2

STRUCTURAL PROPERTIES OF DYNAMIC REASONING

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

ABSTRACT

We characterize the structural properties of dynamic inference in general update models, and show that these are exactly the ones of public update in epistemic logic.

1 DYNAMIC INFERENCE IN ABSTRACTO

Logical dynamics is about actions that change information, such as information update in communication. At an abstract level, we view propositions A as *partial functions* T_A taking input states meeting the preconditions of update with A to output states:

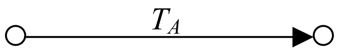

More generally, this process can be modeled by means of *transition models*

$$\boldsymbol{M} = (S, \{T_A \mid A \in \boldsymbol{Prop}\})$$

consisting of the relevant information states S with a family of transition relations T_A over these, one for each proposition A in some abstract index set ***Prop***. Such update propositions suggest the following notion of dynamic inference:

the action of the successive premises enforces the conclusion.

Definition 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 state which is a fixed point for the conclusion:

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

Alternatively, we say *the sequent* $P_1, \ldots, P_k \Rightarrow C$ *is true* in the model – written as:

$\boldsymbol{M} \models P_1, \ldots, P_k \Rightarrow C.$

In what follows, we will use symbols $\boldsymbol{P}$, Q, R to stand for finite sequences of propositions, and A, B, C for single propositions.

Over transition models, dynamic inferential sequents of this form lack the standard structural rules of classical consequence. This situation is analyzed in van Benthem (1996, Chapter 7). Some simple counter-examples establish the following result.

Fact None of the following structural rules hold for dynamic inference: Monotonicity, Contraction, Permutation, Reflexivity, or Cut.

One can show this formally, but the main idea is simply this. Any cooking recipe may be disturbed by inserting arbitrary instructions, deleting repeats of an instruction, interchanging instructions, etc. Even the Cut Rule fails, at least 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 as more often in non-classical logics, some 'substitute rules' turn out to hold.

Fact 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*

Proof For instance, consider Left-Cut. If we move from state s to t via $\boldsymbol{P}$, and then from t to u via $\boldsymbol{Q}$, the first premise $\boldsymbol{P} \Rightarrow A$ tells us that $t\ A\ t$, and so the sequence s, t, t, u fits the action of $\boldsymbol{P}, A, \boldsymbol{Q}$, whence $u\ C\ u$ by the second premise. ■

2 AN ABSTRACT COMPLETENESS THEOREM

The above structural rules are characteristic for dynamic inference with partial update functions. The proper setting for this is the following representation result. Take any set ***Prop*** of propositions, seen as abstract objects – with a binary relation $\Rightarrow$ between finite sequences of propositions and propositions – again written with finite sequents:

Theorem 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, $\{T_A | A \in \textbf{\textit{Prop}}\}$) with partial maps T_A whose relation of dynamic inference as defined above coincides with the given relation $\Rightarrow$ among the abstract propositions A.

Proof The direction from (b) to (a) is the preceding Fact. Now from (a) to (b). For any given abstract structure (***Prop***, $\Rightarrow$), we define a transition model ***M*** as follows. States are finite sequences $X, Y, \ldots$ of propositions. For each proposition A, we then define the following partial function over these states:

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

We must check that the following equivalence holds:

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

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

$X\ T_{p1} <X, P_1>$	(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 must show that the final 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 (in this particular case) that $<P_1, P_2, P_3> \Rightarrow C$ in ***Prop***, and hence by Left-Monotonicity in that structure also

$$<X, P_1, P_2, P_3> \Rightarrow C:$$

Then following the above three transition steps, we can suppress one proposition thanks to the truth of $<X, P_1> \Rightarrow P_2$, by using Left-Cut:

$$<X, P_1, P_3> \Rightarrow C$$

This argument is really completely general. 'Pauses' involve valid sequents that can be used to cut out items in the left sequence $P_1, \ldots, P_k$ at the right places.

'Only if'. This direction essentially involves the remaining structural rule. Again, one example demonstrates the general procedure. Suppose $<P_1, P_2, P_3>$ dynamically imply C in the above

transition structure ***M***. Start with the empty sequence – . We choose three particular transitions for the premises. If $– \Rightarrow P_1$ in ***Prop***, our first transition is –, – ; otherwise, we go to an extended sequence $<P_1>$; etc. Suppose that, for our three propositions, this yields the following sequence of transformations:

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

Now by the assumption of this case, the final state is a fixed point for T_C, i.e.,

$P_1, P_3 \Rightarrow C$ is true in ***Prop***

But then by the fact that $P_1 \Rightarrow P_2$ plus 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. ■

This representation also yields a completeness theorem for sequents interpreted as above on models ***M***. For this purpose, we need to define valid consequence among sequents on transition models, for which we introduce a new arrow:

"from set of sequents Σ to sequent σ" $\quad \Sigma ➔ \sigma$

Definition We have *valid consequence* between a set Σ and a sequent σ if σ is true in all transition models where all sequents from Σ are true.

