Contents

1 Introduction

Traditional puzzles about knowledge often deal with *changes of knowledge* that come about in various ways. Perhaps the most popular examples are the puzzles revolving around the fact that a declaration of ignorance of some sentence S may well lead to knowledge of S. We have in mind the scenarios that go by names such as the Muddy Children, the Cheating Spouses, the Three Wisemen, and the like. The standard treatment of these matters (a) introduces the Kripke semantics of modal logic so as to formalize the informal notions of *knowledge* and *common knowledge*; (b) formalizes one of the scenarios as a particular model; (c) and finally shows how the formalized notions of knowledge and common knowledge illuminate some key aspects of the overall scenario. The informal notion of knowledge which is closest to what is captured in the traditional semantics is probably *justified belief.* And the successful treatment of a host of puzzles leads naturally to the following

Thesis I Let s be a social situation involving the intuitive concepts of knowledge (as justified belief) and common knowledge. Assume that s is presented in such a way that the all of the relevant features of s pertaining to knowledge and common knowledge are completely determined. Then we may associate to s a mathematical model S (S is a multi-agent Kripke model; we call these *state models.*) The association is correct in the sense that all intuitive judgements concerning s correspond to formal assertions concerning S, and vice-versa.

We are not aware of any previous formulations of this thesis. Nevertheless, some version of this thesis is probably responsible for the appeal of epistemic logic.

We shall not be concerned in this paper with a defense of this thesis, but instead we return to our opening point related to *change.* Dynamic epistemic logic, dynamic doxastic logic, and related formalisms, attempt to incorporate *change from model to model* in the syntax and semantics of a logical language. We are especially concerned with changes that result from information-updating actions of various sorts. Our overall aim is to formally represent *epistemic actions* by what we shall call *updates.* By an "epistemic action" we mean a well-defined, systematic change in the information states of agents in a distributed system. By "updates" we shall mean operations defined on the space of all state models.

Perhaps the paradigm case of an epistemic action is a public announcement. The first goal is to say in a general way what the effect of a (public) announcement should be on a model. It is natural to model announcement by the logical notion of *relativization*: announcing a sentence causes all agents to restrict attention to the worlds where the sentence is true (before the announcement). Note that the informal notion of announcement takes situations to situations, and the formal notion of relativization is an operation taking models to models.

In this paper, we wish to consider many types of epistemic actions that are more difficult to handle than public announcements. These include *half-transparent, half-opaque* types of actions, such as announcements to groups in a completely private way, announcements that include the possibility that outsiders suspect the announcement but this suspicion is lost on the insiders, etc. We may also consider types of actions exhibiting information-loss and misinformation, where agents are deceived by others or by themselves.

Thesis II Let o be an operation taking situations s to situations $o(s)$. Assume that o is presented by assertions concerning knowledge and common knowledge facts about s and $o(s)$, and that o is completely determined by these assertions. Then we may associate to o a mathematical model O (an update) in such a way that for all s: if s corresponds to S as in Thesis I, then again $o(s)$ corresponds to $O(S)$ in the same way: all intuitive judgements concerning $o(s)$ correspond to formal assertions concerning $O(S)$, and vice-versa.

Our aim in this paper is not to defend this thesis. Instead, we are concerned with the more modest goal of defining and studying specific models O which we feel "get the facts right" about specific types of actions such as those we mentioned above. We also develop a more general theory, of course.

Our method of defining updates is quite general and leads to *logical languages* extending standard systems of epistemic logic by adding *updates* as new operators. These logical languages also incorporate features of propositional dynamic logic. Special cases of our logic, dealing only with public or semi-public *announcements to mutually isolated groups*, have been considered in Plaza [18], Gerbrandy [7, 8], and Gerbrandy and Groeneveld [9]. But our overall setting is much more liberal, since it allows for all the above-mentioned types of actions. We feel it would be interesting to study further examples with an eye towards applications, but we leave this to other papers.

In our logical systems, we capture only the *epistemic aspect* of these real actions, disregarding other (intentional) aspects. In particular, to keep things simple we only deal with "'purely epistemic" actions; i.e., the ones that do not change the facts of the world, but affect only the agents' beliefs about the world. However, this is *not* an essential limitation, as our formal setting can be easily adapted to express fact-changing actions (see Section 7.6).

On the semantic side, the main original technical contribution of our paper lies in our decision to represent not only the epistemic states, but also (for the first time) the epistemic actions. While for states, these structures represent in the usual way the uncertainty of each agent concerning the current state of the system, we similarly use *action signatures* to represent the uncertainty of each agent concerning the current action taking place. The intuition is that we are dealing with potentially "half-opaque/half-transparent" actions, about which the agents may be incompletely informed, or even completely misinformed. Besides the structure, actions have *preconditions*, defining their domain of applicability: not every action is possible in every state. We model the update of a state by an action as a partial *update* operation, given by a restricted product of the two structures: the uncertainties present in the given state and the given action are multiplied, while the "impossible" combinations of states and actions are eliminated (by testing the actions' preconditions on the state). The underlying intuition is that the agent's uncertainties concerning the state and the ones concerning the action are mutually independent, except that we insist on an applying action only to those states which satisfy its precondition.

On the syntactic side, we use a mixture of dynamic and epistemic logic, with dynamic modalities associated to each action signature, and with common-knowledge modalities for various groups of agents (in addition to the usual individual-knowledge operators). We give a complete and decidable axiomatization for this logic, and we prove various expressivity results. From a proof-theoretical point of view, the main originality of our system is the presence of our Action Rule, an inference rule capturing what might be called a notion of "epistemic (co)recursion". We understand this rule and our Knowledge-Action Axiom (a generalization of Ramsey's axiom to half-opaque actions) as expressing fundamental formal features of the interaction between action and knowledge in multi-agent systems, features that we think have not been formally expressed before.

For impatient readers The main logical systems of the paper are presented in Section 4.2, and to read that one would only have to read the definition in Section 4.1 first. To understand the semantics, one should read in addition Sections 2.1, 2.3, and 3.2–3.4. But we know that our systems and technical results would not be of any interest if the motivation were not so great. For this reason, we have included many examples and discussions, particularly in the sections of the paper preceding the introduction of the logics. Readers may read as much or as little of that material as they desire. Indeed, some readers may find our examples and discussion of more interest than the logical systems themselves. But from Section 4 onward, most of the work in the paper is directed towards the logical systems themselves. The main logical systems are presented in Section 5, with further basic results in Section 6. Section 7 is devoted to completeness results; and this work depends on all previous sections of the paper. Section 8 studies expressive power results for the logical systems. Much of it may be read as soon as the reader understands the semantics, but Section 8.1 requires most of the paper prior to it.

Acknowledgements

We began work on this topic in 1998, and this paper was completed in 2003. So the work has gone through numerous editions (as have we). We have found the presentation of the work to be thorny: there are many details that must be done correctly, there are many options for the presentation. We fear that the mathematically easiest way to present the work would not be understandable to most readers. Since we really want to have results that are not only true but usable, we went "back to the drawing board" several times.

We are therefore grateful to many people for discussions at various points, and we express our gratitude knowing that most of them have not even seen the end product. We especially thank Jelle Gerbrandy, Rohit Parikh, and Johan van Benthem for extremely useful conversations. We also thank Balder ten Cate, Hans van Ditmarsch, Melvin Fitting, Joseph Miller, Yde Venema, and an anonymous referee who feels the work is "somewhat in bad taste" for their comments, questions, and advice. We have learned much from presenting this work in seminars, and we thank the Indiana University Logic Group, the CUNY Applied Logic Seminar, and students in a course at the 2002 North American Summer School in Logic, Language, and Information.

L. Moss thanks the Johns Hopkins University Cognitive Science Department and the University of Maryland Institute for Advanced Computer Studies during part of the the preparation of this paper.

1.1 Scenarios

Probably the best way to enter our our overall subject is to consider some "epistemic scenarios." These give the reader some idea of what the general subject is about, and they also provide us with test problems at different points.

SCENARIO 1: THE CONCEALED COIN. A and B enter a large room containing a remote-control mechanical coin flipper. One presses a button, and the coin spins through the air, landing in a small box on a table. The box closes. The two people are much too far to see the coin.

The main contents of any representation of the relevant knowledge states of A and B are that (a) there are two alternatives, Heads and Tails; (b) neither party knows which alternative holds; and (c) that (a) and (b) are common knowledge. The need for the notion of common knowledge in reasoning about multi-agent interaction is by now standard in applied epistemic logic, and so we take it as unproblematic that one would want (c) to come out in any representation of this scenario.

Perhaps the clearest way to represent this scenario is with a diagram:

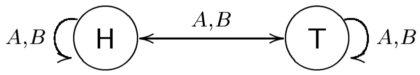

In more standard terms, we have a *set of two alternatives*, call them x and y. We also have some *atomic information*, that x represents the possible fact of the coin is lying Heads up and y that y represents the other possible fact. Finally, we also have some extra apparatus needed that indicates that A and B think both alternatives are possible. Following the standard practice in epistemic logic, we take this apparatus to be *accessibility relations* between the alternatives. The diagram should be read as saying that were the actual state of the world to be x, say, then A and B would still entertain *both* alternatives.

Notice that we did not justify our representation above. This is not so important for this scenario, but as the stories become more involved the need for justification becomes more pressing. We shall indicate much later, in Section 2.7, how one could justify these models. Readers familiar with *bisimulation* and the standard presentations of *epistemic logic* (including the notion of *common knowledge*, can look ahead at that section while reading this one.

SCENARIO 2: THE COIN REVEALED TO SHOW HEADS. A and B sit down. One opens the box and puts the coin on the table for both to see. It's H. The result of this scenario is a model which again is easy to grasp. It consists of one state; call it s. Each agent *knows* the state in the sense that they think s is the only possible state.

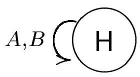

We also shall be interested in keeping track of the relation of the model in Scenario 1 and the model just above. We indicate this in the following way:

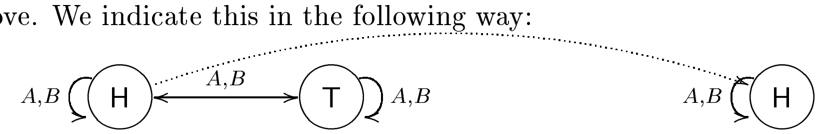

The first thing to note is that the dotted connection is a *partial function*; as we shall see later, this is the hallmark of a *deterministic* epistemic action. But we also will see quite soon that the before-after relation is not always a partial function; so there are epistemic actions which are not deterministic. Another thing to note at this point is that the dotted connection is *not* subscripted with an agent or a set of agents. It does not represent alternative possibility with respect to anyone but instead stands for the before-after relation between two models. We claim that this relation can and should be represented in way which does not vary with the agents. Obviously we shall justify this claim later. But in this section we are previewing our technical apparatus, and it would take us too far afield to delve into the justification.

Getting back to the dotted connection, it is a partial function whose domain is the set of states which could possibly be subject to the action of revealing H. This is only possible in one of the two starting states.

SCENARIO 2.1: THE COIN REVEALED TO SHOW TAILS. As a variation on Scenario 2, there is a different Scenario in which the coin is revealed in the same way to both A and B but with the change that T shows. Our full representation is:

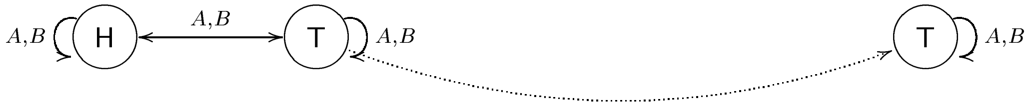

SCENARIO 2.2: THE COIN REVEALED. Finally, we can consider the *non-deterministic sum* of publicly revealing Heads and publicly revealing Tails. The coin is revealed to both A and B, but all that we as external modelers can say is that either they learned that Heads shows, or that they learned that Tails shows. Our representation is

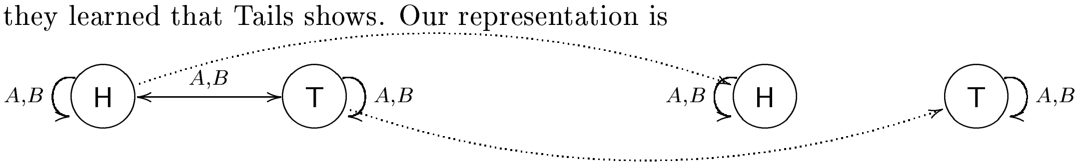

SCENARIO 3: A SEMI-PRIVATE VIEWING OF HEADS. The following is an alternative to the scenarios above in which there is a public revelation. After Scenario 1, A opens the box herself. The coin is lying Heads up. B observes A open the box but does not see the coin. And A also does not disclose whether it is Heads or Tails.

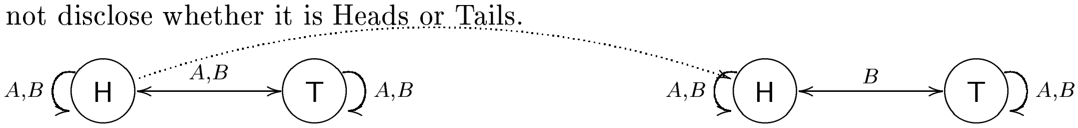

No matter which alternative holds, B would consider both as possible, and A would be certain which was the case

SCENARIO 3.1: B'S TURN. After Scenario 3, B takes a turn and opens the box the same way. We expect that after both have individually opened the boxes they see the same things; moreover, they know this will happen. This time, we begin with the end of Scenario 3, and we end with the same end as in the public revelation of H:

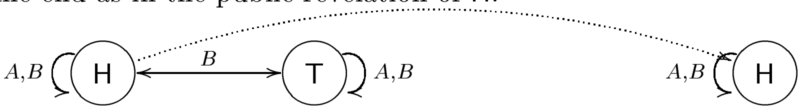

SCENARIO 4: CHEATING. After Scenario 1, A secretly opens the box herself. B does not observe A open the box, and indeed A is certain that B did not suspect that anything happened after they sat down. This is substantially more difficult conceptually, and the representation is accordingly more involved.

One might at first think that we should represent what happens before and after as in the figure below:

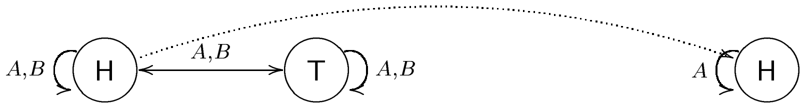

(1)

That is, cheating might be like an announcement that results in A's being sure that the coin lies H up, and B learns nothing. But the problem is in that saying that B knows nothing (new) is not appropriately represented by the absence of all arrows for B. In the state on the right of (1), B does not think it possible that the coin is lying H up. This seems wrong, since B's possibilities should be unchanged by A's cheating. To make the point another way, *deleting* arrows corresponds to *increasing* information. So what is wrong with our representation in (1) is that it corresponds to an increase in information to B, counter to our intuitions. Instead, our representation of this action is given below:

(2)

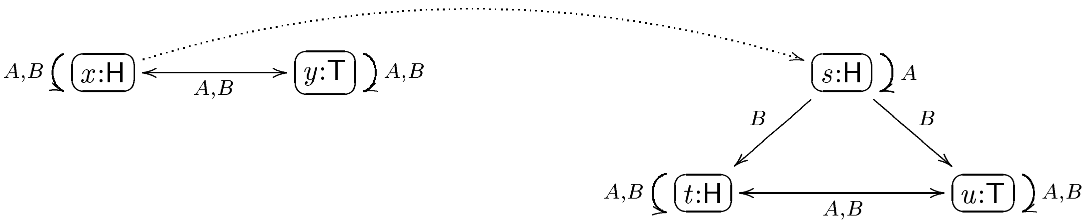

The main point is that after A cheats, B does not think that the actual state of the world is even possible. The states that B thinks are possible should only be two states which are the same as the "before" states, or states which are in some important way similar to those. In addition to the way the state actually is, where A knows the coin lies Heads up *or rather as an aspect of that state*, we need to consider other states to represent the possibilities that B must entertain.

We have indicated names s, t, and u of the possibilities in (2). (We could have done this with the scenarios we already considered, but there was less of a reason.) State s has a different status from t and u: while t and u are the states those that B thinks could hold, s is the one that A *knows* to be the case. Note that the substructure determined by t and u is isomorphic to the "before" structure. This is the way that we formalize the intuition that the states available to B after cheating are essentially the same as the states before cheating.

SCENARIO 5: MORE CHEATING. After Scenario 4, B does the very same thing. That is, B opens the box quite secretly. We leave the reader the instructive exercise of working out a representation. We shall return to this matter in Section 3.5, where we solve the problem based on our general machinery. We merely mention now that part of the goal of the paper is precisely to develop tools to build representations of complex epistemic examples such as this one.

Scenario 6: pick a card. As another alternative to Scenario 3, C walks in and tells both A and B at the same time that she has a card which either says H, T, or is blank. In the first two cases the card describes truly the state of the coin in the box, and in the last case the intention is that no information is given. Then C gives the card to A in the presence of B. Here is our representation:

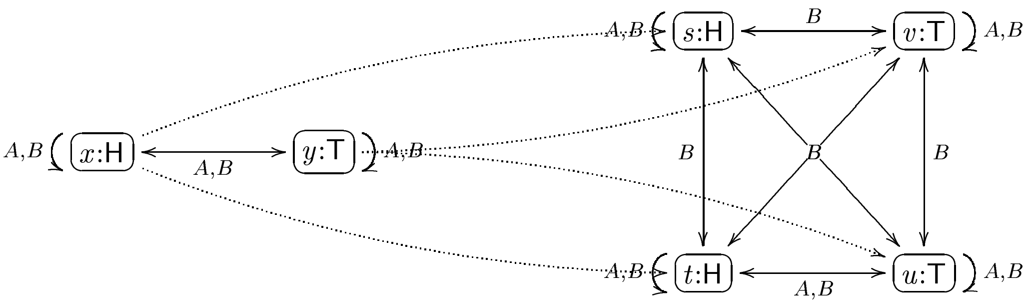

The idea here is that t and u represent states where A did not peek, and s and v states where A did. It is important that our representation not omit any B-arrows. One way to see this is to argue that the action in question leaves B with no knowledge of anything. After walking in the room and suddenly suspecting A, then all B can do is to list possible states of the world. The point is that B should then be indifferent between all the various possibilities. Another way to argue this point is to consider sentences of epistemic logic at t and u and to see whether they are in accord with our intuitions. We shall pursue this line in Section 2.7, and although we shall not consider this example there, we could also use that method to justify the representation.

SCENARIO 7: TAKE IT FROM THE TOP. As another alternative to Scenario 3, suppose that after A and B make their original entrance, A has not looked, but B suspects that she (A) possibly has cheated, and that moreover there is common knowledge of this suspicion. That is, we want to think of one single action (a knowing glance by A, for example) that results in all of this. The representation of "before" and "after" is as follows:

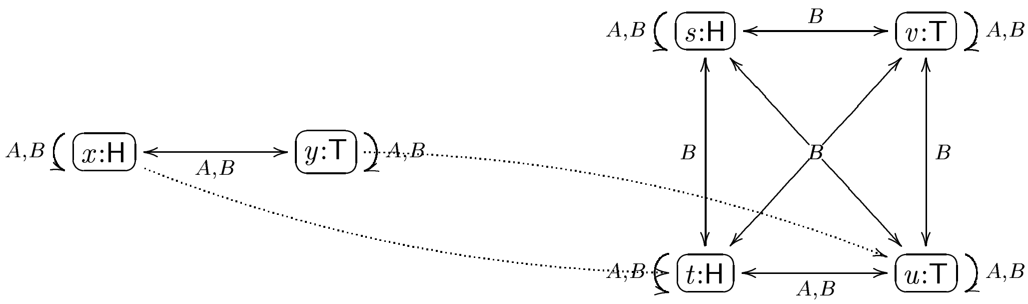

One should compare this with Scenario 6. The blank card there is a parallel to no looking here; card with H is parallel to A's looking and seeing H; and similarly for T. This accounts for the similarity in the models. The main difference is that Scenario 6 was described in such a way that we do not know what the card says; in this scenario we stipulate that A definitely did not look. This accounts for the difference in dotted lines between the two figures.

SCENARIO 7.1: PRIVATE COMMUNICATION ABOUT THE OTHER. After Scenario 7, someone stands in the doorway behind A and raises a finger. This indicates to B that A has not cheated. The communication between B and the third party was completely private. Again, we leave the representation as an exercise to the reader.

SCENARIO 8: DECEPTION. After Scenario 7, A suggests to B that she knows the state of the coin. In fact, she does not know anything whatsoever. But B believes A, A knows this, B knows that A knows it, etc. The difference between this and Scenario 7 is that here, B comes to believe that A knows rather than to merely suspect. Now our representation is

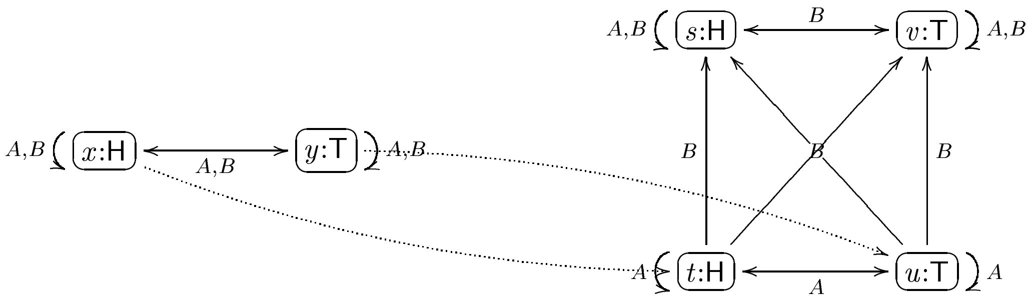

Our first promise We did not explain how we got the representations in this section. As mentioned above, one goal of the paper is to propose a theoretical understanding of these representations. It turns out that there are simple operations and general insights which allow one to construct quite complex representations. One operation allows one to pass, for example, from the representation of Scenario 1 to that of Scenario 4; moreover, the operation is so applicable that it allows us to add private knowledge for one agent in a "private" way to *any* representation.

