Action Emulation

Jan van Eijck
CWI, Amsterdam and Uil-OTS, Utrecht

Ji Ruan
University of Liverpool

Tomasz Sadzik
Stanford University

Abstract. The effects of public announcements, private communications, deceptive messages to groups, and so on, can all be captured by a general mechanism of updating multi-agent models with update action models [3], now in widespread use (see [10] for a textbook treatment). There is a natural extension of the definition of a bisimulation to action models. Surely enough, updating with bisimilar action models gives the same result (modulo bisimulation). But the converse turns out to be false: update models may have the same update effects without being bisimilar. We propose action emulation as a notion of structural equivalence more appropriate for action models, and generalizing standard bisimulation. It is proved that action emulation provides a full characterization of update effect, provided we confine attention to 'smooth' action models. We also give a recipe for turning any action model into a smooth one with the same update effect. Together, this yields a simplification procedure for action models, and it gives designers of multi-agent systems a useful tool for comparing different ways of representing a particular communicative action.

1. Introduction

Knowledge and (lack of) knowledge about knowledge plays a key role in the interaction of agents. In systems that describe the effects of public announcements and group announcements (such as emails with lots of cc's), common knowledge is created: everyone on a cc list knows that the content of the email is now common knowledge. To reason about the effects of such communicative actions one needs a powerful logic that can (at least) express common knowledge, and the effects of communication on common knowledge.

In epistemic logic [12] knowledge is represented with multi-agent Kripke models (or possible world models) that contain for each agent an accessibility relation pointing at the situations that the agent considers possible. To talk about what is the case in such models, a logical language is used that allows one to express things like 'agent a considers ϕ possible' (this would express that ϕ is consistent with what a knows or believes), or 'in all states that are linked to the current state via a and

© 2008 *Kluwer Academic Publishers. Printed in the Netherlands.*

b accessibilities, ϕ is the case' (this would express common knowledge of a and b that ϕ).

While standard epistemic logics do not directly represent acts of communication, Dynamic Epistemic Logic (DEL) does: it introduces the representation of *actions*, and a method of updating a situation with these actions. For an overview of developments in these areas, consult Gerbrandy [14], van Ditmarsch [9], van Benthem [5, 6], and Baltag, Moss and coworkers [3, 1, 2]. Perhaps the most streamlined version of DEL so far is the logic of communication and change (LCC) of [7], which is the version of DEL that we will adopt here. LCC is in fact propositional dynamic logic [17, 16] with the programs interpreted as composite epistemic modalities, and with modal operators added for describing the effects of communicative actions (epistemic PDL with action modalities).

The basic insight of DEL is from [3]: a wide variety of information updates can be treated using a formal product construction with an action model, which is nothing but a multi-agent Kripke model with the valuations replaced by precondition formulas. The reason for this to work is that actions with epistemic effects are quite similar to situations with epistemic aspects. The uncertainty of agents about which action takes place is a lot like the uncertainty of agents about what is the case.

If you receive a message ϕ and I am left in the dark, then this is modeled as an action that allows you to distinguish the ϕ situations from the rest, while I am not allowed to make that distinction. If the two of us get the ϕ message, and some outsider does not, then it makes a real difference whether the two of us know of each other that we get the same information, and this again is encoded in the action model.

Since action model updating is rapidly becoming the standard for modeling communicative action, it is important for multi-agent system design to have means of comparing different ways of representing a particular communicative action. In this paper, we study equivalence of action models: two action models are equivalent if they always produce non-distinguishable results. Our contribution is a concept called *action emulation*, and a proof that this precisely characterizes this equivalence.

Just as bisimulation allows to compute bisimulation minimal models, action emulation allows the computation of minimal forms for action models, at least in principle: our results imply that minimal models under action emulation exist, and it allows us to compute minimal versions of action models in principle. Computation may be costly, for it involves satisfiability checks for conjunctions of preconditions, and the satisfiability problem for our language is the same as that for PDL: it is in EXPTIME [8]. Computation of minimal action models is important for epistemic model checking. In fact, the initial motivation of our quest was the need for a means to simplify action models in the epistemic model checking tool DEMO [11].

The structure of the paper is as follows. In Section 2, we review the version of Dynamic Epistemic Logic we work with, motivate our choice, and define our basic notions. Section 3 gives a definition of equivalence or 'same update effect' for action models that we want to capture, compares this notion to that of bisimulation for action models, and gives examples to show that these notions do not quite match. Section 4 presents our proposal for a structural notion to match action model equivalence, and shows that this definition coincides with equivalence for the case of action models with purely propositional preconditions. Most of our everyday communications are like this. We exchange factual information, deciding whether to send cc's or not, we decide to keep some facts to ourselves, or only tell them to a few close friends. The epistemic pattern of *how* the information is conveyed may be incredibly complex, as when we decide to send private letters of invitation to a large group of acquaintances, but with a cc to our spouse. The section ends with examples of action models where the characterization fails. In Section 6 we define the important notion of a smooth action model, and then, after some preliminaries in Section 7, give a recipe for turning action models into equivalent smooth ones (in Section 8), and show that for all smooth action models, equivalence implies existence of an action emulation. Section 9 gives discussion and questions for further research.

2. Dynamic Epistemic Logic

In this section we formally introduce epistemic models (or multi agent Kripke models), followed by definitions of action models and a suitable epistemic language. Next, we define the process of updating with and action model and the notion of truth in a model.

Epistemic models capture a static description of what agents know about the world and about each other, action models capture the instructions for modifying these static systems. In all definitions we assume that a finite set of agents *Ag* and a set of propositional variables *Prop* are given.

DEFINITION 1. (Epistemic model). *An epistemic model is a triple* $M = (W, V, R)$ *where* W *is a set of worlds,* $V : W \to \mathcal{P}(\mathrm{Prop})$ *assigns a valuation to each world* $w \in W$*, and* $R : \mathrm{Ag} \to \mathcal{P}(W^2)$ *assigns an accessibility relation* $\stackrel{i}{\to}$ *to each agent* $i \in \mathrm{Ag}$.

If M is an epistemic model, we use W_M t‘¡o refer to its set of worlds, V_M to refer to its valuation function, and R_M to refer to its accessibility function.

DEFINITION 2. (Distinctive epistemic model). *A distinctive epistemic model is a pair* (M, U) *with* M *an epistemic model and* $U \subseteq W_M$*. The intended interpretation of the distinction is that the actual world is among* U*.* $\mathbf{M}, \mathbf{N}$ *will be used to refer to distinctive epistemic models.*

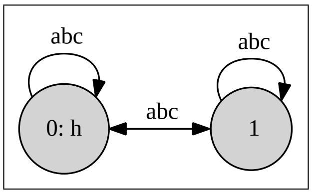

Figure 1. Epistemic model representing the result of a hidden coin toss.

Figure 1 gives an example of an epistemic model that describes the result of a hidden coin toss, with three onlookers, Alice, Bob and Carol. The model has two distinguished situations, marked in grey. Presence of proposition letter h in a world indicates that the valuation makes h true in that world, absence of h in a world indicates that the valuation makes h false in that world, so the picture reveals that the coin has landed heads up in world 0, tails up in world 1. The epistemic indiscernibility relations are indicated by arrows, with labels indicating the agents. None of the agents can tell these two worlds apart. The distinction tells us that the actual world could be either world. The example illustrates the convenience of allowing more than one distinguished world. If one insists on a single distinguished world per model (the actual world), then the situation of our example cannot be described by a single model.

Figure 2 gives a situation like that of Figure 1, but where the coin actually has landed heads up, although nobody knows this yet. A test

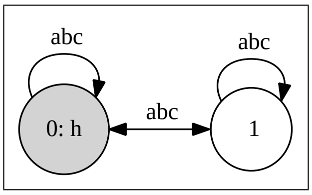

Figure 2. Epistemic model representing that the result of a hidden coin toss is *heads*.

with the claim that the coin has landed *heads* up can be viewed as an action of making the *tail* situation non-distinguished. It changes the model from Figure 1 into that of Figure 2.

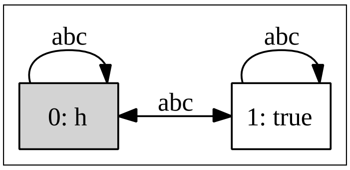

Figure 3. Action model for a hidden test that the coin landed *heads* up.

Baltag, Moss and Solecki [3] proposed to model update actions on epistemic models as taking a product with action models, where action models are like epistemic models, but with valuations replaced by precondition formulas. In the example of Figure 3, the actual action (in grey) is that formula h is checked. The agents a, b and c cannot distinguish this action from an action where formula *true* is checked (nothing happens at all). The result of updating with this action model should be that non-h situations get removed from the set of distinguished worlds. In other words, updating the epistemic model in Figure 1 with the action model in Figure 3 should yield the epistemic model in Figure 2.

DEFINITION 3. (Action model for language $\mathcal{L}$). *An action model for* $\mathcal{L}$ *is a triple* $A = (W, pre, R)$ *where* W *is a set of action states,* $pre : W \to \mathcal{L}$ *assigns a precondition to each action state, and* $R : \mathrm{Ag} \to \mathcal{P}(W^2)$ *assigns an accessibility relation* $\stackrel{i}{\to}$ *to each agent* $i \in \mathrm{Ag}$.