Here is the more general thrust of the preceding theorem.

Corollary 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.

Proof From right to left, this follows from the soundness of the structural rules. Going from left to right requires a small modification of the above construction. Suppose that σ is not derivable from the set Σ. Take the structure (***Prop***, $\Rightarrow$) with the relation $\Rightarrow$ holding only for

sequents derivable from the set Σ *using the three given structural rules*.

Now represent this structure just as above. The result is a transition model ***M*** where all sequents in Σ are true, while σ is false. ■

There are also other notions of dynamic inference, placing other requirements on the update action associated with the conclusion. Analogous abstract characterizations for their structural properties may be found in van Benthem (1996).

3 CONCRETE MODELS: PUBLIC UPDATE IN EPISTEMIC LOGIC

Epistemic logic and information models

One of the most concrete update systems has public announcements transforming multi-*S5* models for epistemic logic. Here is an illustration. Two players draw from a set of red and white cards. Each can see the colour of their own card, the other cannot. Also, it is known that no two white cards have been drawn. Here is the epistemic model, which may be viewed as an information state for the group {***1***, ***2***}. Its three 'worlds' are tuples like *rw*, standing for the physical situation in which "***1*** has a red card, and ***2*** has a white one". The bold-face tuple ***rr*** represents what actually happened, as seen by an outside observer: both players drew a red card.

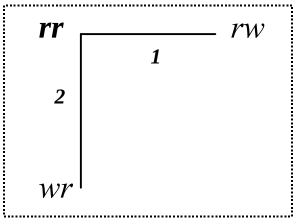

The indexed lines give the usual uncertainty equivalence relations for the two agents. E.g., player ***1*** cannot distinguish the situations *rr* and *rw*. This model ***M*** interprets epistemic assertions in an obvious language with atomic propositions

c_j agent ***j*** has color *c*

and the usual operators

$K_i\phi$ agent ***i*** knows that ϕ
$<i>\phi$ agent ***i*** holds it possible that ϕ

interpreted as universal and existential modalities, respectively. Thus, for instance,

$\mathbf{M}, rw \models K_2 (red_1 \,\&\, white_2)$
$\mathbf{M}, rr \models \neg K_1\, red_2$

Public update by world elimination

Now consider communication between players. Note that ***1*** does not know the color of the other person's card. But this ignorance itself is something she can usefully report (short of broadcasting her own card colour):

"I don't know if ***2*** has the red card"

She knows this, so it is a fair statement. This provides information, as this assertion is false in the bottommost world *wr*. Hence the corresponding public update eliminates this bottom world, and we get an update to the following new epistemic model:

from 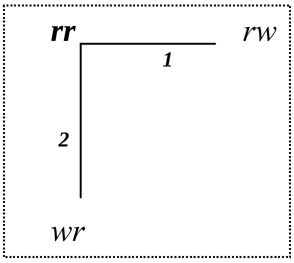

to 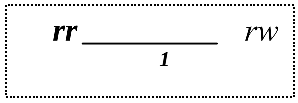

Interestingly, thanks to this communication of ignorance, ***2*** now knows the real situation, as there are no uncertainty lines for him leading out of *rr*. Hence he can now state "I have the red card', leading to a further update eliminating the world *rw*:

from 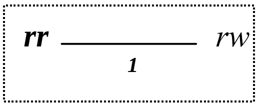

to 

The final model is a information state where the group {***1***, ***2***} has achieved *common knowledge* of the facts, in the technical sense of epistemic logic. (They might just have told each other their cards – but our example illustrates the more interesting forms of communication in general.) In this setting, states of a communication process are epistemic models *(**M**, s)* for the relevant language, with *s* the actual world. These models are described by the usual multi-*S5* epistemic logic, serving as an account of ordinary truth at successive snap-shots of the update process.