1.2 What do we learn from public announcements?

As a further step in the introduction of our subject, we want to consider the matter of public announcements. These are the primary example of what we later call *epistemic actions.* And even though public announcements are too simple to indicate a theory of epistemic actions, this simplicity makes it easier to study them.

Public announcements are perhaps most prominent from the elimination puzzles that go under names like the "Muddy Children" the "Cheating Spouses", or the "Three Wisemen." Since the literature on these puzzles is large, the points we make in this section are not going to be new.

One point that comes up again and again in these puzzles is that publicly announcing a fact does not necessarily make the fact true. For example, telling someone that they do not know φ might well result in their knowing φ. The point is that telling them something (such as the fact that they don't know φ) in general leads to more information. And this extra information might imply that they (then) *do* know φ.

Perhaps the easiest concrete example of this is the two-person Muddy Children scenario. We have two children, A and B. A can see B but not herself, and B can see A and not himself. Let $\boldsymbol{p}$ be the proposition stating that A is dirty, and let $\boldsymbol{q}$ be the proposition stating that B is dirty. We naturally represent this by the figure on the left below.

$$
\begin{tikzcd}
{}_{A,B}\boxed{\boldsymbol{p}, \boldsymbol{q}} \arrow[r, leftrightarrow, "A"] \arrow[d, leftrightarrow, "B"'] & \boxed{\neg\boldsymbol{p}, \boldsymbol{q}}_{A,B} \arrow[d, leftrightarrow, "B"] \\
{}_{A,B}\boxed{\boldsymbol{p}, \neg\boldsymbol{q}} \arrow[r, leftrightarrow, "A"'] & \boxed{\neg\boldsymbol{p}, \neg\boldsymbol{q}}_{A,B}
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
{}_{A,B}\boxed{\boldsymbol{p}, \boldsymbol{q}} \arrow[r, leftrightarrow, "A"] \arrow[d, leftrightarrow, "B"'] & \boxed{\neg\boldsymbol{p}, \boldsymbol{q}}_{A,B} \\
{}_{A,B}\boxed{\boldsymbol{p}, \neg\boldsymbol{q}} &
\end{tikzcd}
$$

Suppose that their father comes and announces that at least one is dirty. Then we update the model as shown in the right. We also have indicated the connections between the states. It is intuitively clear that as a result of the announcement, both A and B do know that at least one of them is dirty. So this is not an example of the kind we are after. But suppose that there is a second announcement to the effect that neither knows that they are dirty.

Now let us suppose that both A and B are dirty to begin with. Then as a result of the second announcement (the one that said that neither knows they are dirty), both A and B would learn that they indeed *do* know that they are dirty. We would represent this by another one-point model

$$
{}_{A,B}\boxed{\boldsymbol{p}, \boldsymbol{q}}
$$

Returning to our issue, if one cannot conclude that a proposition is necessarily true after being announced, what can one conclude? We give an answer to this below, and then illustrate.

> The following are equivalent for all states in which φ holds:
>
> 1. After a successful public announcement of φ, A knows ψ.
>
> 2. In all the states that A considers possible in which φ holds, ψ would hold following a public announcement of φ.

Figure 1: The Announcement-Knowledge Principle

Here is an illustration of the principle at work. We go back to the original model, the one with four states drawn on the left above. Let ψ be a proposition saying that B doesn't know he is dirty. Let φ say that at least one child is dirty. Then our principle tells us that the following are equivalent, for all states in which there is at least one dirty child.

1. Following a public announcement that at least one child is dirty, A knows that B doesn't know he is dirty.

2. In all the states that A considers possible in which at least one child is dirty, after announcement that at least one was dirty, B would not know he is dirty.

The equivalence here implies that in each state in the original model, either both (1) and (2) hold, or they both fail. We examine each state in turn and check (1) and (2) on intuitive grounds.[1]

[1]That is, we read "know" informally. We'll examine this kind of example again, where we read "know" as it

In the state $\boxed{p, q}$, (1) fails. A would reason that perhaps she is clean. In that case, B would see this and realize that *he* must be dirty. So B would know that he is dirty. For the same reason, (2) fails in $\boxed{p, q}$. That is, A considers $\boxed{\neg p, q}$ possible, and in that state in the original model, it would not be the case that after the announcement that at least one child is dirty, B would not know he is dirty.

In contrast, in the state $\boxed{p, \neg q}$, (1) holds. Intuitively, A sees that B is clean and thus knows she is dirty and that he is clean. That is, she knows which state really obtains. In that state, B doesn't know: B sees A is dirty but still thinks that it is possible that he, too, is dirty. This is the reason why (1) holds, and (2) is similar.

The reasoning for the state $\boxed{\neg p, q}$ is similar. We don't have to worry about $\boxed{\neg p, \neg q}$ since it is not true that there is at least one dirty child in that state.

This concludes an example that illustrates the Announcement-Knowledge Principle. It is important to understand this principle, and we encourage the reader to think about it carefully. For example, it is important that (2) cannot be replaced by either of the following variants:

$2'$. In all the states that A considers possible in which φ holds, ψ holds.

$2''$. In all the states that A considers possible, ψ would hold following a public announcement of φ.

(This would be a good exercise!)

Our second promise Another goal of this paper has to do with this Announcement-Knowledge Principle. We want to see what the appropriate versions of it are for other types of epistemic actions. We also want to construct logical systems in which a formalized version of this principle will be a central new axiom.

1.3 Further contents of this paper

Section 2 continues our introduction by both reviewing some of the main background concepts of epistemic logic, and also by situating the logical systems which we'll construct in terms of well-known systems. Section 2.1 presents a new "big picture" of the world of epistemic actions. While the work that we do could be understood without the conceptual part of the big picture, it probably would help the reader to work through our definitions.

Section 3 begins the technical part of the paper in earnest, and here we revisit some of the Scenarios of Section 1.1 and their pictures. The idea is to get a logical system for reasoning with, and about, these kinds of models. Section 4 gives logical systems, and these are studied further in Sections 5, 6, and 7. Section 8 is a study of some of the many expressive power issues that come from our work.

Endnotes We gather at the ends of our sections some remarks about the literature and how our work fits in with that of others.

is formalized in the Kripke semantics. Then, the Announcement-Knowledge Principle becomes a consequence of the semantics. For now, we want to make the case that it holds intuitively.

We mentioned other work in dynamic epistemic logic and dynamic doxastic logic. Much more on these logics and many other proposals in epistemic logic may be found in Gochet and Gribomont [10], and Meyer and van der Hoek [13].

The ideas behind several of the scenarios in Section 1.1 are to be found in several places: see, e.g., Plaza [18], Gerbrandy [7, 8], Gerbrandy and Groeneveld [9], and van Ditmarsch [22]. We shall discuss these papers in more detail later. Our scenarios go beyond the work of these papers. Specifically, our treatment of the actions in Scenarios 7 and 8 seems new. Also, our use of the relation between "before" and "after" (given by the dotted arrows) is new. The Announcement-Knowledge Principle is originally due to Plaza and Gerbrandy (independently).

2 Preview and review

We fix a set AtSen of *atomic sentences*, and a set $\mathcal{A}$ of *agents*. All of our definitions depend on AtSen and $\mathcal{A}$, but for the most part we omit mention of these.

2.1 State models and epistemic propositions

A *state model* is a triple $\mathbf{S} = (S, \xrightarrow{A}_{\mathbf{S}}, \|\cdot\|_{\mathbf{S}})$ consisting of a set S of "states"; a family of binary accessibility relations $\xrightarrow{A}_{\mathbf{S}} \subseteq S \times S$, one for each agent $A \in \mathcal{A}$; and a "valuation" (or a "truth" map) $\|.\|_{\mathbf{S}} : \mathsf{AtSen} \to \mathcal{P}(S)$, assigning to each atomic sentence p a set $\|p\|_{\mathbf{S}}$ of states. When dealing with a single fixed state model $\mathbf{S}$, we often drop the subscript $\mathbf{S}$ from all the notation.

In a state model, the "facts" are thought of as properties of states; the valuation tells us which facts hold at which states. The accessibility relations model the agents' epistemic uncertainty about the current state. That is, to say that $s \xrightarrow{A} t$ in $\mathbf{S}$ means that in the model, in state s, agent A considers it possible that the state is t.

Definition Let SMod be the collection of all state models. An *epistemic proposition* is an operation φ defined on SMod such that for all $\mathbf{S} \in \mathsf{SMod}$, $\varphi_{\mathbf{S}} \subseteq S$.

The collection of epistemic propositions is closed in various ways.

1. For each atomic sentence p we have an atomic proposition $\boldsymbol{p}$ with $\boldsymbol{p}_{\mathbf{S}} = \|p\|_{\mathbf{S}}$.[2]

2. If φ is an epistemic proposition, then so is $\neg\varphi$, where $(\neg\varphi)_{\mathbf{S}} = S \setminus \varphi_{\mathbf{S}}$.

3. If C is a set or class of epistemic propositions, then so is $\bigwedge C$, with
$$(\bigwedge C)_{\mathbf{S}} \quad = \quad \bigcap\{\varphi_{\mathbf{S}} : \varphi \in C\}.$$

4. Taking C above to be empty, we have an epistemic proposition **True**, with $\mathbf{True}_{\mathbf{S}} = S$.

5. We also may take C in part (3) to be a two-element set $\{\varphi, \psi\}$; here we write $\varphi \wedge \psi$ instead of $\bigwedge\{\varphi, \psi\}$. We see that if φ and ψ are epistemic propositions, then so is $\varphi \wedge \psi$, with $(\varphi \wedge \psi)_{\mathbf{S}} = \varphi_{\mathbf{S}} \cap \psi_{\mathbf{S}}$.

6. If φ is an epistemic proposition and $A \in \mathcal{A}$, then $\Box_A\varphi$ is an epistemic proposition, with
$$(\Box_A\varphi)_{\mathbf{S}} \quad = \quad \{s \in S : \text{if } s \xrightarrow{A} t, \text{ then } t \in \varphi_{\mathbf{S}}\}. \tag{3}$$

7. If φ is an epistemic proposition and $\mathcal{B} \subseteq \mathcal{A}$, then $\Box^*_{\mathcal{B}}\varphi$ is an epistemic proposition, with
$$(\Box^*_{\mathcal{B}}\varphi)_{\mathbf{S}} \quad = \quad \{s \in S : \text{if } s \xrightarrow{\mathcal{B}^*} t, \text{ then } t \in \varphi_{\mathbf{S}}\}.$$

   Here $s \xrightarrow{\mathcal{B}^*} t$ iff there is a sequence
$$s = u_0 \quad \to_{A_1} \quad u_1 \quad \to_{A_2} \quad \cdots \quad \cdots \to_{A_n} \quad u_{n+1} = t$$

[2] It is important to see that the *sentence* p is a syntactic object, while the *proposition* $\boldsymbol{p}$ is a semantic object, a function defined on the class of state models. We define $\boldsymbol{p}$ on a state model $\mathbf{S}$ by using the the function $\|.\| : \mathsf{AtSen} \to S$ that is part of the structure of $\mathbf{S}$.

where $A_1, \ldots, A_n \in \mathcal{B}$. In other words, there is a sequence of arrows from the set $\mathcal{B}$ taking s to t. We allow $n = 0$ here, so $\xrightarrow{\mathcal{B}^*}$ includes the identity relation on S. To see that $\Box^*_{\mathcal{B}}\varphi$ is indeed an epistemic proposition, we use parts 3 and 6 above; we may also argue directly, of course.

2.2 Epistemic logic recalled

In this section, we review the basic definitions of modal logic. In terms of our work in Section 2.1 above, the *infinitary modal propositions* are the smallest collection of epistemic propositions containing the propositions $\boldsymbol{p}$ corresponding to atomic sentences p and closed under negation $\neg$, infinitary conjunction $\bigwedge$, and the agent modalities $\Box_A$. The *finitary* propositions are the smallest collection closed the same way, except that we replace $\bigwedge$ by its special cases **True** and the binary conjunction operation.

Recall that we insist on a distinction of syntactic and semantic objects in this paper. Up until now, the only syntactic objects have been the atomic sentences $p \in \mathsf{AtSen}$. But we can build the collections of *finitary* and *infinitary atomic sentences* by the same definitions that we have seen, and then the work of the past section is the *semantics* of our logical systems. For example, we have sentences $p \wedge q$, $\Box_A \neg p$, and $\Box^*_{\mathcal{A}} q$. These then have corresponding epistemic propositions as their semantics: $\boldsymbol{p} \wedge \boldsymbol{q}$, $\Box_A \neg \boldsymbol{p}$, and $\Box^*_{\mathcal{A}} \boldsymbol{q}$, respectively. Note that the latter is a properly infinitary proposition (and so $\Box^*_{\mathcal{A}} q$ is a properly infinitary sentence); it abbreviates an infinite conjunction.

The rest of this section studies examples of the semantics and it also makes the connection of the formal system with the informal notions of *knowledge* and *common knowledge*. We shall study Scenario 3 of Section 1.1, where A opens the box herself to see Heads in a semi-public way: B sees A open the box but not the result, A is aware of this, etc. We want to study the model after the opening. We represented this as

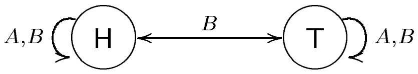

We first must represent this picture as a bona fide state model $\mathbf{S}_3$ in our sense[3]. The picture includes no explicit states, but we must fix some states to have a representation. We choose distinct objects s and t. Then we take as our state model $\mathbf{S}_3$, as defined by

$$
\begin{array}{lcl@{\qquad\qquad}lcl}
S_3 & = & \{s, t\} & & & \\
\xrightarrow{A} & = & \{(s,s), (t,t)\} & \|\mathsf{H}\| & = & \{s\} \\
\xrightarrow{B} & = & \{(s,s), (s,t), (t,s), (t,t)\} & \|\mathsf{T}\| & = & \{t\}
\end{array}
$$

In Figure 2, we list some sentences of English along with their translations into standard epistemic logic. We also have calculated the semantics of the formal sentences in the model $\mathbf{S}_3$. It should be stressed that the semantics is exactly the one defined in the previous section. For example,

$$
\begin{array}{lcl}
\|\Box_A \mathsf{T}\|_{\mathbf{S}_3} & = & \{u \in S_3 : \text{if } u \xrightarrow{A} v, \text{ then } v \in \|\mathsf{T}\|_{\mathbf{S}_3}\} \\
& = & \{u \in S_3 : \text{if } u \xrightarrow{A} v, \text{ then } v = t\} \\
& = & \{t\}
\end{array}
$$

[3]The subscript 3 comes from the number of the scenario; we shall speak of corresponding models $\mathbf{S}_1$, $\mathbf{S}_2$, etc., and each time the models will be the ones pictured in Section 1.1.

English	Formal rendering	Semantics
the coin shows H	H	$\{t\}$
A knows the coin shows H	$\Box_A \mathsf{H}$	$\{s\}$
A knows the coin shows T	$\Box_A \mathsf{T}$	$\{t\}$
B knows that the coin shows H	$\Box_B \mathsf{H}$	$\emptyset$
A knows that B doesn't know it's H	$\Box_A \neg \Box_B \mathsf{H}$	$\{s, t\}$
B knows that A knows that B doesn't know it's H	$\Box_B \Box_A \neg \Box_B \mathsf{H}$	$\{s, t\}$
it is common knowledge that either A knows it's H or A knows that it's T	$\Box^*_{A,B}(\Box_A \mathsf{H} \vee \Box_A \mathsf{T})$	$\{s, t\}$
it is common knowledge that B doesn't know the state of the coin	$\Box^*_{A,B} \neg(\Box_B \mathsf{H} \vee \Box_B \mathsf{T})$	$\{s, t\}$

Figure 2: Examples of translations and semantics

We also emphasize two other points. First, the translation of English to the formal system is based on the rendering of "A knows" as $\Box_A$, and of "it is common knowledge that" by $\Box^*_{A,B}$. Second, the chart bears a relation to intuitions that one naturally has about the states s and t. Recall that s is the state that obtains after A looks and sees that the coin is lying Heads up. The state t, on the other hand, is a state that would have been the case had A seen Tails when she looked.

2.3 Updates

An *epistemic relation* between state models $\mathbf{S}$ and $\mathbf{T}$ is a relation between the sets S and T. We write $r : \mathbf{S} \to \mathbf{T}$ for this.

An *update* $\mathbf{r}$ is a pair of operations on SMod,

$$\mathbf{r} \quad = \quad (\mathbf{S} \mapsto \mathbf{S}(\mathbf{r}), \mathbf{S} \mapsto \mathbf{r_S}),$$

where for each $\mathbf{S} \in \mathsf{SMod}$, $\mathbf{r_S} : \mathbf{S} \to \mathbf{S}(\mathbf{r})$ is an epistemic relation. We call $\mathbf{S} \mapsto \mathbf{S}(\mathbf{r})$ the *update map*, and $\mathbf{S} \mapsto \mathbf{r_S}$ the *update relation.*

Example 2.1 Let φ be an epistemic proposition. We get an update *Pub* φ which represents *the public announcement of* φ. The update map takes the submodel determined by $\varphi_{\mathbf{S}}$. That is, for each $\mathbf{S}$, $\mathbf{S}(\mathit{Pub}\ \varphi)$ is the sub-state-model of $\mathbf{S}$ determined by the states in $\varphi_{\mathbf{S}}$. In this submodel, information about atomic sentences and accessibility relations is simply inherited from the larger model. The update relation $(\mathit{Pub}\ \varphi)_{\mathbf{S}}$ is the inverse of the inclusion.

Example 2.2 We also get a different update $?\varphi$ which represents *testing whether φ is true.* Here we take $\mathbf{S}(?\varphi)$ to be the model whose state set is

$$(\{0\} \times \varphi_{\mathbf{S}}) \;\cup\; (\{1\} \times \mathbf{S}).$$

The arrows are defined by $(i, s) \xrightarrow{A} (j, t)$ iff $s \xrightarrow{A} t$ in $\mathbf{S}$ and $j = 1$. (Note that states of the form $(0, s)$ are never the targets of arrows in the new model.) Finally, we set

$$\|p\|_{\mathbf{S}(?\varphi)} \quad = \quad \{(i, s) \in \mathbf{S}(?\varphi) : s \in \|p\|_{\mathbf{S}}\}.$$

The relation $(?\varphi)_{\mathbf{S}}$ is the set of pairs $(s,(0,s))$ such that $s \in \varphi_{\mathbf{S}}$.

We shall study these two examples (and many others) in the sequel, and in particular we shall justify the names "public announcement" and "test". For now, we continue our general discussion by noting that the collection of updates is closed in various ways.

1. There is an update $\mathbf{1}$ with $\mathbf{S}(\mathbf{1}) = \mathbf{S}$, and $\mathbf{1}_{\mathbf{S}}$ is the identity relation on $\mathbf{S}$.

2. If $\mathbf{r}$ and $\mathbf{s}$ are epistemic updates, then their composition $\mathbf{r} \cdot \mathbf{s}$ is again an epistemic update, where $\mathbf{S}(\mathbf{r} \cdot \mathbf{s}) = \mathbf{S}(\mathbf{r})(\mathbf{s})$, and $(\mathbf{r} \cdot \mathbf{s})_{\mathbf{S}} = \mathbf{r}_{\mathbf{S}} \cdot \mathbf{s}_{\mathbf{S}(\mathbf{r})}$.[4]

3. If X is any set of epistemic updates, then their sum $\sum_X \mathbf{r}$ is again an epistemic update. Call it $\mathbf{s}$. We define $\mathbf{S}(\mathbf{s})$ to be the state model whose set of states is

$$\{(t, \mathbf{r}) : \mathbf{r} \in X \text{ and } t \in \mathbf{S}(\mathbf{r})\}.$$

Each accessibility relation $\xrightarrow{A}$ is defined by

$$(t, \mathbf{r}) \xrightarrow{A} (u, \mathbf{s}) \quad \text{iff} \quad \text{if } \mathbf{r} = \mathbf{s} \text{ and } t \xrightarrow{A} u \text{ in } \mathbf{S}(\mathbf{r}).$$

The interpretations of atomic propositions are inherited from the state models:

$$\|p\|_{\mathbf{S}(\mathbf{s})} \quad = \quad \{(t, \mathbf{r}) \in \mathbf{S}(\mathbf{s}) : t \in \|p\|_{\mathbf{S}(\mathbf{r})}\}.$$

Finally, the update relation between $\mathbf{S}$ and $\mathbf{S}(\mathbf{s})$ is given by

$$s \ \mathbf{s}_{\mathbf{S}} \ (t, \mathbf{r}) \quad \text{iff} \quad s \ \mathbf{r}_{\mathbf{S}} \ t.$$

4. A special case: if $\mathbf{r}$ and $\mathbf{s}$ are epistemic updates, then their sum $\mathbf{r} + \mathbf{s}$ is again an epistemic update, where $\mathbf{S}(\mathbf{r} + \mathbf{s}) = \mathbf{S}(\mathbf{r}) + \mathbf{S}(\mathbf{s})$, and $(\mathbf{r} + \mathbf{s})_{\mathbf{S}} = \mathbf{r}_{\mathbf{S}} + \mathbf{s}_{\mathbf{S}}$. (Again, we are writing $+$ for the disjoint union ("coproduct") operation on sets, extended in the obvious way to relations.)