As in the case of epistemic models, we use W_A for the set of action states of action model A, pre_A for its precondition function, and R_A for its accessibility function.

DEFINITION 4. (Distinctive action model for language $\mathcal{L}$).
A distinctive action model is a pair (A, S) *with* A *an action model and* $S \subseteq W_A$. *The distinction indicates that the action that actually takes place is a member of* S.

Henceforth, we use 'action model' for action model with and without distinctive points. If we wish to stress the presence of a set of distinctive points, we use 'distinctive action model'.

The epistemic language $\mathcal{LANG}$ that we are going to use for the preconditions is epistemic PDL with action modalities. It is defined as follows.

DEFINITION 5. ($\mathcal{LANG}$). *Let* p *range over the set of basic propositions* P *and* i *over the set of agents* Ag. *The formulas of* $\mathcal{LANG}$ *are* *given by:*

$$\begin{aligned}\phi &::= \mathit{true} \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\alpha]\phi \mid \ [A,S]\phi,\\ \alpha &::= i \mid ?\phi \mid \alpha_1 \cup \alpha_2 \mid \alpha_1;\alpha_2 \mid \alpha^*,\end{aligned}$$

where (A, S) is a distinctive finite $\mathcal{LANG}$ action model. Let $\mathcal{LANG}_0$ be the result of removing all formulas of the form $[A, S]\phi$ from the language.

Note that $\mathcal{L}ANG_0$ is in fact epistemic PDL. We employ the usual abbreviations. In particular, *false* is shorthand for $\neg\mathit{true}$, $\phi_1 \vee \phi_2$ for $\neg(\neg\phi_1 \wedge \neg\phi_2)$, $\phi_1 \rightarrow \phi_2$ for $\neg(\phi_1 \wedge \neg\phi_2)$, $\langle\alpha\rangle\phi$ for $\neg[\alpha]\neg\phi$, $\langle A, S\rangle\phi$ for $\neg[A, S]\neg\phi$. Also, we will use $\bigvee\{\phi_1, \ldots, \phi_n\}$ for $\phi_1 \vee \cdots \vee \phi_n$ and $\bigwedge\{\phi_1, \ldots, \phi_n\}$ for $\phi_1 \wedge \cdots \wedge \phi_n$.

This language is more expressive than the usual epistemic languages, even apart from the presence of action models as modal operators. For instance, K_i, the knowledge operator for i, looks like $[i]$, while the common knowledge for i and j is given by $[(i \cup j)^*]$. Relativized common knowledge is expressed by $[(?\phi; i \cup j)^*]$: this epistemic operator talks about paths along i and j links, where everywhere on the path ϕ is true. The advantage of this expressiveness reveals itself when one wants to describe the epistemic effects of (say) public announcements. How can we express epistemically that public announcement of ϕ created the common knowledge among i and j that ψ? For describing the effect of the announcement, we have to describe the situation in the model *before the announcement*, and common knowledge relativized to ϕ expresses just this. A language that has operators for common knowledge and for public announcement cannot be complete unless it can also express relativized common knowledge (see [3]).

The advantage of the availability of regular epistemic programs α turns out to be even more formidable: Van Benthem, Van Eijck and Kooi prove in [7] that for language $\mathcal{LANG}$, the effects of every update action with an action model A, S can be described in purely epistemic terms. This means that in Section 6 we can work with the sub-language $\mathcal{LANG}_0$ instead of the full language $\mathcal{LANG}$, for everything that can be said in $\mathcal{L}ANG$ can also be expressed in $\mathcal{L}ANG_0$.

In distinctive action models, the distinguished points constrain the whereabouts of the actual action, in distinctive epistemic models they constrain the whereabouts of the actual world. There are a number of reasons for employing distinctive models for representing epistemic situations and for updating.

- In the first place, it allows us to generalize over a number of situations in a straightforward way, witness the example of Figure

1, where we could leave it non-determined whether the actual situation is a *heads* or a *tails* situation.

– Similarly for the action models: it allows us to handle choice in a straightforward way. Consider an action of revealing the truth about h as a choice between announcing h if *heads* is the case, and announcing $\neg h$ if *tails* is the case.

– It simplifies the definition of the update process. Updates are always defined, but they may result in epistemic models with an empty domain or with an empty set of distinguished worlds.

In Section 4 we will give a definition of distinctive action emulation that relates the distinctive points of two action models to each other.

Let MOD be the class of distinctive epistemic models and ACT the class of distinctive finite $\mathcal{LANG}$ models. Then $\mathcal{LANG}$-update is an operation of the following type:

$$\otimes : \text{MOD} \times \text{ACT} \to \text{MOD}.$$

The operation $\otimes$ and the truth definition for $\mathcal{LANG}$ are defined by mutual recursion, as follows.

DEFINITION 6. (Update, Truth). *Given a distinctive epistemic model* (M, U) *and a distinctive action model* (A, S)*, we define*

$$(M, U) \otimes (A, S)$$

as

$$((W', V', R'), U'),$$

where

$$\begin{array}{lcl} W' & := & \{(w,s) \mid w \in W_M, s \in W_A, M \models_w pre_s\}, \\ V'(w,s) & := & V_M(w), \\ (w,s) \xrightarrow{i} (w',s') \in R' & :\equiv & w \xrightarrow{i} w' \in R_M, s \xrightarrow{i} s' \in R_A, \\ U' & := & (U \times S) \cap W', \end{array}$$

and where the truth definition is given by:

$$\begin{array}{lcl} M \models_w true & & always \\ M \models_w p & :\equiv & p \in V_M(w) \\ M \models_w \neg\phi & :\equiv & not\ M \models_w \phi \\ M \models_w \phi_1 \wedge \phi_2 & :\equiv & M \models_w \phi_1\ and\ M \models_w \phi_2 \\ M \models_w [\alpha]\phi & :\equiv & for\ all\ w'\ with\ w \xrightarrow{\alpha} w'\ M \models_{w'} \phi \\ M \models_w [A,S]\phi & :\equiv & M' \models_{(w,s)} \phi\ for\ all\ s \in S\ with\ M \models_w \mathrm{pre}_s, \\ & & where\ M' = (M, \{w\}) \otimes (A, S), \end{array}$$

with $\stackrel{\alpha}{\rightarrow}$ *given by*

$$
\begin{array}{lcll}
\stackrel{i}{\rightarrow} & = & R_A(i) & \\
\stackrel{?\phi}{\rightarrow} & = & \{(x,x) \mid M \models_x \phi\} & \\
\stackrel{\alpha_1 \cup \alpha_2}{\rightarrow} & = & \stackrel{\alpha_1}{\rightarrow} \cup \stackrel{\alpha_2}{\rightarrow} & \\
\stackrel{\alpha_1;\alpha_2}{\rightarrow} & = & \stackrel{\alpha_1}{\rightarrow} \circ \stackrel{\alpha_2}{\rightarrow} & \textit{(relational composition of } \stackrel{\alpha_1}{\rightarrow} \textit{ and } \stackrel{\alpha_2}{\rightarrow}\textit{)} \\
\stackrel{\alpha^*}{\rightarrow} & = & (\stackrel{\alpha}{\rightarrow})^* & \textit{(transitive closure of } \stackrel{\alpha}{\rightarrow}\textit{).}
\end{array}
$$

As an illustration, Figure 4 gives the result of updating the epistemic model from Figure 1 with the action model from Figure 3.

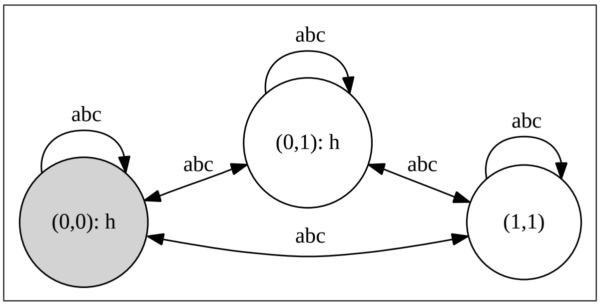

Figure 4. Result of updating model from Figure 1 with action model from Figure 3.

3. Bisimulation versus Action Equivalence

The standard notion of structural equivalence for epistemic models is bisimulation.

DEFINITION 7. (Bisimulation). *Let* M, N *be epistemic models. The relation* $C \subseteq W_M \times W_N$ *is a bisimulation if whenever* wCv *the following hold:*

Invariance $V_M(w) = V_N(v)$,

Zig *for all* $i \in \mathrm{Ag}$, *all worlds* $w' \in W_M$ *with* $w \stackrel{i}{\rightarrow} w'$ *there is a state* $v' \in W_N$ *with* $v \stackrel{i}{\rightarrow} v'$ *and* $w'Cv'$,

Zag *same requirement vice versa.*

DEFINITION 8. (Distinctive Bisimulation). *A distinctive bisimulation between distinctive epistemic models* (M, X) *and* (N, Y) *is a bisimulation* C *between* M *and* N *that is such that for each* $x \in X$ *there is a* $y \in Y$ *with* xCy*, and vice versa.*

Existence of a distinctive bisimulation between $\mathbf{M}$ *and* $\mathbf{N}$ *is indicated by* $\mathbf{M} \underline{\leftrightarrow} \mathbf{N}$.