Dynamic epistemic logic

But our example also showed how communication itself *changes models* through the use of public announcements ***A!*** of assertions *A* from the epistemic language, viewed as actions in their own right. These actions are partial functions, defined only when the assertion is true in the actual world:

> ***A!*** is defined in *(**M**, s)* only when ***M***, *s* |= *A*. In that case,
> the result is the model *(**M**|**A**, s)* resulting from *(**M**, s)* by
> throwing out all worlds of ***M*** where *A* does not hold:

from 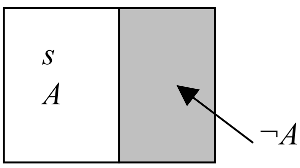

to 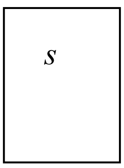

To account for dynamic inference, then, we need a dynamic logic with modalities

$[A!]\phi$, $<A!>\phi$

of the usual kind. E.g., the existential one can be interpreted as follows:

$(\boldsymbol{M}, s) \models <A!>\phi$ iff $\boldsymbol{M}, s \models A$ & $(\boldsymbol{M}|\boldsymbol{A}, s) \models \phi$

A complete logic for this system is known. We state its key axioms here just for the sake of concreteness. These basically show how to 'precompute' effects of updates:

$<A!>p \leftrightarrow p$	basic facts are unaffected by update
$<A!>\neg\phi \leftrightarrow A$ & $\neg<A!>\phi$	updates are partial functions
$<A!>\phi\vee\psi \leftrightarrow <A!>\phi \vee <A!>\psi$	general modal distribution
$<A!>K_a\phi \leftrightarrow A$ & $K_a(A \rightarrow <A!>\phi)$	update through relativization, 1
$<A!>C_G\phi \leftrightarrow A$ & $C_G(A, <A!>\phi)$	update through relativization, 2

The right-hand side of the last clause refers to an appropriate syntactic relativization of the operator C_G of common knowledge for the group G to the formula A. An alternative version of the system has a universal modality $[A!]\phi$ stating that ϕ holds in the new information model updated with A whenever that update *is defined* (i.e., if A is true). Axioms for this are obvious modifications.

Finally, to be precise, we still need a further stipulation. What is the family of all relevant information models? We can take a global perspective here, working with

the *Supermodel* $\boldsymbol{M}$ consisting of all epistemic models.

Then, e.g., the above semantic clause should really read:

$\boldsymbol{M}, (\boldsymbol{M}, s) \models <A!>\phi$ iff $\boldsymbol{M}, (\boldsymbol{M}, s) \models A$ & $\boldsymbol{M}, (\boldsymbol{M}|\boldsymbol{A}, s) \models \phi$

Another option is to restrict the family of all relevant models to 'small corners' of the Supermodel. This approach is natural in modeling games, and other more restricted communicative settings. In that case, all updates will be considered only in so far as they exist within some *special family* $\boldsymbol{M}$ of models. Actual games even have context-dependent restrictions on which updates are possible at which players' turns – but this would take us too far afield here. (Cf. van Benthem (2000).)

4 REPRESENTATION WITH CONCRETE UPDATES

Dynamic inference in communication

Epistemic update logic supports the same notion of dynamic inference that we developed earlier in our general abstract setting. Dynamic propositions are public announcements ***A!*** of epistemic logic formulas A. Again, the preceding choice of settings then returns. Dynamic validity, i.e., truth of a sequent

$$P_1, \ldots, P_k \Rightarrow \phi$$

in the Supermodel says that, starting with any epistemic model whatsoever, successive announcements of the premises result in a model where announcement of ϕ effects no change: i.e., ϕ was already true everywhere even before it was announced. This amounts to validity of the following dynamic formula, which says that the conclusion becomes *common knowledge*:

$$[\mathbf{P_1}!] \ldots [\mathbf{P_k}!]\, C_G\phi$$

But we can also relativize these notions to any smaller family $\boldsymbol{M}$ of epistemic models.

This concrete update system shows the same general phenomena we saw in Sections 1, 2. The three valid structural rules remain valid. More surprisingly, non-validities also show up in this special arena.

Fact All classical structural rules fail for dynamic update inference.

Proof Here are some examples. (a) *Permutation* fails as announcements in the order $\langle\rangle \boldsymbol{p}!;\ \neg \boldsymbol{p}!$ can be consistent, whereas $\neg \boldsymbol{p}!;\ \langle\rangle \boldsymbol{p}!$ is inconsistent: it cannot be announced truthfully. We can cast this in terms of assertions in the dynamic language:

> $[\neg \boldsymbol{p}!][\langle\rangle \boldsymbol{p}!]\ C_G\, q$ is true at every model for any q, as there is no successful update sequence at all, but $[\langle\rangle \boldsymbol{p}!][\neg \boldsymbol{p}!]\ C_G\, q$ fails at some $(\boldsymbol{M}, s)$ for $q \neq p$