5. Another special case: we have the operation of Kleene star on updates:

$$\mathbf{r}^* \quad = \quad \sum \{\mathbf{1}, \mathbf{r}, \mathbf{r} \cdot \mathbf{r}, \ldots, \mathbf{r}^n, \ldots\}$$

6. We can also take $X = \emptyset$ in part 3. This gives an update $\mathbf{0}$ such that $\mathbf{S}(\mathbf{0})$ is the empty model for each $\mathbf{S}$, and $\mathbf{0}_{\mathbf{S}}$ is the empty relation.

A new operation If φ is an epistemic proposition and $\mathbf{r}$ an update, then $[\mathbf{r}]\varphi$ is an epistemic proposition defined by

$$([\mathbf{r}]\varphi)_{\mathbf{S}} \quad = \quad \{s \in S : \text{if } s \ \mathbf{r}_{\mathbf{S}} \ t, \text{ then } t \in \varphi_{\mathbf{S}(\mathbf{r})}\}. \tag{4}$$

We should compare (4) and (3). The point is that we may treat updates in a similar manner to other box-like modalities; the structure given by an update allows us to do this. This point

[4] We are writing relational composition in left-to-right order in this paper.

leads to the formal languages which we shall construct in due course. But we can illustrate the idea even now. Suppose we want to interpret the sentence $[Pub\ \mathsf{H}]\Box_A\mathsf{H}$ in our very first model, the model $\mathbf{S}_1$ pictured again below:

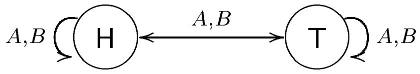

We shall call the two states s (where H holds) and t. Again, we want to determine $[\![[Pub\ \mathsf{H}]\Box_A\mathsf{H}]\!]_{\mathbf{S}_1}$. We already have in Example 2.1 a general definition of $Pub\ \mathsf{H}$ as an update, so we can calculate $\mathbf{S}_1(Pub\ \mathsf{H})$ and also $(Pub\ \mathsf{H})_{\mathbf{S}_1}$. We indicate these in the picture

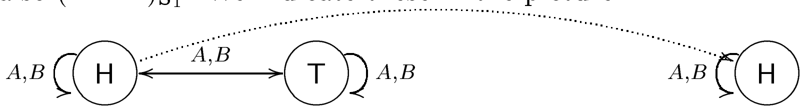

The one-state model on the right is $\mathbf{S}_1(Pub\ \mathsf{H})$, and the dotted arrow shows $(Pub\ \mathsf{H})_{\mathbf{S}_1}$. So we calculate:

$$[\![[Pub\ \mathsf{H}]\Box_A\mathsf{H}]\!]_{\mathbf{S}_1} \quad = \quad \{s \in \mathbf{S}_1 : \text{whenever } s\ (Pub\ \mathsf{H})_{\mathbf{S}_1}\ t, \text{ then also } t \in [\![\Box_A\mathsf{H}]\!]_{\mathbf{S}_1(Pub\ \mathsf{H})}\}.$$

In $\mathbf{S}_1(Pub\ \mathsf{H})$, the state satisfies $\Box_A\mathsf{H}$. Thus $[\![[Pub\ \mathsf{H}]\Box_A\mathsf{H}]\!]_{\mathbf{S}_1} = \{s,t\}$. It might be more interesting to consider $\langle Pub\ \mathsf{H}\rangle\Box_A\mathsf{H}$; recall that this is $\neg[Pub\ \mathsf{H}]\neg\Box_A\mathsf{H}$. Similar calculations show that

$$[\![\langle Pub\ \mathsf{H}\rangle\Box_A\mathsf{H}]\!]_{\mathbf{S}_1} \quad = \quad \{s \in \mathbf{S}_1 : \text{for some } t,\ s\ (Pub\ \mathsf{H})_{\mathbf{S}_1}\ t \text{ and } t \in [\![\Box_A\mathsf{H}]\!]_{\mathbf{S}_1(Pub\ \mathsf{H})}\}.$$

Hence $[\![\langle Pub\ \mathsf{H}\rangle\Box_A\mathsf{H}]\!]_{\mathbf{S}_1} = \{s\}$.

The point here is that we have a general semantics for sentences like $[Pub\ \mathsf{H}]\Box_A\mathsf{H}$, and this semantics crucially uses equation (4). That is, to determine the truth set of a sentence like $[Pub\ \mathsf{H}]\Box_A\mathsf{H}$ in a particular model $\mathbf{S}$, one *applies* the update map to $\mathbf{S}$ and works with the update relation between the $\mathbf{S}$ and the $\mathbf{S}([Pub\ \mathsf{H}])$; one also uses the semantics of $\Box_A\mathsf{H}$ in the new model. This overall point is one of the two leading features of our approach; the other is the update product which we introduce in Section 3.

2.4 The target logical systems

This paper presents a number of logical systems which contain *epistemic operators* of various types. These operators are closely related to aspects of the scenarios of Section 1.1. The logics themselves are presented formally in Section 4, but this work takes a fair amount of preparation. We delay this until after we motivate the subject, and so we turn to an informal presentation of the the syntax and semantics of some logics. The overall idea is that the operators we study correspond roughly to the *shapes* of the action models which we saw in Section 3.5.

THE LOGIC OF PUBLIC ANNOUNCEMENTS. We take a two-place sentential operator $[Pub\ \varphi]\psi$. That is, we want an operation taking sentences φ and ψ to a new sentence $[Pub\ \varphi]\psi$, and we want a logic closed under this operation. The intended interpretation of $[Pub\ \varphi]\psi$ is: assuming that φ holds, then announcing it results in a state where ψ holds. The announcement here should be completely public. The semantics of every sentence χ will be an epistemic proposition $[\![\chi]\!]$ in the sense of Section section-epistemic-prop-definition.

We mention one property of the semantics, having to do with the model $\mathbf{S}_1$ of Section 1.1. The shall arrange that the semantics to work out so that $[\![[Pub\ \mathsf{H}]\Box^*_{A,B}\mathsf{H}]\!]$ has the property that $x \in [\![[Pub\ \mathsf{H}]\Box^*_{A,B}\mathsf{H}]\!]_{\mathsf{S}_1}$. That is, x is a state where H holds, and if we announce this fact publically, it will be common knowledge that H holds. Usually, one would write this as $x \models [Pub\ \mathsf{H}]\Box^*_{A,B}\mathsf{H}$.

Note that we refer to operator as a *two-place* one. This just means that it takes two sentences and returns another sentence. We also think of $[Pub\ \varphi]$ as a modal operator in its own right, on a par with the knowledge operators $\Box_A$. And so we shall think of Pub as something which takes sentences into one-place modal operators.

We also consider a *dual* $\langle Pub\ \varphi\rangle$ to $[Pub\ \varphi]$. As one expects, the semantics will arrange that $\langle Pub\ \varphi\rangle\psi$ and $\neg[Pub\ \varphi]\neg\psi$ be logically equivalent. (That is, they will hold at the same states.) Thus, the intended interpretation of $\langle Pub\ \varphi\rangle\psi$ is: φ holds, and announcing it results in a state where ψ holds. As one indication of the difference, in $\mathbf{S}_1$, $y \models [Pub\ \mathsf{H}]\mathsf{True}$ (vacuously, as it were: it is not possible to make a true announcement of H in y). But we do have $y \models \neg\langle Pub\ \mathsf{H}\rangle\mathsf{True}$; this is how our semantics works out. Once again we only consider true announcements. Our semantics will arrange that $\varphi \to (\langle Pub\ \varphi\rangle\psi \leftrightarrow [Pub\ \varphi]\psi)$. So in working with the example scenarios, the difference between the two modalities will be small.

Further, we *iterate* the announcement operation, obtaining sentences such as $[Pub\, p][Pub\, q]r$. We also want to consider *announcements about announcements*, as in $\langle Pub\ \langle Pub\ \varphi\rangle\psi\rangle\chi$. This says that it is possible to announce publically $\langle Pub\ \varphi\rangle\psi$, and as a result of a true announcement of this sentence, χ will hold.

The logic of completely private announcements to groups. This time, the syntax is more complicated. If φ and ψ are sentences and $\mathcal{B}$ is a set of agents, then $[Pri^{\mathcal{B}}\ \varphi]\psi$ is again a sentence. The intended interpretation of this is: assuming that φ holds, then announcing it publicly to the subgroup $\mathcal{B}$ in such a way that outsiders do not even suspect that the announcement happened results in a state where ψ holds. (The "*Pri*" in the notation stands for "private.") For example, this kind of announcement occurs in the passage from Scenario 1 to Scenario 4; that is, "cheating" is a kind of private announcement to the "group" $\{A\}$. We want it to be the case that in $\mathbf{S}_1$,

$$x \models \langle Pri^{\{A\}}\ \mathsf{H}\rangle(\Box_A\mathsf{H}\ \wedge\ \neg\Box_B\Box_A\mathsf{H}).$$

That is, in x, it is possible to announce H to A privately (since H is true in x), and by so doing we have a new state where A knows this fact, but B does not know that A knows it.

The logic of private announcements to groups allows as modal operators $[Pri^{\mathcal{B}}\ \varphi]$ for all sentences φ of the logic. We shall show that this logical system extends the logic of public announcements, the idea being that when we take $\mathcal{B}$ to be the full set $\mathcal{A}$ of agents, $Pub\ \varphi$ and $Pri^{\mathcal{A}}\varphi$ should be equivalent in the appropriate sense.

The logic of common knowledge of alternatives. If φ and ψ are sentences and $\mathcal{B}$ is a set of agents, then $[Cka^{\mathcal{B}}\ \vec{\varphi}]\psi$ is again a sentence. The intended interpretation of this is: assuming that φ_1 holds, then announcing it publicly to the subgroup $\mathcal{B}$ in such a way it is common knowledge to the set of all agents that the announcement was one of $\vec{\varphi}$ results in a state where ψ holds. For example, consider Scenario scenario-A-opens. In $\mathbf{S}_1$, we have

$$x \models \neg\Box_A\mathsf{H}\ \wedge\ \langle Cka^{\{A\}}\ \mathsf{H},\mathsf{T}\rangle(\Box_A(\mathsf{H}\ \wedge\ \neg\Box_B\Box_A\mathsf{H})).$$

Syntactic	Semantic
sentence φ	epistemic proposition $\boldsymbol{\varphi}$
language $\mathcal{L}$, $\mathcal{L}(\mathbf{\Sigma})$, etc.	state model $\mathbf{S}$
action signature $\mathbf{\Sigma}$	update $\mathbf{r}$
basic action expression $\sigma\vec{\psi}$	epistemic program model $(\mathbf{\Sigma}, \Gamma)(\boldsymbol{\psi}_1, \ldots \boldsymbol{\psi}_n)$
program π, action α	canonical action model $\mathbf{\Omega}$

Figure 3: The main notions in this paper

In words, we can announce Heads to A (alone) but in such a way that B knows that the announcement is either Heads or Tails (and this fact is common knowledge): this after all is just what it means for A to open the box privately in the presence of B. As a result of this opening, A knows the contents of the box (which she did not know before), and she also knows that B does not know that she knows it.

At this point, we have the syntax of some logical languages and also examples of what we intend the semantics to be. We do not yet have the *formal* semantics of any of these languages, of course, and so we cannot present axiomatizations or questions such as completeness. However, even before we give the details on the semantics, we want to be clear that our goal in this paper is to study a very wide class of logical systems, including ones for the representation of all possible "announcement types." The reasons to be interested in this approach rather than to study a few separate logics are as follows: (1) it is more general and elegant to have a unified presentation; and (2) it gives a formal account of the notion of an "announcement type" which would be otherwise lacking and which should be of independent interest. So what we will really be concerned with is:

THE LOGIC OF ALL POSSIBLE EPISTEMIC ACTIONS. If α is an epistemic action and φ is a sentence, then $[\alpha]\varphi$ is again a sentence. Here α is some sort of epistemic action (which of course we shall define in due course); the point is that α might involve arbitrarily complex patterns of suspicion. In this way, we shall recover the four logical systems mentioned above as fragments of the larger logic of all possible epistemic actions.

Our claims in this paper Here are some of the claims of the paper about the logical languages we shall construct and about our notion of updates.

1. Each *type of epistemic action*, such as public announcements, completely private announcements, announcements with common knowledge of suspicion, etc, corresponds to a natural collection of updates as we have defined them above.

2. Each type also gives rise to a logical language with that type of action as a primitive construct. Moreover, it is possible to precisely define the syntax of the language to insure closure under the construct. That is, we should be able to formulate a language with announcements not just about atomic facts, but about announcements themselves, announcements about announcements, etc.

2.5 A guide to the concepts in this paper

In this section, we catalog the main notions that we shall introduce in due course. After this, we turn to a discussion of the particular languages that we study, and we situate them with existing logical languages.

Syntactic and semantic notions It will be important for us to make a sharp distinction between syntactic and semantic notions. For example, we speak of atomic *sentences* and atomic *propositions*. The difference for us is that atomic sentences are entirely syntactic objects: we won't treat an atomic sentence p as anything except an unanalyzed mathematical object. On the other hand, this atomic sentence p also has associated with it the atomic proposition $\boldsymbol{p}$. For us $\boldsymbol{p}$ will be a function whose domain is the (proper class of) state models, and it is defined by

$$\boldsymbol{p}_{\mathbf{S}} \quad = \quad \{s \in S : s \in \|p\|_{\mathbf{S}}\}. \tag{5}$$

This difference may seem pedantic at first, and surely there are times when it is sensible to blur it. But for various reasons that will hopefully become clear, we need to insist on it.

We list in Figure 3 the main notions that we need in this paper. We do this mostly to help readers as they explore the different notions. We mention now that the various notions are not developed in the order listed; indeed, we have tried hard to organize this paper in a way that will be easiest to read and understand. For example, one of our main goals is to present a set of languages (syntactic objects) and their semantics (utilizing semantic objects).

Languages This paper studies a number of languages, and to help the reader we list these in Figure 4. Actually, what we study are not individual languages, but rather *families* of languages parameterized by different choices of primitives. It is standard in modal logic to begin with a set of atomic propositions, and we do the same. The difference is that we shall call these atomic *sentences* in order to make a distinction between these essentially syntactic objects and the *semantic* propositions that we study beginning in Section 2.3. This is our first parameter, a set AtSen of *atomic sentences*. The second is a set $\mathcal{A}$ of *agents*.

Given these, $\mathcal{L}_0$ is ordinary modal logic with the elements of AtSen as atomic sentences and with agent-knowledge modalities $\Box_A$ for $A \in \mathcal{A}$.

We add *common-knowledge operators* $\Box^*_{\mathcal{B}}$ for $\mathcal{B} \subseteq \mathcal{A}$ to get a larger language $\mathcal{L}_1$. In Figure 4, we note the fact that $\mathcal{L}_0$ is literally a subset of $\mathcal{L}_1$ by using the inclusion arrow.

We may also take $\mathcal{L}_0$ and close under countable conjunctions (and hence also disjunctions). We call this language $\mathcal{L}_0^\omega$. Note that $\mathcal{L}_1$ is not literally a subset of $\mathcal{L}_0^\omega$, but there is a translation of $\mathcal{L}_1$ into $\mathcal{L}_0^\omega$ that preserves the semantics. We would indicate this in a chart with a dashed arrow.

PDL is propositional dynamic logic, formulated with atomic programs a corresponding to the agents A that we have fixed. We might note that we can translate $\mathcal{L}_1$ into PDL. The main clauses in the translation are:

$$\begin{array}{lcl} (\Box_A\varphi)^t & = & [a]\varphi^t \\ (\Box^*_{\mathcal{B}}\varphi)^t & = & [(\bigcup_{b\in B} b)^*]\varphi^t \end{array}$$

Beginning in Section 4.2, we study some new languages. These will be based on a third parameter, an *action signature* $\mathbf{\Sigma}$. For each "type of epistemic action" there will be an action

$$\begin{tikzcd}
\mathcal{L}_0 \arrow[r, hook] \arrow[d, hook] & \mathcal{L}_1 \arrow[d, hook] & & & \\
\mathcal{L}_0(\mathbf{\Sigma}) \arrow[r, hook] & \mathcal{L}_1(\mathbf{\Sigma}) \arrow[r, hook] \arrow[dr] & \mathcal{L}(\mathbf{\Sigma}) \arrow[r, dashed] & \mathcal{L}_0^{\omega} \arrow[r, hook] & \mathcal{L}_0^{\infty} \\
& & \mathrm{PDL} \arrow[ur] & &
\end{tikzcd}$$

Figure 4: Languages in this paper

signature $\mathbf{\Sigma}$. For each $\mathbf{\Sigma}$, we'll have languages $\mathcal{L}_0(\mathbf{\Sigma})$, $\mathcal{L}_1(\mathbf{\Sigma})$, and $\mathcal{L}(\mathbf{\Sigma})$. These will be related to the other languages as indicated in the figure.

The idea is then that the modal languages based on $\mathbf{\Sigma}$ figure in to our formalization of modal logic on top of the given type of action. For example, corresponding to the notion of a public announcement is a particular action signature $\mathbf{\Sigma}_{\mathrm{pub}}$, and then the languages $\mathcal{L}_0(\mathbf{\Sigma}_{\mathrm{pub}})$, $\mathcal{L}_1(\mathbf{\Sigma}_{\mathrm{pub}})$, and $\mathcal{L}(\mathbf{\Sigma}_{\mathrm{pub}})$ will have something to do with our notion of a "logic of public announcements."

In Section 8, we compare the expressive power of the languages mentioned so far. It turns out that all of the arrows in Figure 4 correspond to proper increases in expressive power. (There is one exception: $\mathcal{L}_0$ turns out to equal $\mathcal{L}_0(\mathbf{\Sigma})$ in expressive power for all $\mathbf{\Sigma}$.) It is even more interesting to compare expressive power as we change the action signature. For example, we would like to compare logics of public announcement to logics of private announcement to groups. Most of the natural questions in this area are open as of this writing.

2.6 Background: bisimulation

We shall see the concept of bisimulation at various points in the paper, and this section reviews the concept and also develops some of the appropriate definitions concerning bisimulation and updates.

Definition Let $\mathbf{S}$ and $\mathbf{T}$ be state models. A *bisimulation* between $\mathbf{S}$ and $\mathbf{T}$ is a relation $R \subseteq S \times T$ such that whenever $s\ R\ t$, the following three properties hold:

1. $s \in \|p\|_{\mathbf{S}}$ iff $t \in \|p\|_{\mathbf{T}}$ for all atomic sentences p. That is, s and t agree on all atomic information.

2. For all agents A and states s' such that $s \xrightarrow{A} s'$, there is some state t' such that $t \xrightarrow{A} t'$ and $s'\ R\ t'$.

3. For all agents A and states t' such that $t \xrightarrow{A} t'$, there is some state s' such that $s \xrightarrow{A} s'$ and $s'\ R\ t'$.

Example 2.3 This example concerns the update operation $?\boldsymbol{\varphi}$ of Example 2.2. Fix an epistemic proposition $\boldsymbol{\varphi}$ and a state model $\mathbf{S}$. Recall that $(?\boldsymbol{\varphi})_{\mathbf{S}}$ relates each $s \in \boldsymbol{\varphi}$ to the corresponding pair $(0, s)$. We check that the following relation R is a bisimulation between $\mathbf{S}$ and $\mathbf{S}(?\boldsymbol{\varphi})$:

$$R \quad = \quad (?\boldsymbol{\varphi})_{\mathbf{S}} \cup \{(s, (1, s)) : s \in S\}.$$

The definition of $\mathbf{S}(?\boldsymbol{\varphi})$ insures that the interpretations of atomic sentences are preserved by this relation.

Next, suppose that $s \xrightarrow{A} t$ in $\mathbf{S}$ and $s\ R\ (i, s')$. Then we must have $s' = s$. Further, the definition of R tells us that $t\ R\ (1, t)$, and $(i, s) \xrightarrow{A} (1, t)$.

Finally, suppose that $s\ R\ (i, s')$ and $(i, s') \xrightarrow{A} (j, t)$. Then $s' = s$ and $j = 1$. In addition, the definition $\xrightarrow{A}$ in $\mathbf{S}(?\boldsymbol{\varphi})$ implies that $s \xrightarrow{A} t$ in $\mathbf{S}$. And as above, $t\ R\ (1, t)$. This concludes our verifications.

Proposition 2.4 *If there is a bisimulation R such that $s\ R\ t$, then s and t agree on all sentences in* infinitary *modal logic: for all $\varphi \in \mathcal{L}^{\infty}$,*

$$s \models \varphi \textit{ in } \mathbf{S} \quad \textit{iff} \quad t \models \varphi \textit{ in } \mathbf{T}.$$

Recall that SMod is the class of all state models. We have spoken of states as the elements of state models, but we also use the term *states* to refer to the pairs $(\mathbf{S}, s)$ with $s \in S$. The class of all states is itself a state model, except that its collection of states is a proper class rather than a set. Still, it makes sense to talk about bisimulation relations on the class of all states. The largest such is given by

$$(\mathbf{S}, s) \equiv (\mathbf{T}, t) \quad \text{iff} \quad \text{there is a bisimulation } R \text{ between } \mathbf{S} \text{ and } \mathbf{T} \text{ such that } s\ R\ t.$$

This relation $\equiv$ is indeed an equivalence relation.