Note that there is a distinctive bisimulation between the models from Figures 2 and 4. This shows that it is indeed the case that Figure 2 can be viewed as the result of updating the model from Figure 1 with the action model from Figure 3.

Thinking of the action models as 'update programs', the basic semantic notion of equivalence between such programs is that of having the same update effect: two distinctive action models are equivalent if applied to the same epistemic model, they yield bisimilar results. Formally:

DEFINITION 9. (Action equivalence). *Two action models* (A, S) *and* (B, T) *are equivalent, notation* $(A, S) \equiv (B, T)$*, if it holds for all distinctive epistemic models* $\mathbf{M}$ *that* $\mathbf{M} \otimes (A, S) \underline{\leftrightarrow} \mathbf{M} \otimes (B, T)$.

We would like to capture this notion of equivalence by means of a structural definition. Some suitable generalization of bisimulation for the case of action models is an obvious candidate for this. A natural generalization suggests itself: simply replace the invariance requirement of 'having the same valuation' by a requirement of 'having equivalent preconditions'. Since the only difference between epistemic models and action model is in the switch from valuations to preconditions, this seems an obvious choice. A demand of syntactic equality of presuppositions would be too strong, but logical equivalence seems just right. This gives:

DEFINITION 10. (Bisimulation for action models). *Let* A, B *be action models, and let* $\equiv$ *be the appropriate logical equivalence notion for their precondition language. A relation* $C \subseteq W_A \times W_B$ *is a bisimulation if whenever* sCt *the following hold:*

Invariance pre_s *is consistent, and* $\mathrm{pre}_s \equiv \mathrm{pre}_t$,

Zig *for all* $i \in \mathrm{Ag}$ *and all states* $s' \in W_A$ *with* $s \xrightarrow{i} s'$ *there is a state* $t' \in W_B$ *with* $t \xrightarrow{i} t'$ *and* $s'Cs'$,

Zag *same requirement vice versa.*

The consistency requirement on the preconditions ensures that actions with inconsistent preconditions do not end up in the bisimulation relation.

DEFINITION 11. (Distinctive bisimulation for action models).
A distinctive bisimulation between (A, S) and (B, T) is a bisimulation that connects each $s \in S$ to some $t \in T$, and vice versa.
Existence of a distinctive bisimulation between (A, S) and (B, T) is indicated by $(A, S) \underline{\leftrightarrow} (B, T)$.

Surely, bisimilar action models in the sense of Definition 11 are equivalent in the sense of Definition 9:

THEOREM 12. *If (A, S) and (B, T) are action models, then*

$$(A, S) \underline{\leftrightarrow} (B, T) \text{ implies } (A, S) \equiv (B, T).$$

Proof. We have to show that for any epistemic model (M, U) and every (u, s_i) among the distinctive worlds of $(M, U) \otimes (A, S)$ there is a (v, t_j) among the distinctive worlds of $(M, U) \otimes (B, T)$ with $(u, s_i) \underline{\leftrightarrow} (v, t_j)$, and vice versa. This follows immediately from the existence of the distinctive bisimulation $\underline{\leftrightarrow}$ between (A, S) and (B, T), for the relation on $M \otimes A \times M \otimes B$ defined by means of

$$(u, s)C(v, t) \text{ iff } u = v \text{ and } s \underline{\leftrightarrow} t$$

is a bisimulation. □

Henceforth, it is convenient to restrict attention to *consistent action models*, i.e., action models with consistent preconditions. Clearly, for every action model there is a consistent equivalent action model: just restrict the model to the actions with consistent preconditions. Now here is a key observation.

OBSERVATION 13. *The equivalence of two action models does not imply their bisimilarity.*

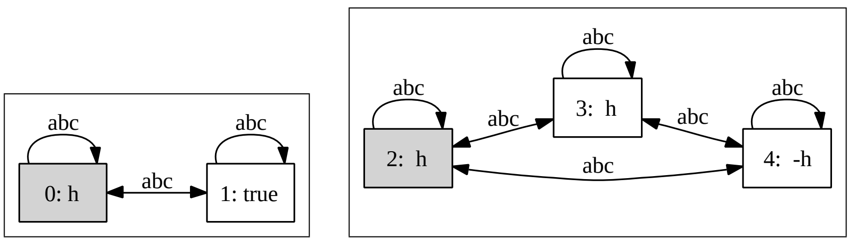

Figure 5. A pair of equivalent, but non-bisimilar action models.

Figure 5 provides an example of two action models for which there is no distinctive bisimulation. The distinctive states of the two action models have the same precondition, but, e.g., the step $0 \xrightarrow{a} 1$ in the left action model cannot be matched by a step from distinctive state 2 in the right action model, for that model has no states with precondition *true*. The update effects of the two action models are the same: they both represent a hidden test for h. Similarly for the two action models

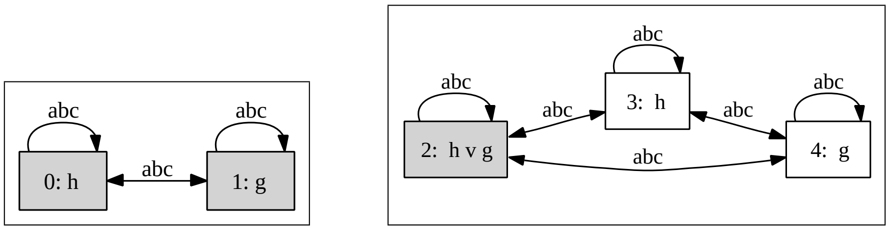

Figure 6. Another pair of equivalent action models that are not bisimilar.

in Figure 6. They are equivalent, for they both have the effect of a public announcement of $h \vee g$. Think of the public announcement that at least one of two tossed coins has landed *heads* up. Again, there is no distinctive bisimulation between the two action models.

Examples like these suggest that the bisimilarity notion in Definition 10 does not quite capture the 'essence' of update actions, and they motivate the quest for a more appropriate structural relation.

4. Action Emulation

One possible reaction to the misfit between action bisimulation and action equivalence would be to conclude that our generalization of bisimulation to the case of action models is flawed. The problem with this view is that a more natural definition of action bisimulation is hard to come by. Also, action bisimulation, the way we defined it, seems a useful notion in itself. For these reasons we regard the notion of action emulation that follows not as a reformulation of bisimulation, but as a genuine alternative.

DEFINITION 14. (Action Emulation). *Given action models A and B, a relation $E \subseteq W_A \times W_B$ is an action emulation if whenever sEt the following hold:*

Invariance *$pre_s \wedge pre_t$ is consistent.*

Zig *If $s \xrightarrow{i} s'$ then there are $t_1, \ldots, t_n$ with*

$$t \xrightarrow{i} t_1, \ldots, t \xrightarrow{i} t_n, s'Et_1, \ldots, s'Et_n \text{ and } pre_{s'} \models \bigvee_{1 \leq j \leq n} \mathrm{pre}_{t_j},$$

Zag *If* $t \xrightarrow{i} t'$ *then there are* $s_1, \ldots, s_n$ *with*

$$s \xrightarrow{i} s_1, \ldots, s \xrightarrow{i} s_n, s_1Et', \ldots, s_nEt' \text{ and } pre_{t'} \models \bigvee_{1 \leq j \leq n} \mathrm{pre}_{s_j}.$$

DEFINITION 15. (Distinctive Action Emulation).
A relation $E \subseteq W_A \times W_B$ *is a distinctive action emulation between distinctive action models* (A, S) *and* (B, T) *if* E *is an action emulation between* A *and* B *such that*

for every $s \in S$ *the set* $T' = \{t \in T \mid sEt\}$ *is non-empty and satisfies* $\mathrm{pre}_s \models \bigvee_{t \in T'} pre_t$,

for every $t \in T$ *the set* $S' = \{s \in S \mid sEt\}$ *is non-empty and satisfies* $\mathrm{pre}_t \models \bigvee_{s \in S'} pre_s$.

Existence of a distinctive action emulation between (A, S) *and* (B, T) *is indicated by* $(A, S) \leftrightarrows (B, T)$.

Examples of (distinctive) action emulations are the relation

$$E_1 = \{(0, 2), (1, 3), (1, 4)\}$$

between the domains of the left and the right action model in Figure 5, and the relation

$$E_2 = \{(0, 2), (1, 2), (0, 3), (1, 4)\}$$

between the domains of the left and the right action model in Figure 6.

It is easily checked that every distinctive bisimulation is a distinctive action emulation: the three conditions of action emulation follow from the three conditions of action bisimulation. This gives:

PROPOSITION 16. *If* (A, S) *and* (B, T) *are action models, then*

$$(A, S) \underline{\leftrightarrow} (B, T) \text{ implies } (A, S) \leftrightarrows (B, T).$$

The following theorem shows that existence of an action emulation is a sufficient condition for equivalence.

THEOREM 17. *For all action models* (A, S) *and* (B, T),

$$(A, S) \leftrightarrows (B, T) \text{ implies } (A, S) \equiv (B, T).$$

Proof. Let (M, X) be an epistemic model. Assume $(A, S) \leftrightarrows (B, T)$, and let E be an action emulation witnessing this.