(b) Likewise, *Contraction* fails as $(\langle\rangle \boldsymbol{p}!;\ \neg \boldsymbol{p}!)\ ;\ (\langle\rangle \boldsymbol{p}!;\ \neg \boldsymbol{p}!)$ is inconsistent, whereas a single occurrence $(\langle\rangle \boldsymbol{p}!;\ \neg \boldsymbol{p}!)$ is not. The dynamic version is similar to the preceding. (c) The *Cut Rule* fails in the form: "if $\boldsymbol{P} \Rightarrow A$ and $\boldsymbol{Q}$, $A \Rightarrow C$, then $\boldsymbol{Q}$, $\boldsymbol{P} \Rightarrow C$". Counter-example: $\boldsymbol{P} = \neg \boldsymbol{p}!$, $A = \neg\langle\rangle \mathbf{p}$, $\boldsymbol{Q} = \langle\rangle \boldsymbol{p}!$, and $C = \bot$ – or in dynamic terms:

> $[\neg \boldsymbol{p}!]C_G\ \neg\langle\rangle p$ and $[\langle\rangle \boldsymbol{p}!][\neg\langle\rangle \boldsymbol{p}!]\ C_G\, q$ are true everywhere, whereas $[\langle\rangle \boldsymbol{p}!][\neg \boldsymbol{p}!]\ C_G\, q$ is false at some epistemic models.

(d) Finally, *Reflexivity* fails as public announcement $\boldsymbol{A}!$ need not always make the announced assertion A true. For instance, announcing

> $(\boldsymbol{p}\ \&\ \langle \boldsymbol{you}\rangle \neg \boldsymbol{p})!$ (" p holds, but you don't know it")

to someone makes p true in the whole new model, so everyone now knows it – and hence the conjunct $\langle \boldsymbol{you}\rangle \neg \boldsymbol{p}$ becomes false by the very update. This failure is one instance of the 'Learning Problem' for updates. This is the question which assertions once announced become automatically common knowledge, and which do not. ■

The definition of validity

The preceding Fact suggests a notion of valid inference via epistemic substitutions. Take any structural inference involving abstract propositions:

"from set of sequents Σ to sequent σ" $\Sigma \Rightarrow \sigma$

This has *substitution instances* with epistemic announcements replacing the abstract propositions, and sequents $\boldsymbol{P} \Rightarrow \phi$ then turning into dynamic formulas of the above form $[\boldsymbol{P_1}!] \ldots [\boldsymbol{P_k}!] C_G\phi$. Now, each such assertion is true or false in the Supermodel $\boldsymbol{M}$ of all epistemic models. But then, so are Horn implications of the form

conjunction of instances for $\Sigma \rightarrow$ instance for σ
#

reflecting these structural inferences. We call a structural inference *update-valid* if all its substitution instances are true implications of this sort. For an illustration, see our counter-examples for the classical structural rules. In their most obvious sense, these refer to the Supermodel, given the earlier substitutions of concrete updates.

This notion of validity might depend on the universe of available models. In a stricter variant of update validity, the implication # must hold in all families $\boldsymbol{M}$ of epistemic models, small or large. Whenever $\boldsymbol{M}$ makes all premise assertions true (insofar as it contains their relevant update sequences), it must do the same for the conclusion.

Fact A sequent inference $\Sigma \Rightarrow \sigma$ is update-valid in the Supermodel iff it is update-valid in arbitrary families of epistemic models.

We could prove this result separately, but it will fall out of the following analysis.

Public update is complete for dynamic inference

Now we will show that the earlier abstract axiomatization of dynamic inference still works in this more concrete setting of one specific form of information update.

Theorem 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.

Proof Soundness in the Supermodel, or indeed any family of epistemic models, is immediate from soundness of these structural rules over all abstract transition models. We must now consider completeness. If an inference $\Sigma \Rightarrow \sigma$ is not derivable by our structural rules, then the construction in Section 2 creates an update model in which the conclusion $\sigma = P_1, \ldots, P_k \Rightarrow C$ fails, while all premise sequents are globally true. Without loss of generality, we can restrict attention to the finite linear submodel consisting of the states involved in refuting the conclusion. This is a linear sequence

$s_1\ T_{p1}\ s_2\ \ldots\ T_{pk}\ s_{k+1}$, but not $s_{k+1}\ C\ s_{k+1}$

In addition to the transformations mentioned here between these states, there may be others – in particular, fixed points for conclusions of true sequents in the premise set. Now, the point is that such an abstract transition model can be represented faithfully inside the Supermodel ***M***, given a suitable map of abstract states s to concrete epistemic models $(\boldsymbol{M}, t)$, and of abstract transformations T_p to public announcements of specific epistemic formulas. The construction is best seen in a concrete case.