We say that a proposition $\boldsymbol{\varphi}$ *preserves bisimulations* if whenever $(\mathbf{S}, s) \equiv (\mathbf{T}, t)$, then $s \in \boldsymbol{\varphi}_{\mathbf{S}}$ iff $t \in \boldsymbol{\varphi}_{\mathbf{T}}$. We also say that an update $\mathbf{r}$ *preserves bisimulations* if the following two conditions hold:

1. If $s\ \mathbf{r}_{\mathbf{S}}\ s'$ and $(\mathbf{S}, s) \equiv (\mathbf{T}, t)$, then there is some t' such that $t\ \mathbf{r}_{\mathbf{T}}\ t'$ and $(\mathbf{S}(\mathbf{r}), s') \equiv (\mathbf{T}(\mathbf{s}), t')$.
2. If $t\ \mathbf{r}_{\mathbf{T}}\ t'$ and $(\mathbf{S}, s) \equiv (\mathbf{T}, t)$, then there is some s' such that $s\ \mathbf{r}_{\mathbf{S}}\ s'$ and $(\mathbf{S}(\mathbf{r}), s') \equiv (\mathbf{T}(\mathbf{s}), t')$.

Proposition 2.5 *Concerning bisimulation preservation:*

1. *The bisimulation preserving propositions include the atomic propositions $\mathbf{p}$, and they are closed under all of the (infinitary) operations.*
2. *The bisimulation preserving updates are closed under composition and (infinitary) sums.*
3. *If $\boldsymbol{\varphi}$ and $\mathbf{r}$ preserve bisimulations, so does $[\mathbf{r}]\boldsymbol{\varphi}$.*

Proof We show the last part. Suppose that $s \in ([\mathbf{r}]\boldsymbol{\varphi})_{\mathbf{S}}$, and let $(\mathbf{S}, s) \equiv (\mathbf{T}, t)$. To show that $t \in ([\mathbf{r}]\varphi)_{\mathbf{T}}$, let $t\ \mathbf{r}_{\mathbf{T}}\ t'$. Then by condition (2) above, there is some s' such that $s\ \mathbf{r}_{\mathbf{S}}\ s'$ and $(\mathbf{S}, s') \equiv (\mathbf{T}, t')$ Since $s \in ([\mathbf{r}]\varphi)_{\mathbf{S}}$, we have $s' \in \varphi_{\mathbf{S}(\mathbf{r})}$. And then $t' \in \varphi_{\mathbf{T}(\mathbf{r})}$, since φ too preserves bisimulation. ⊣

2.7 Justifying our models

At this point, we want to return to the models that we pictured in Section 1.1. The concept of bisimulation allows us to justify our models, at least in part.

Justification of the model from Scenario 1 We did not justify in any way our model from Scenario 1. We know that a full discussion of this one example would involve many important matters that this paper does not address. The point we wish to make is that the models and the logic in some sense justify each other. (More precisely, each could be used as a partial justification of the other.) To see how this works in this case, we consider our intuitions about the concealed coin scenario. They are:

1_0 One possibility, call it x, is when the coin is lying Heads up.

2_0 One (other) possibility, call it y, is when the coin is lying Tails up.

3_0 Neither A nor B knows whether the coin lies Heads up or Tails up.

4_0 Both A and B know (3_0).

5_0 Both A and B know (4_0).

6_0 ... similarly.

Shortly we will restrict attention to some state x as in (1_0). Before this, we have a few comments on intuitions (3_0) and (4_0). Both are intended to follow from our description of the scenario back in Section 1.1: both A and B enter the room together, the coin was flipped while they were too far to see it, and then its box was closed. How can they know whether it lies Head or Tails up? In this paper, knowledge is closest to *justified belief.* So the two agents' ignorance about the state of the coin is then something they can justifiably believe. And so is the higher-order knowledge of *this* fact, etc.

Following standard practice, we summarize (3_0), (4_0), (5_0), etc., by saying that it is *common knowledge* for A and B that neither knows whether the coin lies Heads up or Tails up.

If we focus for a moment on a state x where the coin lies Heads up, then our intuitions about x may be summarized in the following way:

1. The coin is lying Heads up.

2. It is common knowledge to A and B that neither of them knows whether the coin lies Heads up or Tails up.

In addition to these we have another sentence that is part of the overall background in all situations concerning coin flipping:

3. It is common knowledge that the coin lies Heads up iff it does not lie Tails up.

Similarly, if we take a state y where the coin lies Tails up, then our intuitions are the same, with the obvious change in (1). We can formalize these in the following sentences of epistemic logic:

$$\begin{array}{ll} \chi_{\mathsf{H}} & \mathsf{H} \wedge \Box^*_{A,B}(\neg\Box_A\mathsf{H} \wedge \neg\Box_A\mathsf{T} \wedge \neg\Box_B\mathsf{H} \wedge \neg\Box_B\mathsf{T}) \wedge \Box^*_{A,B}(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \\ \chi_{\mathsf{T}} & \mathsf{T} \wedge \Box^*_{A,B}(\neg\Box_A\mathsf{H} \wedge \neg\Box_A\mathsf{T} \wedge \neg\Box_B\mathsf{H} \wedge \neg\Box_B\mathsf{T}) \wedge \Box^*_{A,B}(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \end{array} \tag{6}$$

Proposition 2.6 *Let* $\mathbf{T}$ *be any state model, let* $x \in T$, *and let* $C \in \{A, B\}$. *If* $x \models \chi_{\mathsf{H}} \vee \chi_{\mathsf{T}}$, *then*

1. *there is some* y *such that* $x \xrightarrow{C} u$ *and* $y \models \chi_{\mathsf{H}}$*;*
2. *there is some* y *such that* $x \xrightarrow{C} y$ *and* $y \models \chi_{\mathsf{T}}$*;*
3. *for all* y *such that* $x \xrightarrow{C} y$, $y \models \chi_{\mathsf{H}} \vee \chi_{\mathsf{T}}$.

Proposition 2.7 *Let* $\mathbf{S}_1$ *be the state model of Scenario 3 in Section 2.2. Let* s *be the state where* H *holds and* t *the state where* T *holds.*

1. $s \models \chi_{\mathsf{H}}$ *and* $t \models \chi_{\mathsf{T}}$.
2. *Let* $\mathbf{T}$ *be any state model and suppose* $x \in T$ *satisfies* χ_{H}. *Then there is a bisimulation* $R \subseteq S_1 \times T$ *such that* $s\ R\ x$.

Proof Part (1) is verified directly, and for the most part we did the work in Section 2.2. For part (2), we take R to be $s\ R\ x$ iff $x \in [\![\chi_{\mathsf{H}}]\!]_{\mathbf{T}}$ and $t\ R\ x$ iff $t \in [\![\chi_{\mathsf{T}}]\!]_{\mathbf{T}}$. The verification that this is a bisimulation uses Proposition 2.6 and the structure of the model $\mathbf{S}_1$. ⊣

Proposition 2.7 is our justification of the model $\mathbf{S}_1$. Here is the point. Suppose one accepts the intuitions (1)–(3) that we spelled out earlier, the formalization χ_{H} of these intuitions in modal logic, and the Kripke semantics for χ_{H} itself. Then any mathematical model of the intuitions would have to be a state model $\mathbf{T}$ containing some state x (where χ_{H} holds) which is bisimilar to $\mathbf{S}_1$ in a way that related s and x. In fact, the part of $\mathbf{T}$ accessible from x would have to be completely bisimilar to $\mathbf{S}_1$ itself. So in this sense, if one accepts our intuitions and their formalization, then one should accept our model $\mathbf{S}_1$.

Justification of the model from Scenario 3 The reader should review our description of what happens Scenario 3, and also the example sentences in Section 2.2. It is not necessary to look at the picture in Scenario 3, and indeed our goal is to independently obtain that picture using *only* our intuitions and the concept of bisimulation. We want to focus on the state that obtains when the coin shows Heads and A has learned this after opening the box herself in front of B in the manner we described. Our intuitions are that the following sentences should be true about this state:

1. The coin lies Heads up.
2. It is common knowledge that: if the coin lies Heads up, then A knows this and considers this possible.

3. It is common knowledge that had the coin been lying Tails up, then A would know this and consider it possible.

4. It is common knowledge that B does not know whether the coin lies Heads up or Tails up.

The additional points in (2) that A knows that the coin's disposition is *possible* might strike the reader as unnecessary. However, as we shall see when we discuss the cheating example below, we are going to consider examples where an agent does not think the actual world is possible. But in modeling a case where an agent can *see* that the coin lies (say) Heads up, then not only do they know H in the sense of H being true in all their possible worlds – this might hold vacuously – they also should know that H holds in *some* world. At the same time (2) could be strengthened to say that not only does A know H, but this knowledge would hold introspectively: A would know that A knows this, etc. We agree with this intuition, but we shall not need it. It follows from our formulation of common knowledge.

Point (4) is a counterfactual statement. It conforms to our intuitions about what is happening in the Scenario overall.

In addition to (1)–(4) above, we again have the intuition that in all situations concerning coin flipping it is common knowledge that the coin lies Heads up iff it does not lie Tails up.

Now we take a single sentence of epistemic logic, the conjunction of the formal versions of the all the intuitions above:

$$\begin{array}{lll} \mathsf{H} & \wedge & \Box^*_{A,B}(\mathsf{H} \to \Box_A \mathsf{H} \wedge \Diamond_A \mathsf{H}) \\ & \wedge & \Box^*_{A,B}(\mathsf{T} \to \Box_A \mathsf{T} \wedge \Diamond_A \mathsf{T}) \\ & \wedge & \Box^*_{A,B}\neg(\Box_B \mathsf{H} \vee \Box_B \mathsf{T}) \\ & \wedge & \Box^*_{A,B}(\mathsf{H} \leftrightarrow \neg \mathsf{T}) \end{array}$$

Call this sentence φ. We also consider ψ, the same sentence except that we start with replace the conjunct H (on the first line by itself above) by T.

Proposition 2.8 *Let* $\mathbf{T}$ *be any state model, and let* $x \in T$.

1. *If* $x \models \varphi$, *then there is some* y *such that* $x \xrightarrow{A} y$ *and* $y \models \varphi$*; moreover, all* u *such that* $x \xrightarrow{A} y$ *are such that* $y \models \varphi$.

2. *If* $x \models \psi$, *then there is some* y *such that* $x \xrightarrow{A} y$ *and* $y \models \psi$*; moreover, all* y *such that* $x \xrightarrow{A} y$ *are such that* $y \models \varphi$.

3. *If* $x \models \varphi \vee \psi$, *then*

   (a) *there is some* y *such that* $x \xrightarrow{B} u$ *and* $y \models \varphi$*;*

   (b) *there is some* y *such that* $x \xrightarrow{B} y$ *and* $y \models \psi$*;*

   (c) *for all* y *such that* $x \xrightarrow{B} y$, $y \models \varphi \vee \psi$.

Proof This is routine, using facts like

$$\Box^*_{A,B}\chi \to (\Box_A\Box^*_{A,B}\chi \wedge \Box_B\Box^*_{A,B}\chi). \tag{7}$$

The most interesting parts of the justification concerns assertion (3c) that all B-children y of a state x satisfying $\varphi \vee \psi$ must themselves satisfy this same sentence. The point is that y satisfies either H or T by the formal version of our statement that it is common knowledge that Heads and Tails are exclusive and exhaustive. The rest of the clauses follow from (7). $\dashv$

Proposition 2.9 *Let $\mathbf{S}_3$ be the state model of Scenario 3 in Section 2.2.*

1. *$s \models \varphi$ and $t \models \psi$.*

2. *Let $\mathbf{T}$ be any state model and suppose $x \in T$ satisfies φ. Then there is a bisimulation $R \subseteq S_3 \times T$ such that $s\ R\ x$.*

This last result justifies the model $\mathbf{S}_3$ in the same way that Proposition 2.7 justifies $\mathbf{S}_1$.

Justification of the model from Scenario 4 We next want to repeat the same kind of discussion for the "cheating" example of Scenario 4. Let s be the state that obtains after this scenario; recall that the coin really lies heads up. Here our intuitions about s:

1. The coin lies Heads up.

2. B believes that both Heads and Tails are possible.

3. B believes that it is common knowledge that nobody knows the disposition of the coin.

4. A knows (1)–(3), and in the strong sense that A cannot imagine any possibility that did not conform to (1)–(3). Moreover, A knows that some state is indeed possible (hence (the conjunction of (1)–(3) is possible).

Formalizing this, let NOBODY KNOWS be $\neg(\Box_A\mathsf{H} \vee \Box_A\mathsf{T} \vee \Box_B\mathsf{H} \vee \Box_B\mathsf{T})$, and consider

$$\begin{array}{ll} & \Box^*_A(\mathsf{H} \wedge \Diamond_B\mathsf{H} \wedge \Diamond_B\mathsf{T} \wedge \Box_B\Box^*_{A,B}\text{NOBODY KNOWS}) \\ \wedge & \Diamond_A\mathsf{True} \\ \wedge & \Box^*_{A,B}(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \end{array} \tag{8}$$

We again added the clause that it is common knowledge that Heads and Tails are exclusive and exhaustive. Let ψ be this sentence, let ψ_H be $\mathsf{H} \wedge \neg\mathsf{T} \wedge \Box^*_{A,B}$NOBODY KNOWS, and let ψ_T be $\neg\mathsf{H} \wedge \mathsf{T} \wedge \Box^*_{A,B}$NOBODY KNOWS.

Our justification of $\mathbf{S}_4$ is now based on the following results:

Proposition 2.10 *Let $\mathbf{T}$ be any state model, and let $x \in T$.*

1. *If $x \models \psi$, then there is some y such that $x \xrightarrow{A} y$ and $y \models \psi$; moreover, all such y satisfy $y \models \psi$.*

2. *If $x \models \psi$, then there is some y such that $x \xrightarrow{B} y$ and $y \models \psi_H$; and also some y such that $x \xrightarrow{B} y$ and $y \models \psi_T$; moreover, all such y satisfy $\psi_H \vee \psi_T$.*

3. *If* $x \models \psi_H \vee \psi_T$ *and* ($C = A$ *or* $C = B$)*, then there is some* y *such that* $x \xrightarrow{C} y$ *and* $y \models \psi_H$*; and there is also some* y *uch that* $x \xrightarrow{A} y$ *and* $y \models \psi_T$*. Moreover, all such* y *satisfy* $\psi_H \vee \psi_T$*.*

Proposition 2.11 *Let* $\mathbf{S}_4$ *be the state model of Scenario 4 in Section 2.2.*

1. $s \models \psi$*, and* $t \models \psi_H$ *and* $u \models \psi_T$*.*

2. *Let* $\mathbf{T}$ *be any state model and suppose* $x \in T$ *satisfies* ψ*. Then there is a bisimulation* $R \subseteq S_4 \times T$ *such that* $s\ R\ x$*.*

2.8 Our nearest neighbors

This paper will construct new logical languages and proof systems. We have indicated the names of the new languages in Figure 4 along with some of the better-known ones to serve as mileposts to orient the reader. We want to continue the orientation by further reviewing the known languages at this point and also mentioning their proof systems. None of the facts of this section are used in the sequel.

First, in Figure 5 we see the syntax of the multi-agent version of the modal logic $\mathcal{L}_0$ and the modal logic K presented in its best-known form, as a Hilbert-style proof system. (To read this from Figure 5, drop everything that has a $*$.) Recall that $\mathcal{L}_1$ is $\mathcal{L}_0$ extended with common knowledge operators. Figure 5 again shows the syntax of $\mathcal{L}_1$, and the proof system for $\mathcal{L}_1$ adds to K the Mix Axiom and the Induction Rule.

Another close neighbor of the system in this paper is Propositional Dynamic Logic (PDL). PDL was first formulated by Fischer and Ladner [6], following the introduction of dynamic logic in Pratt [19]. The syntax and the main clauses in the semantics of PDL are shown in Figure 6, along with the *Segerberg Axioms* for it. In PDL, there are two types of syntactic objects, sentences and programs. The programs in PDL are interpreted on a structure by relations on that structure. This is *not* the way our semantics works, and to make this point we compare the standard semantics of (a fragment of) PDL with a language closer to what we shall ultimately study.

Reformulation of Test-only PDL To make this point, we consider the *test-only* fragment of PDL in our terms. This is the fragment built without using atomic programs, so the programs are *skip*, the tests $?\varphi$ and compositions, choices, or iterations of these. We give a reformulation in Figure 7. The point is that in PDL the programs are interpreted by relations on a given model, and in our terms programs are interpreted by *updates*. We have discussed updates of the form $?\boldsymbol{\varphi}$ in Example 2.2. Given that we have an interpretation of the sentence φ as an epistemic proposition $[\![\varphi]\!]$, we then automatically have an update $?[\![\varphi]\!]$.

For the sentences, the main thing to look at is the semantics of sentences $[\pi]\varphi$; here we use the semantic notions from Section 2.3. The way the semantics works is that we have $[\![\pi]\!]$ and $[\![\varphi]\!]$; the former is an update and the latter is an epistemic proposition. Then we use both of these to get an overall semantics, using equation (4). In more explicit terms,

$$\begin{array}{rcl} [\![[\pi]\varphi]\!]_{\mathbf{S}} & = & ([\![\pi]\!][[\![\varphi]\!]])_{\mathbf{S}} \\ & = & \{s \in S : \text{if } s\ [\![\pi]\!]_{\mathbf{S}}\ t, \text{ then } t \in [\![\varphi]\!]_{\mathbf{S}}\} \end{array}$$

Syntax	**Sentences** φ	$p_i \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box_A \varphi \mid \Box^*_{\mathcal{B}}\varphi$
Semantics	**Main Clauses**	$[\![\Box_A\varphi]\!]_{\mathbf{S}} = \Box_A[\![\varphi]\!]_{\mathbf{S}}$ $[\![\Box^*_{\mathcal{B}}\varphi]\!]_{\mathbf{S}} = \Box^*_{\mathcal{B}}[\![\varphi]\!]_{\mathbf{S}}$
Proof System	**Basic Axioms**	
	All sentential validities	
	($\Box_A$-normality)	$\vdash \Box_A(\varphi \to \psi) \to (\Box_A\varphi \to \Box_A\psi)$
$*$	($\Box^*_{\mathcal{C}}$-normality)	$\vdash \Box^*_{\mathcal{C}}(\varphi \to \psi) \to (\Box^*_{\mathcal{C}}\varphi \to \Box^*_{\mathcal{C}}\psi)$
$*$	**Mix Axiom**	$\vdash \Box^*_{\mathcal{C}}\varphi \to \varphi \wedge \bigwedge\{\Box_A\Box^*_{\mathcal{C}}\varphi : A \in \mathcal{C}\}$
	Modal Rules	
	(Modus Ponens)	From $\vdash \varphi$ and $\vdash \varphi \to \psi$, infer $\vdash \psi$
	($\Box_A$-necessitation)	From $\vdash \varphi$, infer $\vdash \Box_A\varphi$
$*$	($\Box^*_{\mathcal{C}}$-necessitation)	From $\vdash \varphi$, infer $\vdash \Box^*_{\mathcal{C}}\varphi$
$*$	**Induction Rule**	From $\vdash \chi \to \psi$ and $\vdash \chi \to \Box_A\chi$ for $A \in \mathcal{C}$, infer $\vdash \chi \to \Box^*_{\mathcal{C}}\psi$

Figure 5: The languages $\mathcal{L}_0$ and $\mathcal{L}_1$, and proof systems for them. For $\mathcal{L}_0$, we drop the $\Box^*_{\mathcal{B}}$ construct in the syntax, and the $*$ material in the proof system.

Soundness of the test axiom We conclude this discussion with a sketch of the verification that the PDL Test Axiom $[?\varphi]\psi \leftrightarrow (\varphi \to \psi)$ is sound for our semantics. This justifies our referring to the updates $?\boldsymbol{\varphi}$ as *tests*. Fix a state model $\mathbf{S}$. Then the following are equivalent:

1. $s \in [\![[?\varphi]\psi]\!]_{\mathbf{S}}$.
2. For all t such that $s(?[\![\varphi]\!]_{\mathbf{S}})t$, $t \in [\![\psi]\!]_{\mathbf{S}(?[\![\varphi]\!])}$.
3. If $s \in [\![\varphi]\!]_{\mathbf{S}}$, then $(0, s) \in [\![\psi]\!]_{\mathbf{S}(?[\![\varphi]\!])}$.
4. If $s \in [\![\varphi]\!]_{\mathbf{S}}$, then $s \in [\![\psi]\!]_{\mathbf{S}}$.
5. $s \in [\![\varphi \to \psi]\!]_{\mathbf{S}}$.

All of the equivalences are immediate from general definitions which we have seen, except for (2)$\Longleftrightarrow$(3) and (3)$\Longleftrightarrow$(4). For (2)$\Longleftrightarrow$(3), note that the domain of the relation $(?[\![\varphi]\!])_{\mathbf{S}}$ is the set $[\![\psi]\!]_{\mathbf{S}}$. Furthermore, if s is in this set, then the only state of $\mathbf{S}(?[\![\varphi]\!])$ which it is related to is $(0, s)$. Concerning the equivalence (3)$\Longleftrightarrow$(4), note first that if $s \in [\![\varphi]\!]_{\mathbf{S}}$, then there is a bisimulation between $\mathbf{S}$ and $\mathbf{S}(?[\![\varphi]\!])$ relating s and $(0, s)$; the bisimulation is given in Example 2.3. Second, an easy induction on the overall languages using Proposition 2.5 shows that each sentence of it is interpreted be a *bisimulation preserving* proposition. Putting these two facts together gives the desired equivalence (3)$\Longleftrightarrow$(4).