Define C between the domains of $M \otimes A$ and $M \otimes B$ by means of: $(w, s)C(v, t) :\equiv w = v \wedge sEt$. In order to show that C is a bisimulation, suppose $(w, s)C(v, t)$.

Invariance From $(w,s)C(v,t)$ we get that $w=v$ and hence $V(w,s) = V(v,t)$.

Zig Let $(w,s) \xrightarrow{i} (w',s')$. Then $w \xrightarrow{i} w'$, $s \xrightarrow{i} s'$, and $M \models_{w'} \text{pre}_{s'}$. From $(w,s)C(v,t)$ we have that sEt. By sEt, there are $t_1,\dots,t_n \in W_B$ with

$$t \xrightarrow{i} t_1,\dots,t \xrightarrow{i} t_n, s'Et_1,\dots,s'Et_n \text{ and } \text{pre}_{s'} \models \bigvee_{1\le j\le n} \text{pre}_{t_j}.$$

Since $M \models_{w'} \text{pre}_{s'}$, it follows from $\text{pre}_{s'} \models \bigvee_{1\le j\le n} \text{pre}_{t_j}$ that there is some t_j with $M \models_{w'} pre_{t_j}$. Thus $(w',s')C(w',t_j)$.

Zag Same reasoning vice versa.

We still have to show that C connects the distinctive points of $(M,X)\otimes(A,S)$ and $(M,X)\otimes(B,T)$. Given $(w,s)\in M\otimes A$ with $w\in X$ and $s\in S$, we have $M\models_w \text{pre}_s$. Since E connects (A,S) and (B,T), there must be $t_1,\dots,t_n$, such that $sEt_1,\dots,sEt_n$ and $\text{pre}_s \models \bigvee_{1\le j\le n}\text{pre}_{t_j}$, hence $M\models_w \bigvee_{1\le j\le n}\text{pre}_{t_j}$. So there must be a t_j such that $M\models_w \text{pre}_{t_j}$, and therefore $(w,s)C(w,t_j)$. The other direction is similar. □

Now for a surprise: the converse of Theorem 17 turns out to be false. Some counterexamples are in Figures 7 and 8.

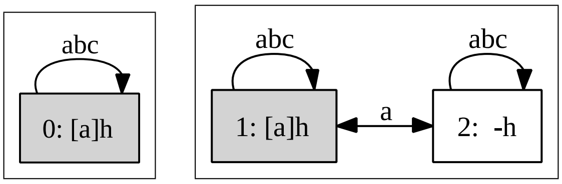

Figure 7. A pair of equivalent action models that do not emulate.

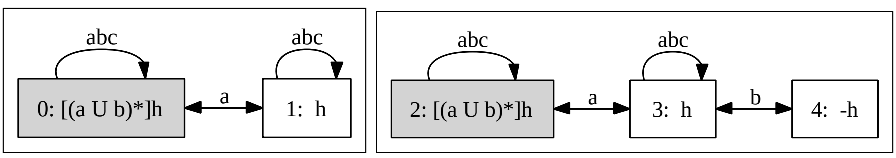

Figure 8. Another pair of equivalent action models that do not emulate.

To see that the action models of Figure 7 are equivalent, first observe that the action model on the left expresses a public announcement $[a]h$ (a public announcement "Alice knows that *heads* has turned up").

The action model on the right describes a communication where $[a]h$ gets announced, but Alice confuses this with the announcement of $\neg h$ (the public announcement "*no heads*", i.e., "*tails* has turned up"). The update result of this is the same as that of the action model on the left, for pairs $(w, 1)$ in the result of updating M with the action model on the right will have to satisfy $M, w \models [a]h$, and therefore $M, v \models h$ will hold for all v with $w \xrightarrow{a} v$. Since the precondition of 2 is $\neg h$, there will be no pairs $(v, 2)$ with $(w, 1) \xrightarrow{a} (v, 2)$ in the update result. There may be arrows $(v, 2) \xrightarrow{a} (w, 1)$, but these will not be reachable from a distinguished state in the update result. In a similar way, it can be shown that the action models of Figure 8 are equivalent. Again, there is no action emulation between these action models.

To end the section in a positive mood, here is a proof that equivalence of action models with purely propositional preconditions implies the existence of an emulation.

THEOREM 18. *Let (A, S) and (B, T) be action models with every precondition occurring in A or B purely propositional. Then*

$$(A, S) \equiv (B, T) \text{ implies } (A, S) \leftrightarrows (B, T).$$

Proof. Let Q be the set of all proposition letters occurring in the preconditions of A and B. Let M be the epistemic model (W, V, R) where $W = \mathcal{P}(Q)$, V is the identity function, and R is the function that assigns the universal relation on W^2 to every agent i. From $(A, S) \equiv (B, T)$ it follows that

$$(M, W) \otimes (A, S) \quad \underline{\leftrightarrow} \quad (M, W) \otimes (B, T).$$

Now define a binary relation $E \subseteq W_A \times W_B$ by means of

$$sEt :\equiv \text{ there is a } v \in W \text{ such that } (v, s) \underline{\leftrightarrow} (v, t).$$

We show that E is an action emulation. Suppose sEt. Then

Invariance By the definition of E, it follows from sEt that for some $v \in W$, $(v, s) \underline{\leftrightarrow} (v, t)$. From this we get that $(w, s) \in W_{M \otimes A}$ and $(w, t) \in W_{M \otimes B}$, and therefore $M \models_w pre_s$ and $M \models_w pre_t$. Thus, $pre_s \wedge pre_t$ is consistent.

Zig Suppose $s \xrightarrow{i} s'$. Every valuation $v \subseteq Q$ has a corresponding formula $\overline{v}$, with $\overline{v}$ given by

$$\overline{v} = \bigwedge_{q \in v} q \wedge \bigwedge_{q \in (Q - v)} \neg q.$$

For every purely propositional formula ϕ with propositional variables in Q there is a non-empty set $\{v_1, \dots, v_n\}$ of worlds in W

(i.e., subsets of Q) with

$$\phi \equiv \bigvee_{1 \leq j \leq n} \overline{v_j}.$$

Applying this to $pre_{s'}$, we get a non-empty set $\{v_1, \ldots, v_n\}$ of worlds in W with $pre_{s'} \equiv \bigvee_{1 \leq j \leq n} \overline{v_j}$. Since the accessibility relation in M is universal, we have $v \xrightarrow{i} v_j$ for all j with $1 \leq j \leq n$. Therefore $(v, s) \xrightarrow{i} (v_j, s')$ for all such j. Now use $(v, s) \underline{\leftrightarrow} (v, t)$ to infer that there must be a non-empty set $\{t_1, \ldots, t_n\}$ with $(v, t) \xrightarrow{i} (v_j, t_j)$ and $(v_j, s') \underline{\leftrightarrow} (v_j, t_j)$, for all j with $1 \leq j \leq n$. Since $M \models_{v_j} pre_{t_j}$ for all j with $1 \leq j \leq n$, it follows that $pre_{s'} \models pre_{t_1} \vee \cdots \vee pre_{t_n}$.

Zag Same reasoning vice versa.

It is easy to check that E connects every state pair $(s, t) \in S \times T$ for which $pre_s \wedge pre_t$ is consistent. □

5. Partial Evaluation in Action Models

Let A be an action model. Let Σ be the set of preconditions occurring in A, and let p range over the proposition letters occurring in Σ. Call this the modal language of A:

$$\phi \ ::= \ true \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [i]\phi$$

We define verification and falsification of formulas from the language of A in A. The definition of smooth action models (Def 24) in Section 6) will be phrased in terms of these notions.

DEFINITION 19. (Verification and Falsification in Action Models).

$A \models_s true$		*always*
$A \ =\!\mid_s true$		*never*
$A \models_s p$	*if*	$\text{pre}_s \models p$
$A \ =\!\mid_s p$	*if*	$\text{pre}_s \models \neg p$ *(i.e.,* $\text{pre}_s \wedge p$ *inconsistent)*
$A \models_s \neg\phi$	*if*	$A \ =\!\mid_s \phi$
$A \ =\!\mid_s \neg\phi$	*if*	$A \models_s \phi$
$A \models_s \phi_1 \wedge \phi_2$	*if*	$A \models_s \phi_1$ *and* $A \models_s \phi_2$
$A \ =\!\mid_s \phi_1 \wedge \phi_2$	*if*	$A \ =\!\mid_s \phi_1$ *or* $A \ =\!\mid_s \phi_2$
$A \models_s [i]\phi$	*if*	*either* $\text{pre}_s \models [i]\phi$ *or*
		for all s' *with* $s \xrightarrow{i} s' : A \models_{s'} \phi$
$A \ =\!\mid_s [i]\phi$	*if*	*either* $\text{pre}_s \models \neg[i]\phi$ *or*
		for some s' *with* $s \xrightarrow{i} s' : A \ =\!\mid_{s'} \phi$

Read $A \models_s \phi$ as "s verifies ϕ in A", and $A \;=\!\!|_s\; \phi$ as "s falsifies ϕ in A". Note that the clauses for truth and falsity of p and of $[i]\phi$ use the concept of logical consequence for the logic of Σ. The evaluation uses the strong Kleene scheme (invented by Kleene [15] to describe the behaviour of partial recursive functions, where the evaluation procedure can loop), extended with the (rather obvious) treatment of the modal operators. Since the case of both true and false is not excluded, we have in fact a Belnap-style four valued system [4].