A worked-out example

Consider the following diagram for an abstract transition model:

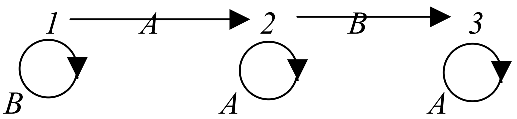

Δ

First consider a multi-S5 model $\boldsymbol{M}_{\boldsymbol{\Delta}}$ with worlds *1, 2, 3* for each point in the diagram, plus one actual world s that will remain constant throughout. Next, take four proposition letters p_1, p_2, p_3 and p_s uniquely true at their corresponding worlds:

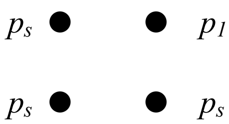

Remark More precisely, one should use Boolean combinations of proposition letters, stipulating that they form a partition of the set of worlds in the usual manner.

Now, the model $\boldsymbol{M}_{\boldsymbol{\Delta}}$ can be described globally by means of one epistemic formula $\delta_{\boldsymbol{M}}$:

$<>p_1\ \&\ <>p_2\ \&\ <>p_3\ \&\ <>p_s\ \&\ [](p_1 \vee p_2 \vee p_3 \vee p_s)$

This formula holds only in $\boldsymbol{M}_\Delta$ and all models bisimilar to it. Now the idea is that B is going to update models, starting from this one, in a way mimicking its action in the given diagram $\boldsymbol{\Delta}$. This will be as a public announcement of a *disjunction* of separate formulas, one for each B-arrow shown. For a start, in the initial model B has a fixed point. We write the following conjunction

$B_1\ \ (<>p_1\ \&\ <>p_2\ \&\ <>p_3\ \&\ <>p_s\ \&\ [](p_1 \vee p_2 \vee p_3 \vee p_s))\ \&\ (p_1 \vee p_2 \vee p_3 \vee p_s)$

This has a global conjunct to the left which is only true in the initial model, and a local conjunct to the right picking out specific worlds inside that model. Therefore, in any family of epistemic models, and even in the full Supermodel, an announcement ***B1!*** only leads to an update on the given model $\boldsymbol{M}_{\Delta}$ and its bisimulation invariants, not even on any of its submodels (as it fails there) – and it holds for all worlds, changing nothing. Next, we assume that the first *A* has already reduced the model to the subdomain {*s, 2, 3*}, and then, *B* has to properly change the model to the universe {*s, 3*}. This is done via a disjunct whose global conjunct can only be true for worlds in the new *{s, 2, 3}*–model, and whose specific conjunct picks out the right worlds:

$B_2 \;\; (\langle\rangle p_2 \,\&\, \langle\rangle p_3 \,\&\, \langle\rangle p_s \,\&\, [\,](p_2 \vee p_3 \vee p_s)) \;\&\; (p_3 \vee p_s)$

There are no further *B*-arrows in the diagram, so we say

B is the disjunction $B_1 \vee B_2$

Likewise, we define *A* as the disjunction $A_1 \vee A_2 \vee A_3$ with

$A_1 \;\; (\langle\rangle p_1 \,\&\, \langle\rangle p_2 \,\&\, \langle\rangle p_3 \,\&\, \langle\rangle p_s \,\&\, [\,](p_1 \vee p_2 \vee p_3 \vee p_s)) \;\&\; (p_2 \vee p_3 \vee p_s)$
$A_2 \;\; (\langle\rangle p_2 \,\&\, \langle\rangle p_3 \,\&\, \langle\rangle p_s \,\&\, [\,](p_2 \vee p_3 \vee p_s)) \;\&\; (p_2 \vee p_3 \vee p_s)$
$A_3 \;\; (\langle\rangle p_3 \,\&\, \langle\rangle p_s \,\&\, [\,](p_3 \vee p_s)) \;\&\; (p_3 \vee p_s)$

Given this construction, it is easy to show the following:

The public announcements ***A!*** and ***B!*** produce an update pattern in the Supermodel ***M*** of all epistemic models, starting from the model $\boldsymbol{M}_{\Delta}$, which is exactly like the given transition diagram **Δ** .

The reason is that, given the above epistemic substitutions, the relevant sequents true in the diagram are true in the Supermodel, while the conclusion sequent fails there.

The general reasoning

Now let us state more generally what is happening here. We want to show that

If a sequent inference $\Sigma \Rightarrow \sigma$ is not derivable by our structural rules,
then it has some false substitution instance in the Supermodel.

First, by our abstract completeness theorem, the non-derivability implies the existence of a transition model Δ where Σ holds, while σ fails. The above construction then provides epistemic formulas ϕ_a for each relevant atomic action *a* in Δ. Replacing all these actions in our sequents by their epistemic counterparts in an obvious manner, and passing to the corresponding dynamic formulas as earlier in this Section, we get

$sub(\Sigma)$, $sub(\sigma)$.