Syntax	**Sentences** φ	$p_i \mid \neg\varphi \mid \varphi \wedge \psi \mid [\pi]\varphi$
	Programs π	$\pi \mid ?\varphi \mid \pi;\sigma \mid \pi + \sigma \mid \pi^*$
Semantics	**Main Clauses**	$[\![[\alpha]\varphi]\!]_{\mathbf{S}} = \{s : \text{if } s[\![\alpha]\!]_{\mathbf{S}}t, \text{ then } t \in [\![\varphi]\!]_{\mathbf{S}}\}$
		$[\![?\varphi]\!]_{\mathbf{S}} = \{(s,s) : s \in [\![\varphi]\!]_{\mathbf{S}}\}$
		$[\![\pi;\sigma]\!]_{\mathbf{S}} = [\![\pi]\!]_{\mathbf{S}} \cdot [\![\sigma]\!]_{\mathbf{S}}$
		$[\![\pi + \sigma]\!]_{\mathbf{S}} = [\![\pi]\!]_{\mathbf{S}} \cup [\![\sigma]\!]_{\mathbf{S}}$
		$[\![\pi^*]\!]_{\mathbf{S}} = ([\![\pi]\!]_{\mathbf{S}})^*$
Proof System	**Basic Axioms**	
	All sentential validities	
	($[\pi]$-normality)	$\vdash [\pi](\varphi \to \psi) \to ([\pi]\varphi \to [\pi]\psi)$
	Program Axioms	
	(Test)	$\vdash [?\varphi]\psi \leftrightarrow (\varphi \to \psi)$
	(Skip)	$\vdash [\mathit{skip}]\varphi \leftrightarrow \varphi$
	(Composition)	$\vdash [\pi;\sigma]\varphi \leftrightarrow [\pi][\sigma]\varphi$
	(Choice)	$\vdash [\pi + \sigma]\varphi \leftrightarrow ([\pi]\varphi \wedge [\sigma]\varphi)$
	Mix Axiom	$\vdash [\pi^*]\varphi \to \varphi \wedge [\pi][\pi^*]\varphi$
	Modal Rules	
	(Modus Ponens)	From $\vdash \varphi$ and $\vdash \varphi \to \psi$, infer $\vdash \psi$
	($[\pi]$-necessitation)	From $\vdash \varphi$, infer $\vdash [\pi]\varphi$
	Induction Rule	From $\vdash \chi \to \psi$ and $\vdash \chi \to [\pi]\chi$ infer $\vdash \chi \to [\pi^*]\psi$

Figure 6: Propositional Dynamic Logic (PDL)

Syntax	**Sentences** φ	$p_i \mid \neg\varphi \mid \varphi \wedge \psi \mid [\pi]\varphi$
	Programs π	$?\varphi \mid skip \mid \pi;\sigma \mid \varphi + \psi \mid \pi^*$
Semantics	**Main Clauses**	$[\![[\pi]\varphi]\!]_{\mathbf{S}} = [\![\pi]\!][[\![\varphi]\!]]$
		$[\![skip]\!] = \mathbf{1}$ from Section 2.3
		$[\![?\varphi]\!] = ?[\![\varphi]\!]$ as in Example 2.2
		$[\![\pi;\sigma]\!] = [\![\pi]\!] \cdot [\![\sigma]\!], [\![\pi + \sigma]\!] = [\![\pi]\!] + [\![\sigma]\!]$
		$[\![\pi^*]\!] = [\![\pi]\!]^*$
Proof System	**Inherited from PDL**	

Figure 7: Test-only PDL, with a semantics in our style

Endnotes As far as we know, the first paper to study the interaction of communication and knowledge in a formal setting is Plaza's paper "Logics of Public Communications" [18]. As the title suggests, the epistemic actions studied are public announcements. In essence, he formalized the logic of public announcements. (In addition to this, [18] contains a number of results special to the logic of announcements which we have not generalized, and it also studies an extension of the logic with non-rigid constants.) The same formalization was made in Gerbrandy [7, 8], and also Gerbrandy and Groeneveld [9]. These papers further formalize the logic of completely private announcements. (However, their semantics use non-wellfounded sets rather than arbitrary Kripke models. As pointed out in [15], this restriction was not necessary.) The logic of common knowledge of alternatives was formulated in [2] and also in van Ditmarsch's dissertation [22].

Our introduction of updates is new here, as are the observations on the test-only fragment of PDL.

Readers familiar with characterizing sentences in modal logic (sometimes called Scott sentences) will recognize that what our main move in Section 2.7 is to write out a characterizing sentence for the models in question in the language $\mathcal{L}_1$. (Every finite Kripke model has a characterizing sentence in $\mathcal{L}_1$.) We wrote our characterizing sentences after putting down some intuitions about the models. This kind of justification of models in epistemic logic seems natural, and we suspect others have done it before us.

In recent times, one very active arena for work on knowledge is distributed systems, and the main source of work in recent times on knowledge in distributed systems is the book **Reasoning About Knowledge** [5] by Fagin, Halpern, Moses, and Vardi. We shall depart from [5] by introducing operators whose semantics are updates as we have defined them, and by doing without temporal logic operators. In effect, our Kripke models are simpler, since they do not incorporate all of the runs of a system; the new operators can be viewed as a compensation for that. We have not made a detailed comparison of our work with the large body of work on knowledge on distributed systems, and such a comparison would require both technical and conceptual results. On the technical side, we suspect that neither framework is translatable into the other. One way to show this would be by expressivity results. Perhaps another way would use complexity results. In this direction, we note that Halpern and Vardi [11] examine ninety-six logics of knowledge and time. Thirty-two of these contain common knowledge operators, and of these, all but twelve of these are undecidable. But overall, our logics are based on differing conceptual points and intended applications, and so we are confident that they differ.

3 The update product operation

In this section, we present the centerpiece of the formulation of our logical systems. The section formulates some new notions, namely *simple action structures*, *epistemic action models*, and an *update product* operation. The leading idea is that epistemic actions, like state models, have the property that different agents think different things are possible in the various states. To model effect of an epistemic update on a state, we use a kind of *product of alternatives.*

3.1 Simple action structures

Let Φ be a the collection of all epistemic propositions. A *simple action structure* is a triple $\mathbf{\Sigma} = (\Sigma, \xrightarrow{\mathcal{A}}, \mathrm{pre})$, where Σ is a set of *simple actions*, $\xrightarrow{\mathcal{A}}$ is an $\mathcal{A}$-indexed family of relations on Σ, and $\mathrm{pre} : \Sigma \to \Phi$. A simple action structure is similar to a state model. We use different notation and terminology because of a technical difference and a bigger conceptual point. The technical difference is that $\mathrm{pre} : \Sigma \to \Phi$ (that is, the codomain is the collection of all epistemic propositions). The conceptual point is that we think of "simple" actions as being deterministic actions. We call the members of the set Σ "simple actions" (instead of states). For a simple action $\sigma \in \Sigma$, we interpret $\mathrm{pre}(\sigma)$ as giving the *precondition of* σ; this is the fact that needs to hold at a state (in a state model) in order for action σ to be "accepted" in that state. So σ will be executable in s iff its precondition $\mathrm{pre}(\sigma)$ holds at s. At this point we have mentioned the ways in which simple action structures and state models differ. What they have in common is that they use accessibility relations to express each agent's uncertainty concerning something. For state models, the uncertainty has to do with which state is the real one; for simple action structure, it has to do with which action is taking place.

Example 3.1 Here is a simple action structure:

$$\begin{tikzcd} \boxed{\sigma{:}\mathsf{H}} \arrow[loop left, "A"] \arrow[rr, "B"] & & \boxed{\tau{:}\mathbf{True}} \arrow[loop right, "A,B"] \end{tikzcd}$$

Formally, $\Sigma = \{\sigma, \tau\}$; $\sigma \xrightarrow{A} \sigma$, $\sigma \xrightarrow{B} \tau$, $\tau \xrightarrow{A} \tau$, $\tau \xrightarrow{B} \tau$; $\mathrm{pre}(\sigma) = \mathsf{H}$, and $\mathrm{pre}(\tau) = \mathbf{True}$.

As we shall see, this simple action structure will be used in the modeling of a *completely private announcement to A that the coin is lying Heads up.* Further examples may be found later in this section.

3.2 Epistemic action models

To model non-determinism, we define an *epistemic action model* to be a pair $(\mathbf{\Sigma}, \Gamma)$ consisting of a simple action structure $\mathbf{\Sigma}$ and a set Γ of *designated simple actions.* Each of the simple actions $\gamma \in \Gamma$ can be thought as being a possible "deterministic resolution" of the non-deterministic action α. As announced above, the intuition about the map pre is that an action is executable in a given state only if all its preconditions hold at that state. We often spell out an epistemic action model as $(\Sigma, \xrightarrow{\mathcal{A}}, \mathrm{pre}, \Gamma)$ rather than $((\Sigma, \xrightarrow{\mathcal{A}}, \mathrm{pre}), \Gamma)$.

Example 3.2 The subset Γ may be arbitrary in an epistemic action model. To mention one choice that we shall use below, let's return to Example 3.1. We would take $\Gamma = \{\sigma\}$.

3.3 The update product of a state model with an epistemic action model

In all of our work in the remainder of this paper, the following operation will prove to play a central role.

Given a state model $\mathbf{S} = (S, \xrightarrow{A}_{\mathbf{S}}, \|\cdot\|_{\mathbf{S}})$ and a simple action structure $\mathbf{\Sigma} = (\Sigma, \xrightarrow{A}, \mathrm{pre})$, we define their *update product* to be the state model

$$\mathbf{S} \otimes \mathbf{\Sigma} \quad = \quad (S \otimes \Sigma, \xrightarrow{A}, \|.\|),$$

given by the following: the new states are pairs of old states s and simple actions σ which are "consistent", in the sense that all preconditions of the action σ "hold" at the state s

$$S \otimes \Sigma \quad = \quad \{(s, \sigma) \in S \times \Sigma : s \in [\![\mathrm{pre}(\sigma)]\!]_{\mathbf{S}}\}. \tag{9}$$

The new accessibility relations are taken to be the "products" of the corresponding accessibility relations in the two frames; i.e., for $(s, \sigma), (s', \sigma') \in \mathbf{S} \otimes \mathbf{\Sigma}$ we put

$$(s, \sigma) \xrightarrow{A} (s', \sigma') \quad \text{iff} \quad s \xrightarrow{A} \sigma \text{ and } s' \xrightarrow{A} \sigma', \tag{10}$$

and the new valuation map $\|.\|_{\mathbf{S}}$: $\mathsf{AtSen} \to \mathcal{P}(S \otimes \Sigma)$ is essentially given by the old valuation:

$$\|p\|_{\mathbf{S}\otimes\mathbf{\Sigma}} \quad = \quad \{(s, \sigma) \in S \otimes \Sigma : s \in \|p\|_{\mathbf{S}}\}. \tag{11}$$

Remark The update operation is defined whenever $\mathbf{S}$ is a state model and whenever the codomain of pre is the collection of all epistemic propositions. A different point concerns the interpretations $[\![\boldsymbol{\varphi}]\!]_{\mathbf{S}\otimes\mathbf{\Sigma}}$ in the product $\mathbf{S} \otimes \mathbf{\Sigma}$. By (11) and (5) from Section 2.5, we see that

$$\boldsymbol{p}_{\mathbf{S}\otimes\mathbf{\Sigma}} \quad = \quad \{(s, \sigma) \in S \otimes \Sigma : s \in \boldsymbol{p}_{\mathbf{S}}\}.$$

We should stress that this last fact holds only for *atomic* propositions $\boldsymbol{p}$. It will not hold in general for arbitrary epistemic propositions $\boldsymbol{\varphi}$.

Intended interpretation The update product restricts the full Cartesian product $S \times \Sigma$ to the smaller set $S \otimes \Sigma$ in order to insure that *states survive actions* in the appropriate sense.

For each agent A, the product arrows $\xrightarrow{A}$ on the output frame represent agent A's epistemic uncertainty about the output state. The intuition is that for these particularly simple kinds of actions that we take as "simple actions", the uncertainty regarding the action is assumed to be independent of the uncertainty regarding the current (input) state. This independence allows us to "multiply" these two uncertainties in order to compute the uncertainty regarding the output state: if whenever the input state is s, agent A thinks the input might be some other state s', and if whenever the current action happening is σ, agent A thinks the current action might be some other action σ', then whenever the output state (s, σ) is reached, agent A thinks the alternative output state (s', σ') might have been reached. Moreover, these are the only output states that A considers possible.

As for the valuation, we essentially take the same valuation as the one in the input model. If a state s survives an action, then the same facts p hold at the output state (s, σ) as at the input state s. This means that *our actions, if successful, do not change the facts.* This condition can of course be relaxed in various ways, to allow for fact-changing actions. But in this paper we are primarily concerned with purely epistemic actions, such as the earlier examples in this section.

3.4 Updates induced by epistemic action models

Recall that we defined updates and proper epistemic actions in Section 2.3. Right above, in Section 3.2, we defined epistemic action *models*. Note that there is a big difference: the updates are pairs of operations on the class of all state models, and the epistemic action models are typically finite structures. We think of updates as essentially semantic, and epistemic action models as mainly syntactic. (They have a semantic component, the map pre.) Here is a connection.

Definition Let $(\mathbf{\Sigma}, \Gamma)$ be an epistemic action model. We define an update which we also denote $(\mathbf{\Sigma}, \Gamma)$ as follows:

1. $\mathbf{S}(\mathbf{\Sigma}, \Gamma) = \mathbf{S} \otimes \mathbf{\Sigma}$.

2. $s\ (\mathbf{\Sigma}, \Gamma)_{\mathbf{S}}\ (t, \sigma)$ iff $s = t$ and $\sigma \in \Gamma$.

We call this *the update induced by* $(\mathbf{\Sigma}, \Gamma)$.

Bisimulation preservation Before moving to examples, we note a simple result that will be used later.

Proposition 3.3 *Let* $\mathbf{\Sigma}$ *be a simple action structure in which every* $pre(\sigma)$ *is a bisimulation preserving epistemic proposition. Let* $\Gamma \subseteq \Sigma$ *be arbitrary. Write* $\mathbf{r}$ *for the update induced by the epistemic action model* $(\mathbf{\Sigma}, \Gamma)$*. Then* $\mathbf{r}$ *preserves bisimulation.*

Proof Fix $\mathbf{S}$ and $\mathbf{T}$, and suppose that $s \equiv t$. Suppose that $s\ \mathbf{r}_{\mathbf{S}}\ s'$, so $s' \in \mathbf{S}(\mathbf{r})$ is of the form (s, σ) for some $\sigma \in \Gamma$. Then $(t, \sigma) \in \mathbf{T}(\mathbf{r})$, and clearly $t\ \mathbf{r}_{\mathbf{T}}\ (t, \sigma)$, and we need only show that $(s, \sigma) \equiv (t, \sigma)$. But the following relation R is a bisimulation between $\mathbf{S}(\mathbf{r})$ and $\mathbf{T}(\mathbf{r})$:

$$(s', \tau_1)\ R\ (t', \tau_2) \quad \text{iff} \quad s' \equiv t' \text{ and } \tau_1 = \tau_2.$$

The verification of the bisimulation properties is easy. And R shows that $(s, \sigma) \equiv (t, \sigma)$, as desired. $\dashv$

3.5 The coin scenario models as examples of the update product

We return to the coin scenarios of Section 1.1. Our purpose is partly to indicate how one may obtain such models as examples of the update product, and at the same time to exemplify the update product construction itself. In this section, the set $\mathcal{A}$ of agents is $\{A, B\}$, and the set AtSen of atomic propositions is $\{\mathsf{H}, \mathsf{T}\}$.

Example 3.4 We begin with an example worked out with many details, and then the rest of our examples will omit many similar calculations. This example has to do with Scenario 4 from Section 1.1, where A takes a look at the coin in such a way that B does not suspect anything (but A knows this). We take as $\mathbf{S}_1$ and $\mathbf{\Sigma}_4$ the structures shown below:

$$\begin{tikzcd} \boxed{s{:}\mathsf{H}} \arrow[loop left, "A,B"] \arrow[r, leftrightarrow, "A,B"] & \boxed{t{:}\mathsf{T}} \arrow[loop right, "A,B"] \end{tikzcd} \qquad \begin{tikzcd} \boxed{\sigma{:}\mathsf{H}} \arrow[loop left, "A"] \arrow[r, "B"] & \boxed{\tau{:}\mathbf{True}}_{A,B} \end{tikzcd}$$

$\mathbf{S}_1$ is from Scenario 1. In $\mathbf{\Sigma}_4$, we take the set of distinguished states to be $\{\sigma\}$. $\mathbf{\Sigma}$ comes from Scenario 4 and also Examples 3.1 and 3.2. To take the update product, we first form the cartesian product $S_1 \times \Sigma_4$:

$$\{(s,\sigma),(s,\tau),(t,\sigma),(t,\tau)\}$$

Of these four pairs, we only want those whose first component satisfies (in $\mathbf{S}$) the precondition of the second component. We do not want (t,σ), since $\mathrm{pre}(\sigma) = \mathsf{H}$ and $t \notin [\![\mathsf{H}]\!]_{\mathbf{S}}$. But the other three pairs do satisfy our condition. So the state model $\mathbf{S}_1 \otimes \mathbf{\Sigma}_4$ will have three states: (s,σ), (s,τ), and (t,τ). The atomic information is inherited from the first component, so we have $[\![\mathsf{H}]\!]_{\mathbf{S}_1\otimes\mathbf{\Sigma}_4} = \{(s,\sigma),(s,\tau)\}$ and $[\![\mathsf{T}]\!]_{\mathbf{S}_1\otimes\mathbf{\Sigma}_4} = \{(t,\tau)\}$. The accessibility relations $\xrightarrow{A}$ and $\xrightarrow{B}$ are those of the product. For example, we have $(s,\sigma) \xrightarrow{B} (s,\tau)$, because $s \xrightarrow{B} s$ and $\sigma \xrightarrow{B} \tau$. But we do not have $(s,\sigma) \xrightarrow{A} (s,\tau)$, because $\sigma \xrightarrow{A} \tau$ is false.

Now, we rename the states as follows:

$$(s,\sigma) \;\rightsquigarrow\; s \qquad\qquad (s,\tau) \;\rightsquigarrow\; t \qquad\qquad (t,\tau) \;\rightsquigarrow\; u$$

And then we get a picture of this state model, the same one we had in Scenario 4:

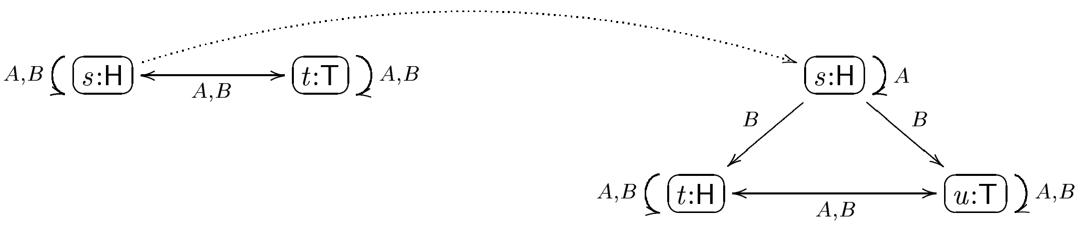

The dotted line here shows the update relation between s and (s,σ). This is the only update relation. For example, we do not relate s and (s,τ) because $\tau \notin \Gamma = \{\sigma\}$. Let $\mathbf{S}_4 = \mathbf{S}_1 \otimes \mathbf{\Sigma}_4$.

Example 3.5 Next we construct the model corresponding to Scenario 5, where B cheats after Scenario 4. We consider the update product of the state model $\mathbf{S}_4$ from Example 3.4 above with the epistemic action model $\mathbf{\Sigma}_5$ shown below:

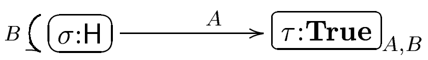

It represents cheating by B. The update product $\mathbf{S}_4 \otimes \mathbf{\Sigma}_5$ has five states: (s,σ), (t,σ), (s,τ), (t,τ), and (u,τ). Notice that (u,σ) is omitted, since $u \notin [\![\mathsf{H}]\!]_{\mathbf{S}_4}$. Again, we rename to simplify the notation:

$$\begin{array}{lllllllll} (s,\sigma) & \rightsquigarrow & v & \quad (s,\tau) & \rightsquigarrow & w & \quad (t,\sigma) & \rightsquigarrow & x \\ (t,\tau) & \rightsquigarrow & y & \quad (u,\tau) & \rightsquigarrow & z & & & \end{array}$$

We leave it to the reader to calculate the accessibility relations in $\mathbf{S}_4 \otimes \mathbf{\Sigma}_5$; they are shown in Figure **??** along with the relation between the models.

Incidentally, we posed in Section 1.1 the exercise of constructing this representation from first principles. Many people are able to construct the five state picture, and some others construct a related picture with seven states. The seven state picture is bisimilar to the one illustrated here.

Example 3.6 We next look back at Scenario 2. The simplest action structure is $\mathbf{\Sigma}_2$: $\boxed{\mathsf{H}}_{A,B}$. It represents a public announcement to A and B that the coin is lying Heads up. Here, the distinguished set Γ is the entire action structure. For the record, we formalize Σ_2 as a singleton set $\{Pub\}$. We have $Pub \xrightarrow{A} Pub$ and $Pub \xrightarrow{B} Pub$. Also, we set $\mathrm{pre}(Pub) = \mathsf{H}$. We did not put the name Pub in the representation above, but in situations where we want the name we would write $\boxed{Pub\text{: } \mathsf{H}}_{A,B}$. Let us call the same structure $\mathbf{S}_2$ when we view it as a state model; formally these are the same object.