For convenience, we work out the verification and falsification rules for $\vee$ and $\langle i \rangle$ from the definitions:

$$
\begin{array}{lll}
A \models_s \phi_1 \vee \phi_2 & \text{if} & A \models_s \phi_1 \text{ or } A \models_s \phi_2 \\
A \;=\!\!|_s\; \phi_1 \vee \phi_2 & \text{if} & A \;=\!\!|_s\; \phi_1 \text{ and } A \;=\!\!|_s\; \phi_2 \\
A \models_s \langle i \rangle \phi & \text{if} & \text{either } pre_s \models \langle i \rangle \phi \text{ or} \\
 & & \text{for some } s' \text{ with } s \xrightarrow{i} s' : A \models_{s'} \phi \\
A \;=\!\!|_s\; \langle i \rangle \phi & \text{if} & \text{either } pre_s \models \neg \langle i \rangle \phi \text{ or} \\
 & & \text{for all } s' \text{ with } s \xrightarrow{i} s' : A \;=\!\!|_{s'}\; \phi
\end{array}
$$

For an example where neither $A \models_s \phi$ nor $A \;=\!\!|_s\; \phi$, think of an action model with a single state s with precondition p. Then $A \not\models_s q$ and $A \not=\!\!|_s q$, for q does not follow from the precondition of s nor is q inconsistent with the precondition of s. But we can also have contradictions! Take an action model A with two states s and s' with $s \xrightarrow{i} s'$ (and no other pairs in R_i), with $pre_s = [i]p$ and $pre_{s'} = \neg p$. Then $A \models_s [i]p$ because $[i]p$ equals the precondition of s, but also $A \;=\!\!|_s\; [i]p$, because for the only i-accessible state, s', it holds that $A \;=\!\!|_{s'}\; p$ (since p is inconsistent with the precondition of s').

For a full treatment we will have to extend strong Kleene evaluation to the full language $\mathcal{L}ANG_0$, but it is rather obvious how to do this.

LEMMA 20. *For all action models A, all A-states s, all ϕ in the language of A:* $\text{pre}_s \models \phi$ *implies* $A \models_s \phi$, *and* $\text{pre}_s \models \neg\phi$ *implies* $A \;=\!\!|_s\; \phi$.

Proof. Induction on the structure of ϕ. If ϕ equals *true* the statement certainly holds. If ϕ equals p, then $pre_s \models p$ implies $A \models_s p$ by the definition of verification, and $pre_s \models \neg p$ implies $A \;=\!\!|_s\; p$, again by the definition of verification.

Assume the statement holds for ϕ_1, ϕ_2. We show that it also holds for $\neg\phi_1, \phi_1 \wedge \phi_2, [i]\phi_1$.

Assume $pre_s \models \neg\phi_1$. We have to show $A \models_s \neg\phi_1$. This follows immediately from $A \;=\!\!|_s\; \phi$ (by induction hypothesis) and the definition of verification for $\neg$. Assume $pre_s \models \neg\neg\phi_1$. Then $pre_s \models \phi_1$, and, by induction hypothesis, $A \models_s \phi$, and hence $A \;=\!\!|_s\; \neg\phi$.

Assume $pre_s \models \phi_1 \wedge \phi_2$. Then $pre_s \models \phi_1$ and $pre_s \models \phi_2$, hence by twice the induction hypothesis, $A \models_s \phi_1$ and $A \models_s \phi_2$. Hence $A \models_s \phi_1 \wedge \phi_2$. Assume $pre_s \models \neg(\phi_1 \wedge \phi_2)$. Then $pre_s \models \neg\phi_1$ or $pre_s \models \neg\phi_2$. Therefore, by induction hypothesis, either $A \dashv_s \phi_1$ or $A \dashv_s \phi_2$. Hence $A \dashv_s \phi_1 \wedge \phi_2$.

Assume $pre_s \models [i]\phi_1$. Then $A \models_s [i]\phi_1$ by definition of verification for $[i]$. Assume $pre_s \models \neg[i]\phi_1$. Then $A \dashv_s [i]\phi_1$ by the clause for falsification of $[i]$. □

If we know that the formulas are purely propositional, we can turn this around:

LEMMA 21. *For all action models A, all A-states s, all ϕ in the purely propositional fragment of the language of A:*

$$A \models_s \phi \textit{ iff } \mathrm{pre}_s \models \phi, \textit{ and } A \dashv_s \phi \textit{ iff } \mathrm{pre}_s \models \neg\phi.$$

Proof. Induction on the propositional structure of ϕ. If ϕ equals *true* the statement certainly holds. If ϕ equals p, then $A \models_s p$ iff $pre_s \models p$ by the definition of verification, and $A \dashv_s p$ iff $pre_s \models \neg p$, again by the definition of verification.

Assume the statement holds for ϕ_1, ϕ_2. We show that it also holds for $\neg\phi_1$ and $\phi_1 \wedge \phi_2$.

$A \models_s \neg\phi_1$ iff (definition of verification) $A \dashv_s \phi_1$ iff (induction hypothesis) $pre_s \models \neg\phi_1$. $A \dashv_s \neg\phi_1$ iff (definition of falsification) $A \models_s \phi_1$ iff (induction hypothesis) $pre_s \models \phi_1$ iff $pre_s \models \neg\neg\phi_1$.

$A \models_s \phi_1 \wedge \phi_2$ iff (definition of verification) $A \models_s \phi_1$ and $A \models_s \phi_2$ iff (induction hypothesis) $pre_s \models \phi_1$ and $pre_s \models \phi_2$ iff $pre_s \models \phi_1 \wedge \phi_2$.

$A \dashv_s \phi_1 \wedge \phi_2$ iff (definition of falsification) $A \dashv_s \phi_1$ or $A \dashv_s \phi_2$ iff (induction hypothesis) $pre_s \models \neg\phi_1$ or $pre_s \models \neg\phi_2$ iff $pre_s \models \neg\phi_1 \vee \neg\phi_2$ iff $pre_s \models \neg(\phi_1 \wedge \phi_2)$. □

Finally, we note for the record that the relations $\models_s$ is monotonic and the relation $\dashv_s$ is anti-monotonic, in the following sense:

PROPOSITION 22. *$A \models_s \phi$ and $\phi \models \psi$ implies $A \models_s \psi$, and $A \dashv_s \phi$ and $\psi \models \phi$ implies $A \dashv_s \psi$.*

6. Smoothness

DEFINITION 23. *Call a formula ϕ a* **constraint formula** *if ϕ is equivalent to a formula built from purely propositional formulas by* *means of* $[i]$, $\wedge$ *and* $\vee$. *I.e, a constraint formula is a formula in the language*

$$\phi ::= \psi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid [i]\phi,$$

where ψ *is purely propositional.*

So here is a way to single out the smooth action models: smooth are the action models where every constraint formula that follows logically from the precondition of a state is consistent with that state.

DEFINITION 24. *Action model* A *is* **smooth** *if for all* $s \in W_A$ *and all constraint formulas* ϕ *in the language of* A*: if* $\text{pre}_s \models \phi$ *then* $A \not\models\!\!\!\dashv_s \phi$.

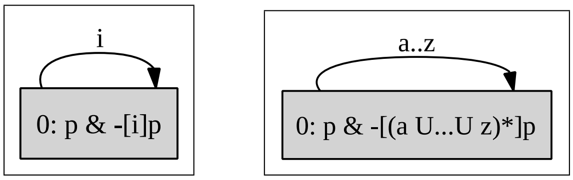

Figure 9. Action models for the Moore announcement (left) and the Al Gore announcement (right).

The reason to restrict the consistency requirement to constraint formulas is to allow action models with preconditions that 'falsify themselves', such as an announcement of the Moore sentence $p \wedge \neg[i]p$. The action model for this (Figure 9 left, with $\wedge$ written as `&`) has formulas that are falsified at state 0 while they also follow from the precondition of 0: we have $p \wedge \neg[i]p \models \neg[i]p$, but also $\dashv_0 \neg[i]p$. Since p is true at 0, and 0 is the only i-accessible state from 0, $\models_0 [i]p$, and therefore $\dashv_0 \neg[i]p$. This does not contradict smoothness; the precondition $p \wedge \neg[i]p$ does not constrain the accessibility relation in the action model itself, for the precondition has no non-trivial modal consequences of the form $[i]\phi_1 \vee \cdots \vee [i]\phi_n$. A similar thing holds for the Al Gore announcement $p \wedge \neg[(a_1 \cup \cdots \cup a_n)^*]p$ in Figure 9, right.

Note that smoothness is a natural extension of consistency:

PROPOSITION 25. *If* A *is smooth then all preconditions in* A *are consistent.*

Proof. Suppose A is smooth. Let s be a state in A and assume pre_s is inconsistent. Then $pre_s \models false$ and $A \dashv_s false$, and since *false* is a constraint formula, it follows that A is not smooth.