Claim $sub(\sigma)$ fails in the Supermodel.

To see this, note that, starting from the epistemic model $(\boldsymbol{M}_\Delta, s)$, announcements of the epistemic formulas ϕ_a produce an isomorphic copy of $\boldsymbol{\Delta}$ itself – where σ failed.

Claim $sub(\Sigma)$ holds in the Supermodel.

This time, we must show that, starting at *any* $(\boldsymbol{M}, s)$ in the Supermodel, a sequence of successive updates for the premises of any of the sequents in $sub(\Sigma)$ ends in a fixed point for its final assertion. Now, the formulas ϕ_a defined above can only hold at *S5*-models that are bisimilar to rooted submodels of $\boldsymbol{M}_\Delta$, at exactly those stages where a could be performed. (The only possible variation is trivial: they might have more copies of the same world around.) In other words, performing the announcements from the premise sequence produces a sequence of models which is still essentially an update sequence inside $\boldsymbol{M}_\Delta$, and hence its final state will be a fixed-point for the conclusion, because all sequents of Σ were true in $\boldsymbol{\Delta}$, and hence also in $\boldsymbol{M}_\Delta$.

We will analyze this still informal argument more formally in the next section. In its general version, it matches evaluation of arbitrary modal formulas in abstract models with evaluation of suitably translated formulas in correlated epistemic update models.

Summarizing, any abstract situation making all sequents in Σ true but the conclusion sequent σ false, can be reproduced exactly with epistemic update in the Supermodel through a choice of suitable substitution instances. The general construction uses exactly the above tricks, with descriptive δ–formulas codifying the stages, and disjunctions of proposition letters cutting down the model at that stage. ■

As a corollary, we get the earlier statement that update validity of sequent inferences in the Supermodel and validity over arbitrary families of epistemic models coincide.

Proof By general soundness, if a sequent inference is refuted in some family of epistemic models, it is not derivable from the given three structural rules. Therefore, it must have a counterexample in the Supermodel by the above argument. ■

Our result shows how curious phenomena may happen with public announcement:

> Every abstract transition diagram satisfying the three structural conditions of Left-Monotonicity, Left-Cut, and Cautious Monotonicity can be mimicked with concrete updates.

These diagrams can be extremely diverse, and hence so can phenomena occurring with public update. In particular, as in the above key illustration, dynamic propositions may be dormant at some stages of a communication process, ruling out nothing at all – and then become active again, following the action of other dynamic propositions.

5 MODAL LOGIC OF INFERENCE

Modal logics of abstract updates

Structural rules are very simple properties of a notion of consequence. There is also a richer modal logic describing such properties. The language has modalities $<a>$, $[a]$ for basic update actions a, interpreted again as *partial functions* over arbitrary models, plus a 'loop modality' (a) saying the following:

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

This seems the minimum required for formalizing the dynamic inference of Section 1. This modal language is decidable with a finite model property, and a simple complete axiomatization (van Benthem (1996, Chapter 7)). Its key loop axioms are these:

$$(a)\phi \leftrightarrow (a)T \;\&\; \phi$$
$$(a)T \rightarrow ([a]\phi \leftrightarrow \phi)$$

This allows us to read dynamic sequents $P_1, \ldots, P_k \Rightarrow C$ as modal formulas

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

Fact In their modal versions, the structural rules of dynamic inference are valid consequences from premises true in a whole model to their conclusions.

Proof We outline the relevant observations.

(a) Left-Monotonicity from $[\boldsymbol{P}](C)T$ to $[A][\boldsymbol{P}](C)T$
is an instance of the modal Necessitation Rule

(b) Left-Cut from $[\boldsymbol{P}](A)T$ and $[\boldsymbol{P}][A][\boldsymbol{Q}](C)T$ to $[\boldsymbol{P}][\boldsymbol{Q}](C)T$
is a modalized consequence of the loop axiom $((A)T \;\&\; [A]\phi) \rightarrow \phi$

(c) Cautious Monotonicity from $[\boldsymbol{P}](A)$T and $[\boldsymbol{P}][\boldsymbol{Q}](C)T$ to $[\boldsymbol{P}][A][\boldsymbol{Q}](C)T$
is a modalized consequence of the loop axiom $((A)T \;\&\; \phi) \rightarrow [A]\phi$

■

But it is easy to see that this richer language can also express properties of a notion of consequence that go beyond mere structural rules, such as more existential modal formats like $([P]<Q>\phi \;\&\; [Q]<R>\psi) \rightarrow [P]<Q>(\phi \;\&<R>\psi))$. And therefore, it is an interesting candidate for a richer, though still abstract theory of dynamic inference.

