). Consider a *ternary* language-dependent notion of consequence $A \models B \mid L$ defined as follows: *A implies every L-consequence of B*. Ordinary valid consequence is $A \models B \mid L_B$, and conservative extension of *A* by *B* is $B \models A \mid L_A \;\&\; A \models B \mid L_A$. This leads to a new calculus with ternary inferences that may also change vocabulary. E.g., $A \models B \mid L$ and $C \models B \mid L'$ imply $A \vee C \models B \mid L \cap L'$. [1]

[1] . Interesting new questions arise in such a setting. E.g., do $A \models B|L$, $A \models B|L'$ imply that $A \models B|L \cup L'$? The answer is "No" in general, but "Yes" for special simple languages.

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.* [2]

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. One can still use the familiar format of structural rules to determine the styles of reasoning which emerge then. Moving beyond that, a modal or dynamic languages 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. They then give rise to interesting new model-theoretic issues, such as generalized interpolation theorems, and new relations of 'boosting' along model changes.

6 References

J. Barwise & J. van Benthem, 1999, 'Interpolation, Preservation, and Pebble Games', *Journal of Symbolic Logic* 64:2, 881–903.

J. Barwise & J. Seligman, 1995, *Information Flow*, Cambridge University Press, Cambridge.

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

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

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

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.

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

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

J. van Benthem, 2007, 'Computation as Conversation', ILLC Amsterdam, to appear in B. Cooper et al. eds., *Computing in Europe 2005.*

P. Lindström, 1966, 'On Relations Between Structures', *Theoria* 32, 172 – 185.

---

[2] E.g., think of 'Ramsey Eliminability' of theoretical terms in scientific theories, whose logical study turns on extension relations between theories with different vocabularies.