For most $\mathbf{S}$, $\mathbf{S} \otimes \mathbf{\Sigma}_2$ is bisimilar to $\mathbf{S}_2$. The only time this fails is when $\mathbf{S}$ contains a simple action s such that either s has no A-successors where H is true, or else no B-successors where H is true.

Example 3.7 Let $\mathbf{\Sigma}_3$ be the following epistemic action model:

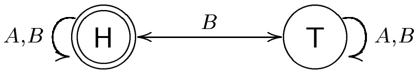

We are using the doubled circle to indicate the set Γ; we maintain this in some of the examples below. $\mathbf{\Sigma}_3$ represents an announcement to A of Heads in the manner of Scenario 3. That is, B knows that A either sees Heads or sees Tails, but not which. Similarly, let let $\mathbf{\Sigma}'_3$ represent the same kind of announcement to B:

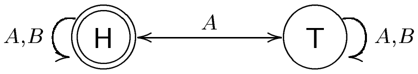

Then we have the following:

$\mathbf{S}_1 \otimes \mathbf{\Sigma}_3 \cong \mathbf{S}_3$, where $\mathbf{S}_3$ is the model in Scenario 3.

$\mathbf{S}_3 \otimes \mathbf{\Sigma}'_3 \cong \mathbf{S}_2$. This conforms with the intuition that successive semi-private viewings by the two parties of the concealed coin amount to a public viewing.

$\mathbf{S}_3 \otimes \mathbf{\Sigma}_3 \cong \mathbf{S}_3$. There is no point for A to look twice.

$\mathbf{S}_4 \otimes \mathbf{\Sigma}_3 \cong \mathbf{S}_3$, where $\mathbf{S}_4$ is the model from Scenario 4.

Example 3.8 Let $\mathbf{\Sigma}_4$ be the following epistemic action model, representing the action where A cheats, learning Heads or Tails in a way which is oblivious to B:

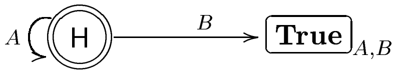

Let $\mathbf{\Sigma}'_4$ be corresponding model with A and B interchanged. Then

$\mathbf{S}_1 \otimes \mathbf{\Sigma}_4 \cong \mathbf{S}_4$.

$\mathbf{S}_4 \otimes \mathbf{\Sigma}'_4$ is the model discussed in Scenario 5.

Incidentally, we would obtain the same results, modulo bisimilarity, if we used not $\mathbf{\Sigma}_4$ but rather the epistemic action model $\mathbf{\Sigma}_4'$:

This epistemic action model consists of two disjoint components. The reason is that there is a bisimulation between $\mathbf{\Sigma}_4$ and $\mathbf{\Sigma}_4'$ relating the distinguished action of Σ_4 to the distinguished H action of the other structure.

Example 3.9 To obtain the model in Scenario 6, we use the following epistemic action model $\mathbf{\Sigma}_6$:

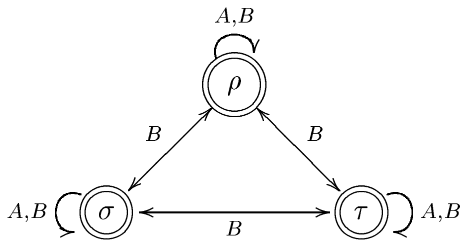

with pre(ρ) = **True**, pre(σ) = H, and pre(τ) = T. As illustrated, the set Γ is the entire set Σ_6 of simple actions. More generally, the general shape of $\mathbf{\Sigma}_6$ is the frame for an action which has three possibilities. A learns which one happens, and B merely learns that one of the three happens. Further, these two aspects are common knowledge. We omit the calculation that shows that $\mathbf{S}_1 \otimes \mathbf{\Sigma}_6$ is the model drawn in Scenario 6.

Example 3.10 The epistemic action model in Scenario 7 is $\mathbf{\Sigma}_7$, shown below:

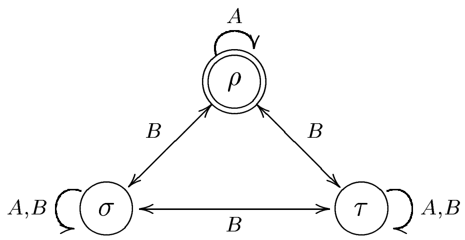

Again we take pre(ρ) = **True**, pre(σ) = H, and pre(τ) = T. Note that this figure has one fewer arrow than $\mathbf{\Sigma}_6$ above, since $\rho \xrightarrow{B} \rho$ is false. Another difference is that instead of $\Gamma = \Sigma_6$, we take Γ to be $\{s\}$. The idea is that B is sure that either t or u happens; that is, A learns either H or T. (We take $\Gamma = \{\rho\}$ to indicate that B is wrong about this.) Then $\mathbf{S}_1 \otimes \mathbf{\Sigma}_7$ is $\mathbf{S}_7$.

Example 3.11 Let φ represent "A knows H or A knows T." Let $\mathbf{\Sigma}_8$ be

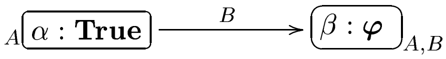

We take $\Gamma = \{\alpha\}$. The idea is that nothing actually happens, but B "learns" that A knows the state of the coin. Then $\mathbf{S}_7 \otimes \mathbf{\Sigma}_8$ is bisimilar to what we drew as $\mathbf{S}_8$ in Scenario 8.

3.6 Composition of epistemic action models

In all settings involving "actions" in some sense or other, composition is a natural operation. In our setting, we would like to define a composition operation on epistemic action models. Here is the relevant definition.

Let $\mathbf{\Sigma} = (\Sigma, \xrightarrow{A}, \text{pre}_\Sigma, \Gamma_{\mathbf{\Sigma}})$ and $\mathbf{\Delta} = (\Delta, \xrightarrow{A}, \text{pre}_\Delta, \Gamma_{\mathbf{\Delta}})$ be epistemic action models. We define the *composition*

$$\mathbf{\Sigma} \bullet \mathbf{\Delta} \quad = \quad (\Sigma \times \Delta, \xrightarrow{A}, \text{pre}_{\Sigma \bullet \Delta}, \Gamma_{\Sigma \bullet \Delta})$$

to be the following epistemic action model:

1. $\Sigma \times \Delta$ is the cartesian product of the sets Σ and Δ.
2. $\xrightarrow{A}$ in the composition $\mathbf{\Sigma} \bullet \mathbf{\Delta}$ is the family of product relations, in the natural way.
3. $\text{pre}_{\Sigma \bullet \Delta}(\sigma, \delta) = \langle(\mathbf{\Sigma}, \sigma)\rangle \text{pre}_\Delta(\delta)$.
4. $\Gamma_{\Sigma \bullet \Delta} = \Gamma_{\mathbf{\Sigma}} \times \Gamma_{\mathbf{\Sigma}}$.

In the definition of pre, $(\mathbf{\Sigma}, \sigma)$ is an abbreviation for the induced update $(\mathbf{\Sigma}, \{\sigma\})$ as defined in Section 3.4.

Example 3.12 We saw $\mathbf{\Sigma}_7$ and $\mathbf{\Sigma}_8$ in Examples 3.10 and 3.11. Then $\mathbf{\Sigma}_7 \bullet \mathbf{\Sigma}_8$ has six actions, since its action set is $\{\rho, \sigma, \tau\} \times \{\alpha, \beta\}$. Its set of distinguished actions is the singleton $\{(\rho, \alpha)\}$. Three of the six are reachable from (ρ, α), and we rename these as $(\rho, \alpha) \rightsquigarrow \gamma$, $(\sigma, \beta) \rightsquigarrow \delta$, and $(\tau, \beta) \rightsquigarrow \epsilon$. Then the picture is

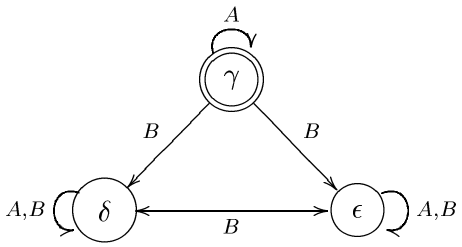

We have $\text{pre}(\gamma) = \langle(\mathbf{\Sigma}_7, \rho)\rangle\mathbf{True}$, $\text{pre}(\delta) = \langle(\mathbf{\Sigma}_7, \sigma)\rangle\boldsymbol{\varphi}$, and $\text{pre}(\epsilon) = \langle(\mathbf{\Sigma}_7, \tau)\rangle\boldsymbol{\varphi}$.

This is a representation of deception as a single action. We may return to $\mathbf{S}_1$ and apply this action. We get the following picture:

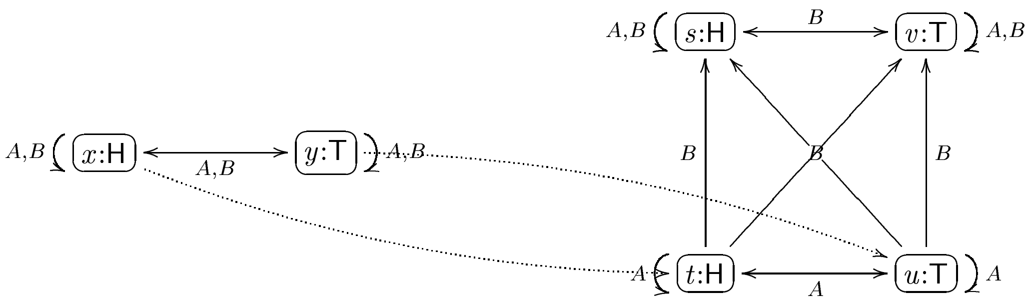

We conclude by verifying that our definition of composition is correct in the sense that the update induced by a composition is the composition of the induced updates. (The composition of updates was defined in Section 2.3, and again, induced updates are defined in Section 3.4.)

Proposition 3.13 $(\mathbf{\Sigma}, \Gamma_{\mathbf{\Sigma}}) \bullet (\mathbf{\Delta}, \Gamma_{\mathbf{\Delta}})$ *is equivalent to* $(\mathbf{\Sigma}, \Gamma_{\mathbf{\Sigma}}) \cdot (\mathbf{\Delta}, \Gamma_{\mathbf{\Delta}})$.

Proof We write $\mathbf{r}$ for $(\mathbf{\Sigma}, \Gamma_{\mathbf{\Sigma}})$, $\mathbf{s}$ for $(\mathbf{\Delta}, \Gamma_{\mathbf{\Delta}})$, and $\mathbf{t}$ for $(\mathbf{\Sigma}, \Gamma_{\mathbf{\Sigma}}) \bullet (\mathbf{\Delta}, \Gamma_{\mathbf{\Delta}})$. Let $\mathbf{S} = (S, \xrightarrow{A}_{\mathbf{S}}, [\![.]\!]_{\mathbf{S}})$ be a state model. Recall that

$$\mathbf{S}(\mathbf{r} \cdot \mathbf{s}) \quad = \quad \mathbf{S}(\mathbf{r})(\mathbf{s}) \quad = \quad (\mathbf{S} \otimes (\mathbf{\Sigma}, \Gamma_{\mathbf{\Sigma}})) \otimes (\mathbf{\Delta}, \Gamma_{\mathbf{\Delta}}).$$

We claim that this is isomorphic to $\mathbf{S} \otimes (\mathbf{\Sigma} \bullet \mathbf{\Delta}, \Gamma_{\mathbf{\Sigma} \bullet \mathbf{\Delta}})$, and indeed the isomorphism is $(s, (\sigma, \delta)) \mapsto ((s, \sigma), \delta)$. We check that $(s, (\sigma, \delta)) \in \mathbf{S} \otimes (\mathbf{\Sigma} \bullet \mathbf{\Delta})$ iff $((s, \sigma), \delta) \in (\mathbf{S} \otimes \mathbf{\Sigma}) \otimes \mathbf{\Delta}$. And the following are equivalent:

1. $(s, (\sigma, \delta)) \in \mathbf{S} \otimes (\mathbf{\Sigma} \bullet \mathbf{\Delta})$.
2. $s \in \|\mathrm{pre}_{\Sigma \bullet \Delta}(\sigma, \delta)\|_{\mathbf{S}}$.
3. $s \in \|\langle(\mathbf{\Sigma}, \sigma)\rangle \mathrm{pre}_{\Delta}(\delta)\|_{\mathbf{S}}$.
4. $(s, \sigma) \in \mathbf{S} \otimes \mathbf{\Sigma}$ and $(s, \sigma) \in \|\mathrm{pre}_{\mathbf{\Delta}}(\delta)\|_{\mathbf{S} \otimes \mathbf{\Sigma}}$.
5. $((s, \sigma), \delta) \in (\mathbf{S} \otimes \mathbf{\Sigma}) \otimes \mathbf{\Delta}$.

The rest of verification of isomorphism is fairly direct.

We also need to check that $\mathbf{t}_{\mathbf{S}}$ and $(\mathbf{r} \cdot \mathbf{s})_{\mathbf{S}}$ are related by the isomorphism. Now

$$\mathbf{t}_{\mathbf{S}} \quad = \quad \{(s, (s, (\sigma, \delta))) \in S \otimes (\Sigma \bullet \Delta) : \sigma \in \Gamma_{\mathbf{\Sigma}}, \delta \in \Gamma_{\mathbf{\Delta}}\}.$$

Recall that $(\mathbf{r} \cdot \mathbf{s})_{\mathbf{S}} = \mathbf{r}_{\mathbf{S}} \cdot \mathbf{s}_{\mathbf{S}(\mathbf{r})}$ and that this is a relational composition in left-to-right order. And indeed,

$$\begin{array}{lcl} \mathbf{r}_{\mathbf{S}} & = & \{(s, (s, \sigma)) : (s, \sigma) \in S \otimes \Sigma, \sigma \in \Gamma_{\mathbf{\Sigma}}\} \\ \mathbf{s}_{\mathbf{S}(\mathbf{r})} & = & \{((s, \sigma), ((s, \sigma), \delta)) \in \mathbf{S}(\mathbf{r}) \otimes \Delta : \delta \in \Gamma_{\mathbf{\Delta}}\}. \end{array}$$

This completes the proof. ⊣

Summary We have seen enough examples of the update product to make the case that complicated representations of epistemic situations, including ones which are not easy to construct by hand, may be obtained from simpler representations of initial situations using the update product with epistemic action models. The notion of an epistemic action model is closely related to the familiar multi-agent Kripke modes. Each epistemic action model induces an update in a natural way. And the collection of epistemic action models has a natural composition operation.

Endnotes The work of this section is intended to make good on our promise at the end of Section 1.1: we gave a principled explanation of how complex representations of naturally occurring scenarios may be computed from state models before the scenario and epistemic action models. There are precursors to our work in special cases, most notably Hans van Ditmarsch's dissertation [22]. That work is about examples like $\mathbf{\Sigma}_6$, where some agent or set of agents knows that the current action belongs to some set, but does not know which action it is. But to our knowledge, the work in this section is the first time that anything like this has been obtained in general.

We have taught the material in this section in several courses at different levels. Our experience is that the material here is highly appealing to students and makes the case for formal representations in the first place (for students not familiar with formal logic) and for the kind of technical work that we pursue in this area (for those who are).

This is an improved version of our earlier treatment of [2].

4 Logical languages based on action signatures

This section presents the centerpiece of our formal work, the logical systems corresponding to *types of epistemic action.* Our overall goal is to abstr from the propositional language earlier, in a way which allows us to fix only the *epistemic structure* of the desired actions, and vary their preconditions.

4.1 Action signatures

We have presented many examples of the update product in Section 3.5. These allow us to represent many natural updates, and this is one of our goals in this paper. But the structures which we have seen so far are not sufficient to get logical languages incorporating updates. For example, we have in Example 3.8 a program model that represents a private announcement to the agent A that a proposition H happens, and this takes place in a way that B learns nothing. The picture of this was

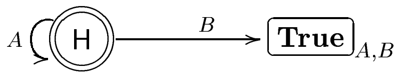

What we want to do now is to vary things a bit, and then to *abstract* them. For example, suppose we want to announce a different proposition to A, say φ. We would use

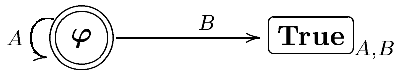

Or, we could announce φ to A in such a way that B is mislead into believing $\neg\varphi$ (and that everyone learns $\neg\varphi$). Now we use

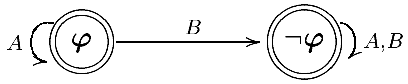

And the previous representations include the information that what *actually* happens is what A hears. To vary this, we need only change which world is designated by the doubled border. We could switch things, or double neither or both worlds. So at this point, we can see how to abstract things, leading to the following definition.

Definition An *action signature* is a structure

$$\mathbf{\Sigma} \quad = \quad (\Sigma, \xrightarrow{\mathcal{A}}, (\sigma_1, \sigma_2, \ldots, \sigma_n))$$

where $\mathbf{\Sigma} = (\Sigma, \xrightarrow{\mathcal{A}})$ is a finite Kripke frame, and $\sigma_1, \sigma_2, \ldots, \sigma_n$ is an enumeration of Σ in a list without repetitions. Just as with simple action structures (see Section 3.1), we call the elements of Σ *simple actions.*

The way this works for us is that an action signature together with an assignment of epistemic propositions to the simple actions in Σ will give us a full-fledged simple action structure. And this is the exact sense in which an action signature is an abstraction of the notion of a simple action structure. We shall shortly use action signatures in constructing logical languages.

Examples 4.1 Let $\mathbf{\Sigma}_{\text{pub}}$ be the following action signature. $\mathbf{\Sigma}_{\text{pub}}$ is intended to represent *public announcements*. It has $\Sigma = \{Pub\}$, and $Pub \xrightarrow{A} Pub$ for every agent A. So Pub is the type of action which is a public announcement. Note once again that we have not said *what* is being announced. We are purposely separating out the structure of the announcement from the content, and $\mathbf{\Sigma}_{Pub}$ is our model of the structure.

The next simplest action signature is the "test" signature $\mathbf{\Sigma}_?$. We take $\mathbf{\Sigma}_? = \{?, skip\}$, with the enumeration $?, skip$. We also take $? \xrightarrow{A} skip$, and $skip \xrightarrow{A} skip$ for all A. This turns out to be a totally opaque form of test: φ is tested on the real world, but nobody knows this is happening. Its function will be to generate tests $?\varphi$, which affect the states precisely in the way dynamic logic tests do.

For each set $\mathcal{B} \subseteq \mathcal{A}$ of agents, we define the action signature $Pri^{\mathcal{B}}$ of *completely private announcements to the group* $\mathcal{B}$. It has $\Sigma = \{Pri^{\mathcal{B}}, skip\}$; $Pri^{\mathcal{B}} \xrightarrow{B} Pri^{\mathcal{B}}$ for all $B \in \mathcal{B}$, $Pri^{\mathcal{B}} \xrightarrow{C} skip$ for $C \notin \mathcal{B}$, and $skip \xrightarrow{A} skip$ for all agents A.

The action signature $Cka_k^{\mathcal{B}}$ is given by: $\Sigma = \{1, \ldots, k\}$; $i \xrightarrow{B} i$ for $i \leq k$ and $B \in \mathcal{B}$; and finally $i \xrightarrow{C} j$ for $i, j \leq k$ and $C \notin \mathcal{B}$. This action signature is called the signature of *common knowledge of alternatives for an announcement to the group* $\mathcal{B}$.

Signature-based program models Now that we have a general abstract notion, we introduce some notation to regain the earlier examples. Let $\mathbf{\Sigma}$ be a action signature, let $\Gamma \subseteq \Sigma$, and let $\vec{\boldsymbol{\psi}} = \boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_n$ be a list of epistemic propositions. We obtain an epistemic program model $(\mathbf{\Sigma}, \Gamma)(\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_n)$ in the following way:

1. The set of actions is Σ, and the accessibility relations are those given by the action signature.

2. For $j = 1, \ldots, n$, $\text{pre}(\sigma_j) = \boldsymbol{\psi}_j$.

3. The set of distinguished actions is Γ.

In the special case that Γ is the singleton set $\{\sigma_i\}$, we write the resulting signature-based program model as $(\mathbf{\Sigma}, \sigma)(\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_n)$.

To summarize: every action signature, set of simple actions in it, and tuple of epistemic propositions gives an epistemic program model in a canonical way.

4.2 The syntax and semantics of $\mathcal{L}(\mathbf{\Sigma})$

Fix a action signature $\mathbf{\Sigma}$. Let n be the number of actions in Σ. We present in Figure 8 a logic $\mathcal{L}(\mathbf{\Sigma})$. Actually, we have three logics there, the full logic $\mathcal{L}(\mathbf{\Sigma})$ and the smaller fragments $\mathcal{L}_0(\mathbf{\Sigma})$ and $\mathcal{L}_1(\mathbf{\Sigma})$.

The first thing to notice about this is that for the first time in a great while, we have a genuine syntax. In this regard, note that in the programs of the form $\sigma\psi_1, \ldots, \psi_n$ we have *sentences* $\vec{\psi}$ rather than *epistemic propositions* (which we had written using boldface letters $\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_n$ in Section 2.3). Also, the signature $\mathbf{\Sigma}$ figures into the semantics exactly in those programs $\sigma\psi_1, \ldots, \psi_n$; in those we require that $\sigma \in \Sigma$.