This modal language expresses *local properties*, true at specific states of a model. But sequent validity in our original sense involves global truth of the modal formulas for the premises throughout a model. We can view this as a case of so-called 'global consequence' for modal logic. Another option is to enrich the modal language with a *universal modality* $U\phi$ stating that

ϕ is true at every world in the current model. Then, we can express our sequent inferences by prefixing the premise formulas with *U*.

Formal definitions again

Again, there is a notion of substitution instance for modal formulas of the local language in the earlier epistemic update calculus. Call such a formula ϕ *update-valid* if

> every formula of epistemic update 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 models.

As before, the latter system completely determines our new abstract update logic.

Public update is also complete for modal update validity

Here is the relevant result.

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

Proof By the finite model property for the modal language with loops, one only needs to consider finite models, and indeed, finite *unraveled trees*. Now, we extend the concrete representation of Section 4. Any modal tree model with labeled actions gives rise to a finite family of epistemic models, with its proposition letters encoded by epistemic *S5* formulas, and its basic actions *a* mimicked by update actions ***A!***.

Again, we start with an example. The tricks employed earlier work just as well for tree models as for 'lines'. Let the abstract model look like this:

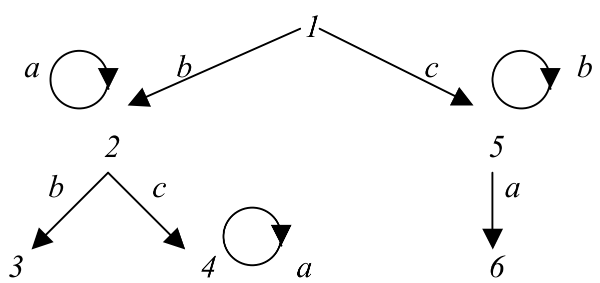

Again, one creates a corresponding epistemic model, with universe *{s, 1, 2, 3, 4, 5, 6}*. As before, the world *s* is thrown in to make sure some 'actual world' persists through the update sequence. Substitution formulas arise as follows. E.g., one codes the action of the given atomic action *b* via public announcement of the disjunction of

(i) the descriptive δ–formula for the whole model **&** the disjunction $(p_s \vee p_2 \vee p_3 \vee p_4)$,
(ii) the δ–formula of the model with domain *{s, 2, 3, 4}* **&** $(p_s \vee p_3)$,

(iii) the δ–formula of the domain $\{s, 5, 6\}$ **&** $(p_s \vee p_5 \vee p_6)$

Proposition letters may be coded even more simply by using the disjunction of the descriptive formulas for the submodels corresponding to the nodes where they hold.

Now for formalities. Consider any abstract tree model $\boldsymbol{\Delta}$ for our poly-modal language. As before, without loss of generality, we may 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, every *S5*-model $\boldsymbol{M}$ has an obvious 'descriptive formula' $\delta(\boldsymbol{M})$ true only in $\boldsymbol{M}$ and its bisimulation invariants (being just versions of $\boldsymbol{M}$ with perhaps duplications of worlds). 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 $\delta(\boldsymbol{M}_{\boldsymbol{\Delta},x})$ for all x such that $\boldsymbol{\Delta}, x \models p$

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

These two translations evidently lift to take arbitrary modal formulas ϕ to obvious counterparts $upd(\phi)$ in the dynamic-epistemic language of updates. Here is our claim, with $\boldsymbol{M}$ again the Supermodel consisting of all epistemic models:

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

The proof is by induction. (a) For proposition letters p, note that $\boldsymbol{M}_{\boldsymbol{\Delta},x}$ validates *upd(p)* if and only if it satisfies some formula $\delta(\boldsymbol{M}_{\boldsymbol{\Delta},y})$ with $\boldsymbol{\Delta}, y \models p$, and these formulas determine their 'source' uniquely. (b) The case of Boolean operations is automatic. (c) For modalities $\langle a \rangle$, we follow the construction of the formulas *upd(a)*. E.g., if $\boldsymbol{\Delta}, x \models \langle a \rangle \phi$, then there is a y with $R_a^{\boldsymbol{\Delta}}\, x, y$ and $\boldsymbol{\Delta}, y \models \phi$. By the inductive hypothesis, $(\boldsymbol{M}_{\boldsymbol{\Delta},y}, \boldsymbol{s}) \models upd(\phi)$. Now, announcing the true disjunction *upd(a)* at the model $(\boldsymbol{M}_{\boldsymbol{\Delta},x}, s)$ will only 'trigger' its unique disjunct with the right prefix $\delta(\boldsymbol{M}_{\boldsymbol{\Delta},x})$, and the worlds remaining after the update are precisely those in $\boldsymbol{M}_{\boldsymbol{\Delta},y}$. (Here we use the fact that all atomic actions are partial functions.) For the converse, assume that $\boldsymbol{M}$, $(\boldsymbol{M}_{\boldsymbol{\Delta},x}, s) \models \langle upd(a)! \rangle\, upd(\phi)$. Then one of the formulas $\delta(\boldsymbol{M}_{\boldsymbol{\Delta},x})$ & $(\bigvee\{p_z \mid z \text{ in } \boldsymbol{M}_{\boldsymbol{\Delta},y})$ from *upd(a)* holds in $(\boldsymbol{M}_{\boldsymbol{\Delta},x}, s)$. This can only be the one for x itself, and therefore the update takes us to a model $\boldsymbol{M}_{\boldsymbol{\Delta},\boldsymbol{y}}$ whose root node y could be reached in $\boldsymbol{\Delta}$ from x via some a-arrow. The rest follows by the inductive hypothesis. (d) Finally, the argument for the loop modality is essentially the same to the preceding one.

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

As we observed earlier, however, sequent inference is not purely local. It involved global truth of premise formulas at all worlds in the Supermodel. The reason why the earlier result worked is that the modal formulas occurring behind that universal modality were very special, namely *universal* statements of the form $[a_1]...[a_k](b)$. It is not hard to see that the above construction will show that formulas of this form true throughout the model $\boldsymbol{\Delta}$. will also be true globally in the Supermodel.

Can we be more ambitious than this, and prove the above theorem for the extended modal language with the universal modality? The answer is negative. The Supermodel satisfies modal laws which do not hold in general. Here is one example, reflecting the fact that there are 'endpoints', viz. epistemic models containing just one world:

$\neg(U(\phi \to \langle A!\rangle\neg\phi)\ \&\ U(\neg\phi \to \langle A!\rangle\phi))$ holds in $\boldsymbol{M}$, for all ϕ, A

But the poly-modal formula $U(p \to \langle a\rangle\neg p)\ \&\ U(\neg p \to \langle a\rangle p)$ *is* satisfiable in

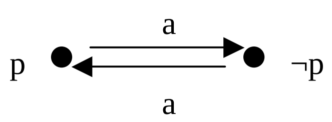

What would be the complete poly-modal theory of the Supermodel in this sense?

6 DISCUSSION

Our representation results show that concrete update systems can be complete for structural theories of abstract logical dynamics. This means, in particular, that all dynamic behaviour that can occur in general transition models will already show with plain public update over epistemic models.

Conversely, we can also look at the above as a new type of technical result in modal logic, viz. *axiomatizing* part of *the meta-theory* of certain operations on models. In our case, ***A!*** is a natural operation of *relativization* of models to *A*–definable submodels.

The above representation of abstract transition models requires only *single-agent epistemic models*. This may have to do with the above trick of defining all needed worlds by unique new proposition letters. If we fix some finite number of these in advance, say just $\{p, q\}$, then many-agent epistemic models with updates involving genuine stackings of *K*–operators for different agents may become indispensable. Related to this is an issue of *complexity*. Our method is a sort of reduction of satisfiability for the minimal modal logic (whose complexity is PSPACE-complete) to that for a dynamic update version of *S5*, which latter system is only NP-complete. But no miracle has happened here. We will get a blow up in the proposition letters needed to encode the worlds, and also, the complexity of the epistemic update system is most probably higher than that of *S5* itself (though this question is still open).

Next, the *Supermodel* of all epistemic models is an unrealistic update universe. One often knows beforehand that only certain announcements can be made – as in games of information. Then the appropriate structures are much more like the above local update models, viewed as *trees of epistemic models*, related by only some out of all possible update actions. In that case, some of the earlier principles of update logic must be modified. In particular, we no longer have a valid update law like $p \rightarrow \langle \boldsymbol{p!} \rangle T$. Even though the atomic fact p is true now, $\langle \boldsymbol{p!} \rangle T$ may fail in case there is no admissible announcement of this fact in our local update model.

Finally, the dynamic update logic with assertions $[A\boldsymbol{!}]\phi$ in Section 3 is *itself* a sort of structural calculus for logical dynamics. Why be even more general than that? What we have tried to find in this paper is a workable level of abstraction just above this, by dropping specific information about the atomic actions of public announcement – without going stratospheric in still more general mathematical approaches.

REFERENCES

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

van Benthem, J. (2000). Logics for Information Update. In: *Proceedings TARK VIII*, pp. 51–88. Morgan Kaufmann, Los Altos.