PROPOSITION 26. *If each precondition of* A *is a consistent purely propositional formula, then* A *is smooth.*

Proof. Let A be an action model with each precondition a consistent purely propositional formula. Let s be a state in A, and let ϕ be a constraint formula such that $pre_s \models \phi$. Then since pre_s is purely propositional, ϕ is equivalent to a purely propositional formula. Because pre_s is consistent, $pre_s \models \phi$ implies $pre_s \not\models \neg\phi$. Lemma 21 gives that $A \not\dashv_s \phi$.

THEOREM 27. *Let A be a smooth action model. Then for no state s in the domain of A and no constraint formula ϕ it holds that $A \dashv_s \phi$ and $A \models_s \phi$.*

Proof. Induction on the structure of the constraint formula. For purely propositional constraints the property holds by Lemma 21 and the consistency of pre_s.

Assume the property holds for ϕ_1 and ϕ_2. Then the property is easy to prove for $\phi_1 \wedge \phi_2$ and $\phi_1 \vee \phi_2$. Finally, suppose $A \dashv_s [i]\phi_1$ and $A \models_s [i]\phi_1$. There are three cases: (i) $pre_s \models [i]\phi$, (ii) $pre_s \models \neg[i]\phi$, and (iii) neither of $pre_s \models [i]\phi$ and $pre_s \models \neg[i]\phi$. In case (i), it follows from the smoothness of A that $A \not\dashv [i]\phi$, and contradiction. In case (ii), it follows from the smoothness of A that $A \not\dashv \neg[i]\phi$. From this, by the clause for falsification, $A \not\models [i]\phi$, and contradiction. In case (iii), we get from $A \dashv_s [i]\phi_1$ by the definition of falsification that there is some s' with $s \xrightarrow{i} s'$ and $A \dashv_{s'} \phi_1$. Similarly, we get from $A \models_s [i]\phi_1$ by the definition of verification that for all s' with $s \xrightarrow{i} s'$ it holds that $A \models_{s'} \phi_1$. Contradiction with the induction hypothesis. □

THEOREM 28. *If A is a smooth action model then for each state s in the domain of A the set*

$$\Phi_s = \{\text{pre}_s\} \cup \{\phi \mid \phi \textit{ is a constraint formula and } A \models_s \phi\}$$

is consistent.

Proof. Let A be a smooth action model and s a state in W_A. Assume Φ_s is inconsistent. Then there is some $\phi \in \Phi_s$ with $\Phi_s \models \neg\phi$. It follows that $A \models_s \neg\phi$. From this, $A \dashv_s \phi$, and contradiction with Theorem 27. □

7. Filtration and Canonical Models

Our goal in Section 8 is to provide a recipe for turning any action model into an equivalent smooth action model. Next, in Section **??**, we will prove the converse of Theorem 17 for all smooth action models, i.e., we will extend Theorem 18 from purely propositional action models to all smooth action models.

For both goals we need a technique (called filtration) for constructing models from sets of formulas. The filtration technique in modal logic is used to construct a finite model for a consistent modal formula ϕ (see [8]). For ordinary modal logic the construction is based on the set of all sub-formulas of ϕ, but in PDL we have to be careful in the handling of formulas with complex modalities α, so we need so-called Fischer/Ladner closures [13].

DEFINITION 29. *Let Σ be a set of $\mathcal{L}ANG_0$ formulas. Then* FL(Σ), *the Fischer/Ladner closure of Σ, is the smallest set of formulas X that has $\Sigma \subseteq X$, that is closed under taking sub-formulas, and that satisfies the following constraints:*

– *if* $[\alpha \cup \alpha']\phi \in X$ *then* $[\alpha]\phi \in X$ *and* $[\alpha']\phi \in X$,

– *if* $[\alpha; \alpha']\phi \in X$ *then* $[\alpha][\alpha']\phi \in X$,

– *if* $[\alpha^*]\phi \in X$ *then* $[\alpha][\alpha^*]\phi \in X$.

Note that the definition handles the actual formulas of the language, not their abbreviations. As an example, consider $\Sigma = \{[(a \cup b)^*]h\}$. Then

$$FL(\Sigma) = \{[(a \cup b)^*]h, [(a \cup b)][(a \cup b)^*]h, [a][(a \cup b)^*]h, [b][(a \cup b)^*]h, h\}.$$

DEFINITION 30. (Closure under single negation). *For any formula ϕ, define $\sim\phi$, the single negation of ϕ, as follows: if ϕ has the form $\neg\psi$ then $\sim\phi = \psi$, otherwise $\sim\phi = \neg\phi$. Then $\sim\phi$ forms the negation of ϕ, while cancelling double negations. A set of formulas X is* closed under single negations *if $\phi \in X$ implies $\sim\phi \in X$.*

DEFINITION 31. (Closure of Σ). *For any formula set Σ, the closure of Σ, notation* $\neg$FL(Σ) *is the smallest set X which contains* FL(Σ) *and is closed under single negations.*

As an example, observe that the closure of $\{[(a \cup b)^*]h\}$ consists of the union of $FL(\{[(a \cup b)^*]h\})$ and the set of all negations of formulas in $FL(\{[(a \cup b)^*]h\})$.

In building epistemic models and action models from sets of formulas Σ we can take worlds (or actions) to be maximal consistent sets of formulas taken from $\neg FL(\Sigma)$.

DEFINITION 32. *Let Σ be a set of formulas. A set of formulas Γ is an* atom *over Σ if Γ is a maximal consistent subset of* $\neg$FL(Σ). *Let* At(Σ) *be the set of all atoms over Σ.*

It is easy to show for every consistent formula $\phi \in \neg FL(\Sigma)$ there is a $\Gamma \in At(\Sigma)$ with $\phi \in \Gamma$ (see [8]). For any finite formula set Γ, let $\widehat{\Gamma} = \bigwedge \Gamma$.

DEFINITION 33. *The canonical model M_Σ over finite formula set Σ is given by*

$$\begin{array}{rcl} W_\Sigma & = & \mathrm{At}(\Sigma), \\ V_\Sigma(\Gamma) & = & \{p \in \mathrm{Prop} \mid p \in \Gamma\}, \\ R_\Sigma(i) & = & \{(\Gamma, \Gamma') \mid \widehat{\Gamma} \wedge \langle i \rangle \widehat{\Gamma'} \textit{ is consistent } \}. \end{array}$$

See [8] for a proof that this canonical model 'works', in the sense that we can prove the following:

LEMMA 34. (Truth Lemma). *For all atoms $\Gamma \in \mathrm{At}(\Sigma)$ and all $\phi \in \neg\mathrm{FL}(\Sigma)$ it is the case that $M_\Sigma \models_\Gamma \phi$ iff $\phi \in \Gamma$.*

8. Construction of Smooth Action Models

Let A be an action model with set of preconditions Σ. We will show how to turn A into a smooth action models, by imposing constraints on the accessibilities, using sets of constraints from $\neg FL(\Sigma)$.

DEFINITION 35. *Let $\neg\mathrm{FLC}(\Sigma)$ be the set of all constraint formulas in $\neg\mathrm{FL}(\Sigma)$. Let Σ be a set of preconditions. For each atom Γ over Σ, call the set $\Gamma \cap \neg\mathrm{FLC}(\Sigma)$ a constrained atom over Σ.*

Note that if Γ is an atom of Σ, then $\Delta = \Gamma \cap \neg FLC(\Sigma)$ satisfies the following closure properties: if $\neg\phi \in \Delta$ then ϕ is purely propositional and $\phi \notin \Delta$, if $\phi_1 \wedge \phi_2 \in \Delta$ then $\phi_1, \phi_2 \in \Delta$, and if $\phi_1 \vee \phi_2 \in \Delta$ then $\phi_1 \in \Delta$ or $\phi_2 \in \Delta$. We will use constrained atoms to build smooth action models.

DEFINITION 36. *Let A be an action model, and let Σ be the set of preconditions of A. Assume without loss of generality that all preconditions in A are consistent. (If this is not the case, just restrict A to the action models with consistent preconditions.) Then $\mathrm{Sm}(A)$ is the action model (W, pre, R) given by*

$$\begin{array}{rcl} W & = & \{(s, \Delta) \mid \\ & & \quad s \in W_A, \Delta \textit{ a constrained atom consistent with } \mathrm{pre}_s\}, \\ \mathrm{pre}(s, \Delta) & = & \mathrm{pre}(s), \\ R(i) & = & \{((s, \Delta), (s', \Delta')) \in W^2 \mid s \xrightarrow{i} s', \widehat{\Delta} \wedge \langle i \rangle \widehat{\Delta'} \textit{ consistent}\}. \end{array}$$

If S is the set of distinctive points of A, then

$$\mathrm{Sm}(S) = \{(s, \Delta) \in W \mid s \in S\}$$

is the set of distinctive points of $\mathrm{Sm}(A)$*.*

LEMMA 37. *Let A be an action model with precondition set Σ, Let ϕ be a constraint formula in $\neg\mathrm{FL}(\Sigma)$ and let (s, Δ) be a state in $\mathrm{Sm}(A)$. Then*

$$\mathrm{Sm}(A) \models_{(s,\Delta)} \phi \text{ implies } \phi \in \Delta, \text{ and } \mathrm{Sm}(A) \;=\!\!|_{(s,\Delta)} \phi \text{ implies } \phi \notin \Delta.$$

Proof. Induction on the structure of constraint formula ϕ. Let ϕ be purely propositional, and assume $Sm(A) \models_{(s,\Delta)} \phi$. Then by Lemma 21, $pre(s, \Delta) \models \phi$. Since $pre(s, \Delta) = pre_s$ it follows by construction of Δ that $\phi \in \Delta$. Assume $Sm(A) =\!\!|_{(s,\Delta)} \phi$. Then by Lemma 21, $pre(s, \Delta) \models \neg\phi$. Since $pre(s, \Delta) = pre_s$ it follows by construction of Δ that $\phi \notin \Delta$.