The second thing to note is that we have two sorts of syntactic object: sentence and programs. We call programs of the form $\sigma\psi_1, \ldots, \psi_n$ *simple programs*. Note that they might not be "atomic" in the sense that the sentences ψ_j might themselves contain actions.

sentences φ	True	p_i	$\neg\varphi$	$\varphi \wedge \psi$	$\Box_A \varphi$	$\Box^*_{\mathcal{B}} \varphi$	$[\pi]\varphi$
programs π	$skip$	$crash$	$\sigma\psi_1, \ldots, \psi_n$	$\pi + \rho$	$\pi \cdot \rho$	π^*	

$$
\begin{array}{lcllcl}
[\![\mathsf{True}]\!] & = & \mathbf{True} & [\![skip]\!] & = & \mathbf{1} \\
[\![p]\!] & = & \boldsymbol{p} & [\![crash]\!] & = & \mathbf{0} \\
[\![\varphi \wedge \psi]\!] & = & [\![\varphi]\!] \wedge [\![\psi]\!] & [\![\sigma_i \psi_1 \ldots \psi_n]\!] & = & (\mathbf{\Sigma}, \sigma_i)([\![\psi_1]\!], \ldots, [\![\psi_n]\!]) \\
[\![\neg\varphi]\!] & = & \neg[\![\varphi]\!] & [\![\pi \cdot \rho]\!] & = & [\![\pi]\!] \cdot [\![\rho]\!] \\
[\![\Box_A \varphi]\!] & = & \Box_A [\![\varphi]\!] & [\![\pi + \rho]\!] & = & [\![\pi]\!] + [\![\rho]\!] \\
[\![\Box^*_{\mathcal{B}} \varphi]\!] & = & \Box^*_{\mathcal{B}} [\![\varphi]\!] & [\![\pi^*]\!] & = & [\![\pi]\!]^* \\
[\![[\pi]\varphi]\!] & = & [[\![\pi]\!]][\![\varphi]\!] & & &
\end{array}
$$

For $\mathcal{L}_1(\mathbf{\Sigma})$, we drop the π^* construct. For $\mathcal{L}_0(\mathbf{\Sigma})$, we drop the π^* and $\Box^*_{\mathcal{B}}$ constructs. In both of these restricted contexts, we usually refer to the programs as *actions* and use letters like α, β, etc.

Figure 8: The language $\mathcal{L}(\mathbf{\Sigma})$ and its semantics

The semantics defines two maps by simultaneous recursion on $\mathcal{L}(\mathbf{\Sigma})$:

1. a map $\varphi \mapsto [\![\varphi]\!]$, taking the sentences of $\mathcal{L}(\mathbf{\Sigma})$ into epistemic propositions; and
2. a map $\pi \mapsto [\![\pi]\!]$, taking the programs of $\mathcal{L}(\mathbf{\Sigma})$ into updates.

The formal definition is given in Figure 8. The main thing to note is that with one key exception, the operations on the right-hand sides are immediate applications of our general definitions of the closure conditions on epistemic propositions and updates which we noted back in Section 2.3. A good example to explain this is the clause for the semantics of sentences $[\alpha]\varphi$. Assuming that we have an epistemic program $[\![\alpha]\!]$ and an epistemic proposition $[\![\varphi]\!]$, we can form the epistemic proposition $[[\![\alpha]\!]][\![\varphi]\!]$ (see equation (4) in Section 2.3). Note that we have overloaded the square bracket notation; this is intentional, and we have done the same with other notation as well.

We also discuss the definition of the semantics for simple actions $\sigma_i \vec{\psi}$. This is precisely where the structure of the action signature $\mathbf{\Sigma}$ enters. For this, recall that we have general definition of a signature-based epistemic program model $(\mathbf{\Sigma}, \Gamma)(\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_n)$, where $\Gamma \subseteq \Sigma$ and the $\boldsymbol{\psi}$'s are any epistemic propositions. What we have in the semantics of $\sigma_i \vec{\psi}$ is the special case of this where Γ is the singleton $\{\sigma_i\}$ and $\boldsymbol{\psi}_i$ is $[\![\psi_i]\!]$, a proposition which we already have when we come to define $[\![\sigma_i \vec{\psi}]\!]$.

At this point, it is probably good to go back to our earlier discussions in Section 2.1 of epistemic propositions and updates. What we have done overall is to give a fully syntactic presentation of languages of these epistemic propositions and updates. The constructions of the language correspond to the closure properties noted in Section 2.1. (To be sure, we have restricted to the finite case at several points because we are interested in a syntax, and at the same time we have re-introduced some infinitary aspects via the Kleene star.)

Preservation of bisimulation and atomic propositions We note two basic facts about the interpretations of sentences and programs in the logics $\mathcal{L}(\mathbf{\Sigma})$.

Proposition 4.2 *The interpretations of the sentences and programs of $\mathcal{L}(\mathbf{\Sigma})$ preserve bisimulation.*

Proof By induction on $\mathcal{L}(\mathbf{\Sigma})$, using Propositions 2.5 and 3.3. $\dashv$

Proposition 4.3 *The interpretations of programs of $\mathcal{L}(\mathbf{\Sigma})$ preserve atomic sentences in the sense that if s $[\![\pi]\!]_{\mathbf{S}}$ t, then for all atomic sentences p, $s \in \|p\|_{\mathbf{S}}$ iff $t \in \|p\|_{\mathbf{S}([\![\pi]\!])}$.*

Proof By induction on π. $\dashv$

4.3 Formalization of the target logics as fragments

We shall formalize the target logics of Section 2.4 in terms of appropriate action signatures $\mathbf{\Sigma}$. Our goal is that each of the target logics may be thought of as $\mathcal{L}(\mathbf{\Sigma})$ for some $\mathbf{\Sigma}$. As we shall see, this is not exactly possible, because the languages $\mathcal{L}(\mathbf{\Sigma})$ often give us sentences that are not in the target logics. However, we introduce the notion of a *fragment* in order to close that gap.

We shall use the action signatures of Examples 4.1 and the notation from there.

The logic of public announcements This is formalized as $\mathcal{L}(\mathbf{\Sigma}_{Pub})$. We have sentences $[Pub\ \varphi]\psi$, just as we described in Section 2.4. Note that $\mathcal{L}(\mathbf{\Sigma}_{Pub})$ allows announcements inside announcements. If φ, ψ, and χ are sentences, then so is $\langle Pub\ [Pub\ \varphi]\psi\rangle\chi$.

We check that $\mathcal{L}(\mathbf{\Sigma}_{Pub})$ is a good formalization of the logic of public announcements. Fix a sentence φ of $\mathcal{L}(\mathbf{\Sigma}_{Pub})$ and a state model $\mathbf{S}$. We calculate:

$$\begin{array}{rcl} \mathbf{S}([\![Pub\ \varphi]\!]) & = & \mathbf{S}((\mathbf{\Sigma}_{Pub}, Pub)([\![\varphi]\!])) \\ & = & \mathbf{S} \otimes ((\mathbf{\Sigma}_{Pub}, Pub)([\![\varphi]\!])) \\ & = & \{(s, Pub) : s \in [\![\varphi]\!]_{\mathbf{S}}\} \end{array}$$

The state model has the structure given earlier in terms of the update product operation. The update relation $[\![Pub\ \varphi]\!]_{\mathbf{S}}$ relates s to (s, Pub) whenever the latter belongs to the state model $\mathbf{S}([\![Pub\ \varphi]\!])$. The model itself is isomorphic to the sub-state-model of $\mathbf{S}$ induced by $\{s \in S : s \in [\![\varphi]\!]_{\mathbf{S}}\}$. Under this isomorphism, the update relation is then the inverse of inclusion. This is just how the action of public announcement was described when we first encountered it, in Example 2.1 of Section 2.3.

Test-only PDL was introduced in Section 2.8. Although it was not one of the target languages of Section 2.4, it will be instructive to see how it is formalized in our setting. Recall that test-only PDL has actions of the form $?\varphi$. We want to use our action signature $\mathbf{\Sigma}_?$. The simple actions of it are $?$ and *skip*. So technically we have sentences of the following forms:

$$[?\ \varphi, \psi]\chi \quad \text{and } [skip\ \varphi, \psi]\chi \tag{12}$$

(We also have sentences $[skip]\chi$ coming directly from the syntax of $\mathcal{L}(\mathbf{\Sigma}_?)$. Hopefully, the reader will not be confused at this point.) So in a sense, our syntax has overgenerated. We get around this by restricting attention to those sentences in (12) where ψ is True. We shall make this clearer in our notion of a *fragment* below. However, at this point, we can study the semantics of the basic actions $[?\ \varphi, \mathsf{True}]$ and $[skip\ \varphi, \mathsf{True}]$. Fix a state model $\mathbf{S}$. We calculate:

$$
\begin{array}{rcl}
\mathbf{S}([\![?\ \varphi, \mathsf{True}]\!]) & = & \mathbf{S}((\mathbf{\Sigma}_?, ?)([\![\varphi]\!], \mathbf{True})) \\
& = & \mathbf{S} \otimes ((\mathbf{\Sigma}_{Pub}, Pub)([\![\varphi]\!], \mathbf{True})) \\
& = & \{(s, ?) : s \in [\![\varphi]\!]_{\mathbf{S}}\} \cup \{(s, skip) : s \in S\}
\end{array}
$$

The structure of $\mathbf{S}([\![?\ \varphi, \mathsf{True}]\!])$ is that for each A. $(s, ?) \xrightarrow{A} (t, skip)$ provided $s \xrightarrow{A} t$ in $\mathbf{S}$. Also $(s, skip) \xrightarrow{A} (t, skip)$ under the same condition, and there are no other arrows. The update relation $[\![?\ \varphi, \mathsf{True}]\!]_{\mathbf{S}}$ relates s to $(s, ?)$ whenever the latter belongs to the updated structure. Overall, this update is isomorphic to what we described in Example 2.2 of Section 2.3.

Turning to the update map of $[skip\ \varphi, \mathsf{True}]$, we again fix $\mathbf{S}$. The model $\mathbf{S}([\![skip\ \varphi, \mathsf{True}]\!])$ is again literally the same as what we calculated above. However, the update relation $[\![skip\ \varphi, \mathsf{True}]\!]_{\mathbf{S}}$ now relates each $s \in S$ to the pair $(s, skip)$. This relation is a bisimulation. We shall formulate a notion of action equivalence later, and it will turn out that the update $[\![skip\ \varphi, \mathsf{True}]\!]$ is equivalent to $\mathbf{1}$. For now, we can also consider the semantics of sentences of the form $[skip\ \varphi, \mathsf{True}]\psi$. We have

$$
\begin{array}{rcl}
[\![[skip\ \varphi, \mathsf{True}]\psi]\!]_{\mathbf{S}} & = & \{s \in S : \text{if } s\ [\![skip\ \varphi, \mathsf{True}]\!]_{\mathbf{S}}\ t, \text{ then } t \in [\![\psi]\!]_{\mathbf{S}([\![skip\ \varphi, \mathsf{True}]\!])}\} \\
& = & \{s \in S : (s, skip) \in [\![\psi]\!]_{\mathbf{S}([\![skip\ \varphi, \mathsf{True}]\!])}\} \\
& = & \{s \in S : s, \in [\![\psi]\!]_{\mathbf{S}}\}
\end{array}
$$

This last equality is by Proposition 4.2 on bisimulation preservation. The upshot is that $[\![[skip\ \varphi, \mathsf{True}]\psi]\!]_{\mathbf{S}} = [\![\psi]\!]_{\mathbf{S}}$. So for this reason, we might as well identify the sentences $[skip\ \varphi, \mathsf{True}]\psi$ and ψ. Or to put things differently, we might as well identify the basic action $skip\ \varphi, \mathsf{True}$ and (the constant of $\mathcal{L}(\mathbf{\Sigma}_?)$) *skip*. Since we are modifying the syntax, we might also abbreviate $[?\ \varphi, \mathsf{True}]$ to $[?\ \varphi]$. Doing all this leads to a language which is exactly test-only PDL as we had it in Section 2.8, and our semantics there agrees with what we have calculated in the present discussion.

At this point we know that test-free PDL corresponds to a subset of the language $\mathcal{L}(\mathbf{\Sigma}_?)$. The latter language is larger, and indeed sentences like $[?\ p, q]\Box_A^* r$ are probably not expressible in test-only PDL. (We have not checked this.) This leads to the following definition:

Definition A *fragment* of $\mathcal{L}(\mathbf{\Sigma})$ is a set F of sentences and programs such that the following hold:

1. F is closed under subsentences and subprograms.
2. F is closed under the sentence forming operations of $\neg$, $\wedge$, $\Box_A$, $\Box_{\mathcal{B}}^*$, and $[\pi]\varphi$.
3. F is closed under the program forming operations of *skip*, $+$, $\cdot$, and iteration π^*.
4. If $\sigma\psi_1, \ldots, \psi_n$ belongs to F and $\sigma \xrightarrow{A} \tau$ for some A, then also $\tau\psi_1, \ldots, \psi_n$ belongs to F.

A fragment of $\mathcal{L}_1(\mathbf{\Sigma})$ is defined in the same way, except that we do not require closure under the program iteration operation π^*.

So we see that test-only PDL may be formalized as a fragment of $\mathcal{L}(\mathbf{\Sigma}_?)$. It is the fragment F where in each basic action $?\ \varphi, \psi$ we have $\psi = \mathsf{True}$.

The logic of totally private announcements will be formalized as a fragment of $\mathcal{L}_1(\mathbf{\Sigma})$, where $\mathbf{\Sigma}$ is the disjoint union of the action signatures $Pri^{\mathcal{B}}$ for $\emptyset \neq \mathcal{B} \subseteq \mathcal{A}$. (These action signatures $Pri^{\mathcal{B}}$ were defined in Section 4.1.) For example, in the case of $\mathcal{A} = \{A, B\}$, $\mathbf{\Sigma}$ would have $2 \times (2^2 - 1) = 6$ actions. We might write these as $Pri^A, Pri^B, Pri^{A,B}, skip^A, skip^B, skip^{A,B}$. Let us assume that $\mathbf{\Sigma}$ includes the information that this is the order of these actions. The fragment of interest is the one where in basic actions $\sigma\ \psi_1, \ldots, \psi_6$, we have the following: (1) if σ is Pri^A or $skip^A$, then all ψ except ψ_1 are True; (2) if σ is Pri^B or $skip^B$, then all ψ except ψ_2 are True; (3) if σ is $Pri^{A,B}$ or $skip^{A,B}$, then all ψ except ψ_3 are True. When dealing with this fragment, it is sensible to adopt abbreviations just as we did with the formalization of the test-only fragment of PDL. We would abbreviate, for example,

$$[Pri^B\ \mathsf{True}, \varphi, \mathsf{True}, \mathsf{True}, \mathsf{True}, \mathsf{True}]\psi$$

as $[Pri^B\ \varphi]\psi$. We also would abbreviate any action $skip^X \vec{\psi}$ to just $skip$, for the same reasons we saw with $\mathbf{\Sigma}_?$.

Combining logics Given two or more action signatures and their associated languages, it is natural to consider the *combined language* where we have action construction operations $\sigma\vec{\psi}$ for all σ in any of the given signatures. We may use the notion of a fragment to study this. Specifically, let $\mathbf{\Sigma}$ be the disjoint union of the given signatures. Then the language of interest is a fragment of $\mathcal{L}(\mathbf{\Sigma})$.

The logic of common knowledge of alternatives uses the disjoint union over all $\mathcal{B}$ and k of the action signatures $Cka^{\mathcal{B}}_k$.

4.4 Other logics

Logics based on frame conditions Many other types of logics are easily representable in our framework. For example, consider the logic of all S4 programs. To formalize this, we need only consider the disjoint union of all finite action signatures whose underlying accessibility relations are preorders. For the logic of S5 programs, we change preorders to equivalence relations. Another important class is given by the KD45 axioms of modal logic. These systems are particularly important because the KD45 and S5 conditions are so prominent in epistemic logic.

Announcements by Particular Agents Our modeling of the notion of the *public announcement that* φ is impersonal in the sense that the announcement does not come from anywhere in particular. It might be best understood as coming externally, as if someone shouted φ into the room where the agents were standing.

We also want a notion of the *public announcement by* A *of* φ. We shall write this as $Pub^A\varphi$. For this, we identify $Pub^A\varphi$ with the (externally-made) public announcement that φ holds and A knows this. This identification does not represent the fact that A intends to inform the others of φ. But as we know, intentions are not modeled at all in our system. We claim, however, that on a purely propositional level, the identification works. And using it, we can represent announcements by A. One way to do this is via abbreviation: we take $\langle Pub^A\ \varphi\rangle\psi$ to be an abbreviation for $\langle Pub\ \varphi \wedge \Box_A\varphi\rangle\psi$. (A different way to formalize $Pub^A\ \varphi$ would be to use a special signature constructed just for that purpose. But the discussion here shows that there is no need to do this. One can use Σ_{Pub}.)

Lying We can also represent misleading epistemic actions, such as lying. Again, we want to understand the action of A *lying that* φ externally. To all other announcements, this program should be the same thing as an announcement of φ by A. But to say that A *lies* about φ, we want to assume that φ is actually false. Further, we want to assume that A moreover knows that φ is false. (For if φ just happened to be false by A said φ, we would not really want to call that "lying.")

The technical details on the representation go as follows. We take a signature $\Sigma_{Pubcond}$ given as follows.

$$\Sigma_{Pubcond} \quad = \quad \{Pubcond^A : A \in \mathcal{A}\} \cup \{Pub^A : A \in \mathcal{A}\}.$$

($Pubcond$ stands for "conditional public announcement.") We take $\text{pre}(Pubcond^A) = 2$ and $\text{pre}(Pub^A) = 1$ for all A. Further, we take the following, for all A and B: $Pubcond^A \xrightarrow{A} Pubcond^A$, $Pub^A \xrightarrow{B} Pub^A$; and for $B \neq A$, $Pubcond^A \xrightarrow{B} Pub^A$. So we have sentences like $[Pubcond\ \varphi, \psi]\chi$. The extra argument ψ is a kind of secret condition. And we can use $[Lie^A\varphi]\chi$ as an abbreviation of

$$[Pubcond^A\ \varphi \wedge \Box_A\varphi, \neg\varphi \wedge \Box_A\neg\varphi]\chi.$$

That is, for A to lie about φ there is a condition that $\neg\varphi \wedge \Box_A\neg\varphi$. But the other agents neither need to know this ahead of time nor do they in any sense "learn" this from the announcement. Indeed, for the other agents, $Lie^A\varphi$ is just like a truthful public announcement by A.

Gratuitous Suspicion We can also represent gratuitous suspicion ("paranoia"): maybe no "real" action has taken place, except that some people start suspecting some action (e.g., some private communication) has taken place. **from Larry: we need more here. My feeling is that there might be a better account of lying based on what is going on in Scenarios 7 and 8.**

Endnote All of the work in this section is new here.

5 Logical systems

With all logical languages and their interpretations, it is a matter of central interest to obtain a *sound and complete* proof system for the validities. We are of course interested in the languages $\mathcal{L}(\mathbf{\Sigma})$ and their sublanguages $\mathcal{L}_0(\mathbf{\Sigma})$ and $\mathcal{L}_1(\mathbf{\Sigma})$. We write $\models \varphi$ to mean that for all state models $\mathbf{S}$ and all $s \in S$, $s \in [\![\varphi]\!]_{\mathbf{S}}$.

In Figure 9 below we present a logic for $\mathcal{L}(\mathbf{\Sigma})$. We write $\vdash \varphi$ if φ can be obtained from the axioms of the system using its inference rules. We often omit the turnstile $\vdash$ when it is clear from the context.

Most of the system will be quite standard from modal logic. The Action Axioms are new, however. These include the in the Atomic Permanence axiom; note that in this axiom p is an atomic sentence. The axiom says that announcements do not change the brute fact of whether or not p holds. This axiom reflects the fact that our actions do not change any kind of local state. (We discuss an extension of our system in Section 7.6 where this axiom is not sound.) Another axiom in this group is the Action-Knowledge Axiom, giving a criterion for knowledge after an announcement.

As it happens, the logic without the $**$ axioms is complete for $\mathcal{L}_1(\mathbf{\Sigma})$. Dropping both the $*$ and $**$ axioms gives a system complete for $\mathcal{L}_0(\mathbf{\Sigma})$. Our completeness results appear in Section 7.

In this section, we prove the soundness of all the axioms and rules of the system except for the Action Rule. We delay discussion of the Action Rule mostly because it takes a fair amount or work to prove the soundness. Also, delaying things emphasizes the significance of adding the infinitary operators $\Box^*_{\mathcal{B}}$ to $\mathcal{L}_0(\mathbf{\Sigma})$ and then the program repetition construct π^* to this.

5.1 Soundness, except for the Action Rule

Proposition 5.1 *The Atomic Permanence Axiom* $[\sigma_i\vec{\psi}]p \leftrightarrow (\psi_i \to p)$ *is sound.*

Proof Recall how $[\![\sigma_i\vec{\psi}]\!]$ works; call this update $\mathbf{r}$. Fix a state model $\mathbf{S}$. The following are equivalent:

1. $s \in [\![[\sigma_i\vec{\psi}]p]\!]_{\mathbf{S}}$.
2. $s \in ([\mathbf{r}]\boldsymbol{p})_{\mathbf{S}}$.
3. If $s\ \mathbf{r}_{\mathbf{S}}\ t$, then $t \in \boldsymbol{p}_{\mathbf{S}(\mathbf{r})}$.
4. If $(s, \sigma_i) \in \mathbf{S}(\mathbf{r})$, then $(s, \sigma_i) \in \boldsymbol{p}_{\mathbf{S}(\mathbf{r})}$.
5. If $s \in [\![\psi_i]\!]_{\mathbf{S}}$, then $s \in \|p\|_{\mathbf{S}} = \boldsymbol{p}_{\mathbf{S}}$.
6. $s \in [\![\psi_i \to p]\!]_{\mathbf{S}}$.