Now assume the property holds for constraint formulas ϕ_1, ϕ_2. We show that it also holds for $\phi_1 \wedge \phi_2, \phi_1 \vee \phi_2, [i]\phi_1$.

Assume $Sm(A) \models_{(s,\Delta)} \phi_1 \wedge \phi_2$. Then by the definition of verification, $Sm(A) \models_{(s,\Delta)} \phi_1$ and $Sm(A) \models_{(s,\Delta)} \phi_2$. By the induction hypothesis, $\phi_1 \in \Delta$ and $\phi_2 \in \Delta$. It follows that $\phi_1 \wedge \phi_2 \in \Delta$. Assume $Sm(A) =\!\!|_{(s,\Delta)} \phi_1 \wedge \phi_2$. Then by the definition of falsification, $Sm(A) =\!\!|_{(s,\Delta)} \phi_1$ or $Sm(A) =\!\!|_{(s,\Delta)} \phi_2$. By the induction hypothesis, $\phi_1 \notin \Delta$ or $\phi_2 \notin \Delta$. It follows that $\phi_1 \wedge \phi_2 \notin \Delta$.

Assume $Sm(A) \models_{(s,\Delta)} \phi_1 \vee \phi_2$. Then by the definition of verification, $Sm(A) \models_{(s,\Delta)} \phi_1$ or $Sm(A) \models_{(s,\Delta)} \phi_2$. By the induction hypothesis, $\phi_1 \in \Delta$ or $\phi_2 \in \Delta$. It follows that $\phi_1 \vee \phi_2 \in \Delta$. Assume $Sm(A) =\!\!|_{(s,\Delta)} \phi_1 \vee \phi_2$. Then by the definition of falsification, $Sm(A) =\!\!|_{(s,\Delta)} \phi_1$ and $Sm(A) =\!\!|_{(s,\Delta)} \phi_2$. By the induction hypothesis, $\phi_1 \notin \Delta$ and $\phi_2 \notin \Delta$. It follows that $\phi_1 \vee \phi_2 \notin \Delta$.

Assume $Sm(A) \models_{(s,\Delta)} [i]\phi_1$. Then by the definition of verification, either $pre(s, \Delta) \models [i]\phi_1$ or for all (s', Δ') with $(s, \Delta) \xrightarrow{i} (s', \Delta')$ it holds that $Sm(A) \models_{(s',\Delta')} \phi_1$. In the first case, $pre_s \models [i]\phi_1$, and therefore $[i]\phi \in \Delta$. In the second case, the induction hypothesis yields $\phi_1 \in \Delta'$ for all (s', Δ') with $(s, \Delta) \xrightarrow{i} (s', \Delta')$, It follows that all Δ' with $\widehat{\Delta} \wedge \langle i \rangle \widehat{\Delta'}$ consistent satisfy $\phi_1 \in \Delta'$. Therefore, $[i]\phi_1$ is consistent with Δ. It follows by the construction of Δ that $[i]\phi_1 \in \Delta$.

Assume $Sm(A) =\!\!|_{(s,\Delta)} [i]\phi_1$. Then by the definition of falsification, either $pre(s, \Delta) \models \neg[i]\phi_1$ or there is a (s', Δ') with $(s, \Delta) \xrightarrow{i} (s', \Delta')$ and $Sm(A) =\!\!|_{(s',\Delta')} \phi_1$. In the first case, $pre_s \models \neg[i]\phi_1$, and therefore $[i]\phi \notin \Delta$, by the construction of Δ. In the second case, the induction hypothesis yields $\phi_1 \notin \Delta'$. $(s,\Delta)\xrightarrow{i}(s',\Delta')$ implies $\widehat{\Delta} \wedge \langle i\rangle\widehat{\Delta'}$ is consistent. Therefore, $\langle i\rangle{\sim}\phi$ is consistent with Δ, i.e., ${\sim}[i]\phi$ is consistent with Δ. It follows by construction of Δ that $[i]\phi \notin \Delta$. □

THEOREM 38. *For each action model A it holds that* $\mathrm{Sm}(A)$ *is smooth.*

Proof. Let ϕ be a constraint formula. We assume without loss of generality that $\phi \in \neg FL(\Sigma)$, where Σ is the set of preconditions in A. Asssume $Sm(A) \models_{(s,\Delta)} \phi$. We will prove by induction on the structure of ϕ that $Sm(A) \not\dashv_{(s,\Delta)} \phi$.

Basis. If ϕ is purely propositional then $Sm(A) \models_{(s,\Delta)} \phi$ implies $pre(s,\Delta) \models \phi$, by Lemma 21, and therefore $pre(s,\Delta) \not\models \neg\phi$, by consistency of the preconditions, and again by Lemma 21, $Sm(A) \not\dashv_{(s,\Delta)} \phi$.

Induction step. Suppose the property holds for constraint formulas ϕ_1 and ϕ_2. Let $\phi = \phi_1 \wedge \phi_2$. Then by the verification definition, it follows from

$$Sm(A) \models_{(s,\Delta)} \phi_1 \wedge \phi_2$$

that

$$Sm(A) \models_{(s,\Delta)} \phi_1 \text{ and } Sm(A) \models_{(s,\Delta)} \phi_2.$$

By the induction hypothesis, $Sm(A) \not\dashv_{(s,\Delta)} \phi_1$ and $Sm(A) \not\dashv_{(s,\Delta)} \phi_2$. By the definition of falsification, $Sm(A) \not\dashv_{(s,\Delta)} \phi_1 \wedge \phi_2$.

The case of $\phi \vee \phi_2$ is similar.

Let $\phi = [i]\phi_1$, Then by the verification definition, it follows from $Sm(A) \models_{(s,\Delta)} [i]\phi_1$ that either $pre(s,\Delta) \models [i]\phi_1$ or for all (s',Δ') with $(s,\Delta)\xrightarrow{i}(s',\Delta')$ it holds that $Sm(A) \models_{(s,\Delta)} \phi_1$.

In the first case, $pre(s,\Delta) \models [i]\phi_1$ implies $pre_s \models [i]\phi_1$. Therefore $A \models_s [i]\phi_1$, and hence $[i]\phi_1 \in \Delta$, by Lemma 37. Now suppose there is a pair (s',Δ') with $(s,\Delta)\xrightarrow{i}(s',\Delta')$. Then from the definition of the accessibilities in $Sm(A)$, it follows that $\widehat{\Delta} \wedge \langle i\rangle\widehat{\Delta'}$ is consistent. Thus, $\phi_1 \in \Delta'$. Therefore, by Lemma 37, $Sm(A) \not\dashv_{(s',\Delta')} \phi_1$. By the definition of falsification, $Sm(A) \not\dashv_{(s,\Delta)} [i]\phi_1$.

In the second case, we can assume $pre(s,\Delta) \not\models [i]\phi_1$. The induction hypothesis gives that for all (s',Δ') with $(s,\Delta)\xrightarrow{i}(s',\Delta')$ it holds that $Sm(A) \not\dashv_{(s,\Delta)} \phi_1$. By the definition of falsification and the fact that $pre(s,\Delta) \not\models [i]\phi_1$, it follows that $Sm(A) \not\dashv_{(s,\Delta)} [i]\phi_1$.

Question: can we extend this to the case of PDL formulas, where there is no notion of modal degree?

THEOREM 39. *For every distinctive action model (A,S) it holds that* $(A,S) \equiv (\mathrm{Sm}(A),\mathrm{Sm}(S))$.

Proof. Let M be an arbitrary epistemic model. The relation

$$C \subseteq W_{M \otimes A} \times W_{M \otimes Sm(A)}$$

given by

$$(w, s)C(w', (s', \Delta)) :\equiv w = w' \wedge s = s' \wedge M \models_w \widehat{\Delta}$$

is a bisimulation that connects S to $Sm(S)$. Thus,

$$(A, S) \equiv (Sm(A), Sm(S)).$$

LEMMA 40. *(A, S) is smooth iff $(A, S) \leftrightarrows (\mathrm{Sm}(A), \mathrm{Sm}(S))$.*

Proof. Let (A, S) be smooth. Let Σ be the preconditions in A. We show that the relation

$$C \subseteq W_A \times W_{Sm(A)}$$

given by

$$sC(s', \Delta) \text{ iff } s = s'$$

is a bisimulation.

Assume $sC(s, \Delta)$. Then $pre_s = pre(s, \Delta)$, so invariance holds,

For the zig condition, suppose $s \xrightarrow{i} s'$. Let Δ' be any constrained atom over Σ with

$$\{\phi \in \neg FLC(\Sigma) \mid pre_{s'} \models \phi\} \subseteq \Delta'$$

and

$$\{\delta \mid [i]\delta \in \Delta\} \subseteq \Delta'.$$

Such a constrained atom Δ' exists, for suppose the set

$$\{\phi \in \neg FLC(\Sigma) \mid pre_{s'} \models \phi\} \cup \{\delta \mid [i]\delta \in \Delta\}$$

is inconsistent. Since by consistency of $pre_{s'}$ the set

$$X = \{\phi \in \neg FLC(\Sigma) \mid pre_{s'} \models \phi\}$$

is consistent, it follows that there is some $[i]\delta \in \Delta$ such that δ is inconsistent with X, i.e., such that $\neg\delta$ follows from X. This means, by Lemma 20, that $A \models_{s'} \neg\delta$, and therefore $A \;\dashv_{s'} \delta$. By the smoothness of A, it follows that $pre_{s'} \not\models \delta$, and contradiction with $\delta \in X$.

Since $s \xrightarrow{i} s'$ and $\widehat{\Delta} \wedge \langle i \rangle \widehat{\Delta'}$ is consistent by the construction of Δ', we have that $(s, \Delta) \xrightarrow{i} (s', \Delta')$. Obviously, we also have $s' C(s', \Delta')$.

For the zag direction, assume $(s, \Delta) \xrightarrow{i} (s', \Delta')$. This implies $s \xrightarrow{i} s'$. Obviously, also $s' C(s', \Delta')$.

THEOREM 41. *Let (A,S) and (B,T) be smooth action models. Then*

$$(A,S) \equiv (B,T) \text{ implies } (A,S) \leftrightarrows (B,T).$$

Proof. Let Σ be the set of preconditions occurring in A, and Π the set of preconditions occurring in B.

By the fact that A and B are smooth, we can use Lemma 40. From a bisimulation between $(M,W)\otimes(A,S)$ and $(M,W)\otimes(B,T)$ this lemma gives us a bisimulation between $(M,W)\otimes(\mathit{Sm}(A),\mathit{Sm}(S))$ and $(M,W)\otimes(\mathit{Sm}(B),\mathit{Sm}(T))$. So first replace A by $\mathit{Sm}(A)$ and B by $\mathit{Sm}(B)$.

Let $M = M_{\Sigma\cup\Pi}$ be the canonical model built from these preconditions, and let $W = W_{\Sigma\cup\Pi}$. From $(A,S)\equiv(B,T)$ it follows that

$$(M,W)\otimes(A,S) \quad \underline{\leftrightarrow} \quad (M,W)\otimes(B,T).$$

From this, with Lemma 40:

$$(M,W)\otimes(\mathit{Sm}(A),\mathit{Sm}(S)) \quad \underline{\leftrightarrow} \quad (M,W)\otimes(\mathit{Sm}(B),\mathit{Sm}(T)).$$

Now define a binary relation $E \subseteq W_{\mathit{Sm}(A)} \times W_{\mathit{Sm}(B)}$ by means of

$$\begin{array}{rcl}(s,\Gamma^A)E(t,\Gamma^B) & :\equiv & \text{there is a } \Gamma\in W_M \text{ such that} \\ & & \Gamma^A \text{ is the restriction of } \Gamma \text{ to } \neg FLC(\Sigma), \\ & & \Gamma^B \text{ is the restriction of } \Gamma \text{ to } \neg FLC(\Pi), \\ & & \text{and } (\Gamma,(s,\Gamma^A)) \underline{\leftrightarrow} (\Gamma,(t,\Gamma^B)).\end{array}$$

We show that E is an action emulation. Suppose $(s,\Gamma^A)E(t,\Gamma^B)$. Then

Invariance By the definition of E, it follows from $(s,\Gamma^A)E(t,\Gamma^B)$ that for some $\Gamma\in W_M$: $(\Gamma,(s,\Gamma^A)) \underline{\leftrightarrow} (\Gamma,(t,\Gamma^B))$. From the fact that $(\Gamma,(s,\Gamma^A))$ is in the update, $M \models_\Gamma \mathit{pre}_s$, whence by the truth lemma, $\mathit{pre}_s\in\Gamma$. By the same reasoning we get that $\mathit{pre}_t\in\Gamma$. Since Γ is consistent, it follows that $\mathit{pre}_s \wedge \mathit{pre}_t$ is consistent.

Zig Suppose $(s,\Gamma^A)\xrightarrow{i}(s',\Gamma'^A)$. By the definition of $\mathit{Sm}(A)$, it follows that $s\xrightarrow{i}s'$ and that $\widehat{\Gamma^A} \wedge \langle i\rangle\widehat{\Gamma'^A}$ is consistent.

This means that there is some $\Gamma'\in W_M$ with $\mathit{pre}_{s'}\in\Gamma'$ and $\Gamma\xrightarrow{i}\Gamma'$, and Γ'^A the restriction of Γ' to $\neg FLC(\Sigma)$.

Now applying $(\Gamma,(s,\Gamma^A)) \underline{\leftrightarrow} (\Gamma,(t,\Gamma^B))$, we find a non-empty set G given by:

$$\begin{array}{rl} G = & \{(\Gamma'',(t',\Gamma''^B)) \in W_{M\otimes \mathit{Sm}(B)} \mid \\ & (\Gamma,(t,\Gamma^B))\xrightarrow{i}(\Gamma'',(t',\Gamma''^B)) \\ & \text{and } (\Gamma',(s',\Gamma'^A)) \underline{\leftrightarrow} (\Gamma'',(t',\Gamma''^B))\}.\end{array}$$

Since $M \models_{\Gamma''} pre_{t'}$ for all $(\Gamma'', (t', \Gamma''^B)) \in G$, it follows that $pre_{s'} \models pre_{t_1} \vee \cdots \vee pre_{t_n}$.

Zag Same reasoning vice versa.

Check that E connects every state pair $((s, \Gamma^A), (t, \Gamma^B)) \in Sm(S) \times Sm(T)$ for which $pre_s \wedge pre_t$ is consistent. $\Box$

9. Conclusion and Further Issues

***

Acknowledgements

Thanks to Johan van Benthem, Hans van Ditmarsch and Albert Visser for inspiring discussions, and thanks to two anonymous reviewers for helpful comments. We are grateful to our editor Wiebe van der Hoek for encouragement, patience. trust and flexibility. The first author is grateful to the Netherlands Institute for Advanced Studies in Wassenaar (NIAS) for providing the opportunity to complete this paper as Fellow-in-Residence.

References

1. Baltag, A. A logic for suspicious players: epistemic action and belief-updates in games. *Bulletin of Economic Research 54*, 1 (2002), 1–45.
2. Baltag, A., and Moss, L. Logics for epistemic programs. *Synthese 139*, 2 (2004), 165–224.
3. Baltag, A., Moss, L., and Solecki, S. The logic of public announcements, common knowledge, and private suspicions. In *Proceedings of TARK'98* (1998), I. Bilboa, Ed., pp. 43–56.
4. Belnap, N. A useful four-valued logic. In *Modern Uses of Multiple-valued Logic* (1977), J. Dunn and G. Epstein, Eds., Reidel, Dordrecht, pp. 8–37.
5. Benthem, J. v. Language, logic, and communication. In *Logic in Action*, J. van Benthem, P. Dekker, J. van Eijck, M. de Rijke, and Y. Venema, Eds. ILLC, 2001, pp. 7–25.
6. Benthem, J. v. One is a lonely number: on the logic of communication. Tech. Rep. PP-2002-27, ILLC, Amsterdam, 2002.
7. Benthem, J. v., van Eijck, J., and Kooi, B. Logics of communication and change. *Information and Computation 204*, 11 (2006), 1620–1662.
8. Blackburn, P., de Rijke, M., and Venema, Y. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.

9. Ditmarsch, H. v. *Knowledge Games*. PhD thesis, ILLC, Amsterdam, 2000.
10. Ditmarsch, H. v., van der Hoek, W., and Kooi, B. *Dynamic Epistemic Logic*, vol. 337 of *Synthese Library*. Springer, 2006.
11. Eijck, J. v. DEMO — a demo of epistemic modelling. In *Interactive Logic — Proceedings of the 7th Augustus de Morgan Workshop*, J. v. Benthem, D. Gabbay, and B. Löwe, Eds., no. 1 in Texts in Logic and Games. Amsterdam University Press, 2007, pp. 305–363.
12. Fagin, R., Halpern, J., Moses, Y., and Vardi, M. *Reasoning about Knowledge*. MIT Press, 1995.
13. Fischer, M., and Ladner, R. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences 18*, 2 (1979), 194–211.
14. Gerbrandy, J. *Bisimulations on planet Kripke*. PhD thesis, ILLC, 1999.
15. Kleene, S. On notation for ordinal numbers. *Journal of Symbolic Logic 3* (1938), 150–155.
16. Kozen, D., and Parikh, R. An elementary proof of the completeness of PDL. *Theoretical Computer Science 14* (1981), 113–118.
17. Segerberg, K. A completeness theorem in the modal logic of programs. In *Universal Algebra and Applications*, T. Traczyck, Ed. Polish Science Publications, 1982, pp. 36–46.