$\dashv$

	Basic Axioms	
	All sentential validities	
	($[\pi]$-normality)	$\vdash [\pi](\varphi \to \psi) \to ([\pi]\varphi \to [\pi]\psi)$
	($\Box_A$-normality)	$\vdash \Box_A(\varphi \to \psi) \to (\Box_A\varphi \to \Box_A\psi)$
$*$	($\Box^*_{\mathcal{C}}$-normality)	$\vdash \Box^*_{\mathcal{C}}(\varphi \to \psi) \to (\Box^*_{\mathcal{C}}\varphi \to \Box^*_{\mathcal{C}}\psi)$
	Action Axioms	
	(Atomic Permanence)	$\vdash [\sigma_i\vec{\psi}]p \leftrightarrow (\psi_i \to p)$
	(Partial Functionality)	$\vdash [\sigma_i\vec{\psi}]\neg\chi \leftrightarrow (\psi_i \to \neg[\sigma_i\vec{\psi}]\chi)$
	(Action-Knowledge)	$\vdash [\sigma_i\vec{\psi}]\Box_A\varphi \leftrightarrow (\psi_i \to \bigwedge\{\Box_A[\sigma_j\vec{\psi}]\varphi : \sigma_i \xrightarrow{A} \sigma_j \text{ in } \mathbf{\Sigma}\})$
$**$	**Action Mix Axiom**	$\vdash [\pi^*]\varphi \to \varphi \wedge [\pi][\pi^*]\varphi$
$*$	**Epistemic Mix Axiom**	$\vdash \Box^*_{\mathcal{C}}\varphi \to \varphi \wedge \bigwedge\{\Box_A\Box^*_{\mathcal{C}}\varphi : A \in \mathcal{C}\}$
	Skip Axiom	$\vdash [\mathit{skip}]\varphi \leftrightarrow \varphi$
	Crash Axiom	$\vdash [\mathit{crash}]\mathsf{False}$
	Composition Axiom	$\vdash [\pi][\rho]\varphi \leftrightarrow [\pi \cdot \rho]\varphi$
	Choice Axiom	$\vdash [\pi + \rho]\varphi \leftrightarrow [\pi]\varphi \wedge [\rho]\varphi$

	Modal Rules	
	(Modus Ponens)	From $\vdash \varphi$ and $\vdash \varphi \to \psi$, infer $\vdash \psi$
	($[\pi]$-necessitation)	From $\vdash \psi$, infer $\vdash [\pi]\psi$
	($\Box_A$-necessitation)	From $\vdash \varphi$, infer $\vdash \Box_A\varphi$
$*$	($\Box^*_{\mathcal{C}}$-necessitation)	From $\vdash \varphi$, infer $\vdash \Box^*_{\mathcal{C}}\varphi$
$**$	**Program Induction Rule**	From $\vdash \chi \to \psi \wedge [\pi]\chi$, infer $\vdash \chi \to [\pi^*]\psi$
$*$	**Action Rule**	

Let ψ be a sentence, let α be an action, and let $\mathcal{C}$ be a set of agents. Let there be sentences χ_β for all β such that $\alpha \to^*_{\mathcal{C}} \beta$ (including α itself), and such that

1. $\vdash \chi_\beta \to [\beta]\psi$.
2. If $A \in \mathcal{C}$ and $\beta \xrightarrow{A} \gamma$, then $\vdash (\chi_\beta \wedge \textsc{Pre}(\beta)) \to \Box_A\chi_\gamma$.

From these assumptions, infer $\vdash \chi_\alpha \to [\alpha]\Box^*_{\mathcal{C}}\psi$.

Figure 9: The logical system for $\mathcal{L}(\mathbf{\Sigma})$. For $\mathcal{L}_1(\mathbf{\Sigma})$, we drop the $**$ axioms and rule; for $\mathcal{L}_0(\mathbf{\Sigma})$, we also drop the $*$ axioms and rules.

Proposition 5.2 *The Partial Functionality Axiom* $[\sigma_i\vec{\psi}]\neg\chi \leftrightarrow (\psi_i \to \neg[\sigma_i\vec{\psi}]\chi)$ *is sound.*

Proof Again, let $\mathbf{r} = [\![\sigma_i\vec{\psi}]\!]$. Also, let $\boldsymbol{\chi} = [\![\chi]\!]$. Fix a state model $\mathbf{S}$. The following are equivalent:

1. $s \in [\![[\sigma_i\vec{\psi}]\neg\chi]\!]_{\mathbf{S}}$.
2. $s \in ([\mathbf{r}]\neg\boldsymbol{\chi})_{\mathbf{S}}$.
3. If $s\ \mathbf{r}_{\mathbf{S}}\ t$, then $t \in \neg\boldsymbol{\chi}_{\mathbf{S}(\mathbf{r})}$.
4. If $s\ \mathbf{r}_{\mathbf{S}}\ t$, then $t \notin \boldsymbol{\chi}_{\mathbf{S}(\mathbf{r})}$.
5. If $s \in [\![\psi_i]\!]_{\mathbf{S}}$, then $(s, \sigma_i) \notin \boldsymbol{\chi}_{\mathbf{S}(\mathbf{r})}$.
6. If $s \in [\![\psi_i]\!]_{\mathbf{S}}$, then $(s, \sigma_i) \in \neg\boldsymbol{\chi}_{\mathbf{S}(\mathbf{r})}$.
7. $s \in [\![\psi_i \to \neg[\sigma_i\vec{\psi}]\chi]\!]_{\mathbf{S}}$.

The crucial equivalence here is (3)$\Longleftrightarrow$(4). The reason this holds is that in the action model for $\sigma_i\vec{\psi}$, there is just one distinguished world. So for each fixed $s \in S$, there is at most one t such that $s\ \mathbf{r}_{\mathbf{S}}\ t$. $\dashv$

Proposition 5.3 *The Action-Knowledge Axiom*

$$[\sigma_i\vec{\psi}]\Box_A\varphi \leftrightarrow (\psi_i \to \bigwedge\{\Box_A[\sigma_j\vec{\psi}]\varphi : \sigma_i \xrightarrow{A} \sigma_j \ \textit{in}\ \mathbf{\Sigma}\})$$

is sound.

Proof Again let $\mathbf{r} = [\![\sigma_i\vec{\psi}]\!]$ (see Section 4.1). Fix a state model $\mathbf{S}$. The following are equivalent:

1. $s \in [\![[\sigma_i\vec{\psi}]\Box_A\varphi]\!]_{\mathbf{S}}$.
2. $s \in [\mathbf{r}][\![\Box_A\varphi]\!]_{\mathbf{S}}$.
3. If $s\ \mathbf{r}_{\mathbf{S}}\ t$, then $t \in [\![\Box_A\varphi]\!]_{\mathbf{S}}$.
4. If $s \in [\![\psi_i]\!]_{\mathbf{S}}$, then $(s, \sigma_i) \in [\![\Box_A\varphi]\!]_{\mathbf{S}}$.
5. If $(s, \sigma_i) \in \mathbf{S}(\mathbf{r})$, then for all (t, σ_j) such that $(t, \sigma_j) \in \mathbf{S}(\mathbf{r})$, $s \xrightarrow{A} t$ and $\sigma_i \xrightarrow{A} \sigma_j$, we have $(t, \sigma_j) \in [\![\varphi]\!]_{\mathbf{S}}$.
6. If $(s, \sigma_i) \in \mathbf{S}(\mathbf{r})$, then for all t such that $s \xrightarrow{A} t$, if $\sigma_i \xrightarrow{A} \sigma_j$ and $(t, \sigma_j) \in \mathbf{S}(\mathbf{r})$, then $(t, \sigma_j) \in [\![\varphi]\!]_{\mathbf{S}}$.
7. If $(s, \sigma_i) \in \mathbf{S}(\mathbf{r})$, then for all t such that $s \xrightarrow{A} t$ and $\sigma_i \xrightarrow{A} \sigma_j$: if $(t, \sigma_j) \in \mathbf{S}(\mathbf{r})$, then $(t, \sigma_j) \in [\![\varphi]\!]_{\mathbf{S}}$.
8. If $(s, \sigma_i) \in \mathbf{S}(\mathbf{r})$, then for all t and σ_j such that $s \xrightarrow{A} t$ and $\sigma_i \xrightarrow{A} \sigma_j$: $t \in [\![[\sigma_j\vec{\psi}]\varphi]\!]_{\mathbf{S}}$.
9. If $s \in [\![\psi_i]\!]_{\mathbf{S}}$, then for all t such that $s \xrightarrow{A} t$, $t \in [\![\bigwedge\{[\sigma_j\vec{\psi}]\varphi : \sigma_i \xrightarrow{A} \sigma_j\}]\!]_{\mathbf{S}}$.
10. $s \in [\![\psi_i \to \bigwedge\{\Box_A[\sigma_j\vec{\psi}]\varphi : \sigma_i \xrightarrow{A} \sigma_j\}]\!]_{\mathbf{S}}$.

$\dashv$

Proposition 5.4 *The Composition Axiom* $[\pi \cdot \rho]\varphi \leftrightarrow [\pi][\rho]\varphi$ *is sound.*

Proof Write $\mathbf{r}$ for $[\![\pi]\!]$ and $\mathbf{s}$ for $[\![\rho]\!]$. Fix a state model $\mathbf{S}$. The following are equivalent:

1. $s \in [\![[\pi \cdot \rho]\varphi]\!]_{\mathbf{S}}$.
2. If $s\ [\![\pi \cdot \rho]\!]_{\mathbf{S}}\ u$, then $u \in [\![\varphi]\!]_{\mathbf{S}(\mathbf{r}\cdot\mathbf{s})}$.
3. If $s\ (\mathbf{r} \cdot \mathbf{s})_{\mathbf{S}}\ u$, then $u \in [\![\varphi]\!]_{\mathbf{S}(\mathbf{r}\cdot\mathbf{s})}$.
4. If $s\ \mathbf{r}_{\mathbf{S}}\ t$ and $t\ \mathbf{s}_{\mathbf{S}(\mathbf{r})}\ u$, then $u \in [\![\varphi]\!]_{\mathbf{S}(\mathbf{r}\cdot\mathbf{s})}$.
5. If $s\ \mathbf{r}_{\mathbf{S}}\ t$, then $t \in [\![[\rho]\varphi]\!]_{\mathbf{S}(\mathbf{r})}$.
6. $s \in [\![[\pi][\rho]\varphi]\!]_{\mathbf{S}}$.

The equivalence of (3) and (4) is by the definition of relational composition. The remaining equivalences are from the semantic definitions. The equivalence of (4) and (5) is by the fact that $\mathbf{S}(\mathbf{r} \cdot \mathbf{s}) = \mathbf{S}(\mathbf{r})(\mathbf{s})$. $\dashv$

Proposition 5.5 *The Skip Axiom* $[skip]\varphi \leftrightarrow \varphi$ *is sound.*

Proof Fix a state model $\mathbf{S}$. Recall that the semantics of *skip* is the update $\mathbf{1}$. So the following are equivalent:

1. $s \in [\![[skip]\varphi]\!]_{\mathbf{S}}$.
2. If $s\ \mathbf{1}_{\mathbf{S}}\ t$, then $t \in [\![\varphi]\!]_{\mathbf{S}(\mathbf{1})}$.
3. $s \in [\![\varphi]\!]_{\mathbf{S}}$.

$\dashv$

Proposition 5.6 *The Crash Axiom* $[crash]\mathsf{False}$ *is sound.*

Proof Recall that $[\![crash]\!] = \mathbf{0}$, the update such that $\mathbf{S}(\mathbf{0})$ is the empty model, and $\mathbf{0}_{\mathbf{S}}$ is the empty relation. Working through the definitions, $[crash]\mathsf{False}$ is easily seen to hold vacuously. $\dashv$

Proposition 5.7 *The Choice Axiom* $[\pi + \rho]\varphi \leftrightarrow [\pi]\varphi \wedge [\rho]\varphi$ *is sound.*

Proof Fix a state model $\mathbf{S}$; we drop $\mathbf{S}$ from the notation in the rest of this proof. Write γ for $\pi + \rho$, $\mathbf{r}$ for $[\![\pi]\!]$, $\mathbf{s}$ for $[\![\rho]\!]$, and $\mathbf{t}$ for $[\![\gamma]\!]$. Then the following are equivalent:

1. $s \in [\![[\pi + \rho]\varphi]\!]_{\mathbf{S}}$.
2. If $s\ \mathbf{t}\ u$, then $u \in [\![\varphi]\!]_{\mathbf{S}(\gamma)}$.
3. If $s\ \mathbf{r}\ t$, then $t \in [\![\varphi]\!]_{\mathbf{S}(\gamma)}$; and if $s\ \mathbf{s}\ t'$, then $t' \in [\![\varphi]\!]_{\mathbf{S}(\gamma)}$.
4. If $s\ \mathbf{r}\ t$, then $t \in [\![\varphi]\!]_{\mathbf{S}(\pi)}$; and if $s\ \mathbf{s}\ t'$, then $t' \in [\![\varphi]\!]_{\mathbf{S}(\rho)}$.

5. $s \in [\![[\pi]\varphi]\!]_{\mathbf{S}}$ and $s \in [\![[\rho]\varphi]\!]_{\mathbf{S}}$.

6. $s \in [\![[\pi]\varphi \wedge [\rho]\varphi]\!]_{\mathbf{S}}$.

The equivalence (2)$\Longleftrightarrow$(3) comes from the fact that each of the u such that $s \ \gamma \ u$ is either either (a) an element t of $\mathbf{S}(\mathbf{r})$ related to s by $\mathbf{r}$, or else (b) an element t' of $\mathbf{S}(\mathbf{s})$ related to s by $\mathbf{s}$. This is by the definition of $\mathbf{r} + \mathbf{s}$.

We show the equivalence (3)$\Longleftrightarrow$(4) by considering the cases (a) and (b) noted just above. We use the fact that the natural injections of $\mathbf{S}(\mathbf{r})$ and $\mathbf{S}(\mathbf{s})$ in $\mathbf{S}(\mathbf{t})$ are bisimulations, and also the fact that $[\![\varphi]\!]$ preserves bisimulations (see Proposition 4.2). $\dashv$

Proposition 5.8 *The Action Mix Axiom $[\pi*]\varphi \to \varphi \wedge [\pi][\pi^*]\varphi$ is sound. The Program Induction Rule is also sound: from $\vdash \chi \to \psi \wedge [\pi]\chi$, infer $\vdash \chi \to [\pi^*]\psi$.*

Proof These are standard. $\dashv$

A simplified form of the Action Rule Our logical system has one additional rule of inference, the Action Rule. We are not going to discuss this rule in full generality at this point (see Section 6 for this), but we do want to mention a simplified form of this that the reader could understand and use at this point.

Let $\mathbf{\Sigma}$ be an action signature, and let $\sigma_1, \ldots, \sigma_n$ be our enumeration of it. Let α be a simple action. Let ψ be a sentence, and let $\mathcal{C}$ be a set of agents. Let there be sentences χ_j for all $1 \leq j \leq n$, and such that

1. For all j, $\vdash \chi_j \to [\sigma_j \vec{\psi}]\varphi$.

2. If $A \in \mathcal{C}$ and $\sigma_j \xrightarrow{A} \sigma_k$, then $\vdash (\chi_j \wedge \psi_j) \to \Box_A \chi_k$.

From these assumptions, infer $\vdash \chi_i \to [\sigma_i \vec{\psi}]\Box^*_{\mathcal{C}}\varphi$.

One can see that the purpose of the Action Rule is to allow us to infer sentences of the form $[\sigma_i \vec{\psi}]\Box^*_{\mathcal{C}}\varphi$. Without this rule, or something like it, we have no way to do this. For the same reason, we shall need a result permitting us to infer sentences of more general forms, such as

$$[\sigma_{i_1}\vec{\psi}^1 \cdot \sigma_{i_2}\vec{\psi}^2 \cdot \sigma_{i_l}\vec{\psi}^l]\Box^*_{\mathcal{C}}\varphi.$$

The more general Action Rule is needed for this. In the case of $l = 0$, our statement statement just above simplifies to a familiar rule:

The Induction Rule From $\vdash \chi \to \psi \wedge \Box_A \chi$ for all A, infer $\vdash \chi \to \Box^*_{\mathcal{A}}\psi$.

This will turn out to be derivable from the Action Rule and the Skip Axiom.

In Sections 5.3 and 5.4 we shall use the simplified form of the Action Rule stated above. That is, we postpone all soundness discussions until Section 6.

5.2 Logical systems for the target logics

As we know, we have logics $\mathcal{L}_0(\mathbf{\Sigma})$ corresponding to the logics presented in Section 2.4. Each type of announcement gives us an action signature, and each action signature a logic. What we want to do here is to spell out what the axioms of $\mathcal{L}_0(\mathbf{\Sigma})$ come to when we specialize to some of those logics. In doing this, we find it convenient to adopt simpler notations tailored for the fragments.

The logic of public announcements is shown in Figure 10. We only included the axioms and rule of inference that specifically used the structure of the signature $\mathbf{\Sigma}_{\mathrm{pub}}$. So we did not include the sentential validities, the normality axiom for $\Box_A$, the composition axiom, modus ponens, etc.

Basic Axioms	
($[Pub\ \varphi]$-normality)	$\vdash [Pub\ \varphi](\psi \to \chi) \to ([Pub\ \varphi]\psi \to [Pub\ \varphi]\chi)$
Action Axioms	
(Atomic Permanence)	$\vdash [Pub\ \varphi]p \leftrightarrow (\varphi \to p)$
(Partial Functionality)	$\vdash [Pub\ \varphi]\neg\psi \leftrightarrow (\varphi \to \neg[Pub\ \varphi]\psi)$
(Action-Knowledge)	$\vdash [Pub\ \varphi]\Box_A\psi \leftrightarrow (\varphi \to \Box_A[Pub\ \varphi]\psi)$
Modal Rules	
($[Pub\ \varphi]$-necessitation)	From $\vdash \psi$, infer $\vdash [Pub\ \varphi]\psi$
Announcement Rule	From $\vdash \chi \to [Pub\ \varphi]\psi$ and $\vdash \chi \wedge \varphi \to \Box_A\chi$ for all A, infer $\vdash \chi \to [Pub\ \varphi]\Box^*_{\mathcal{A}}\psi$

Figure 10: The main points of the logic of public announcements.

Our next logic is the logic of completely private announcements to groups. We discussed the formalization of this in Section 4.3. We have actions $Pri^B\ \varphi$ and (of course) *skip*. The axioms and rules are just as in the logic of public announcements, with a few changes. We must of course consider the relativized operators $[Pri^B\ \varphi]$ instead of their simpler counterparts $[Pub\varphi]$.)

return here and fix the paragraph below

The actions $Pri_t(\varphi, \mathcal{B})$ all have True as their precondition, and since $(\mathsf{True} \to \psi)$ is logically equivalent to ψ, we certainly may omit these actions from the notation in the axioms and rules. The most substantive change which we need to make in Figure 10 concerns the Action-Knowledge Axiom. It splits into two axioms, noted below:

$$\begin{array}{ll} [Pri(\varphi,\mathcal{B})]\Box_A\psi \leftrightarrow (\varphi \to \Box_A[Pri(\varphi,\mathcal{B})]\psi) & \text{for } A \in \mathcal{B} \\ [Pri(\varphi,\mathcal{B})]\Box_A\psi \leftrightarrow (\varphi \to \Box_A\psi) & \text{for } A \notin \mathcal{B} \end{array}$$

The last equivalence says: assuming that φ is true, then after a private announcement of φ to the members of $\mathcal{B}$, an outsider knows ψ just in case she knew ψ before the announcement.

Finally, we study the logic of common knowledge of alternatives. The Action-Knowledge now becomes

$$\begin{array}{ll} [Cka(\vec{\varphi},\mathcal{B})]\Box_A\psi \leftrightarrow (\varphi_1 \rightarrow \Box_A[Cka(\vec{\varphi},\mathcal{B})]\psi) & \text{for } A \in \mathcal{B} \\ [Cka(\vec{\varphi},\mathcal{B})]\Box_A\psi \leftrightarrow (\varphi_1 \rightarrow \bigwedge_{0\leq i\leq k} \Box_A[Cka(\vec{\varphi}^i,\mathcal{B})]\psi) & \text{for } A \notin \mathcal{B} \end{array}$$

where in the last clause $(\varphi_1,\ldots,\varphi_n)^i$ is the sequence $\varphi_i,\varphi_1,\ldots,\varphi_{i-1},\varphi_{i+1},\ldots,\varphi_k$. (That is, we bring φ_i to the front of the sequence.)

5.3 Examples in the target logics

This section studies some examples of the logic at work.

We begin with an application of the Action Rule in the logic of public announcements. We show

$$\vdash \Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \rightarrow [Pub\ \mathsf{H}]\Box^*\neg\mathsf{T}.$$

That is, on the assumption that it is common knowledge that Heads and Tails are mutually exclusive, then as a result of a public announcement of Heads it will be common knowledge that the state is not Tails.

We give this application in detail. Recall that $\mathbf{\Sigma}_{\text{pub}}$ has one simple action which we call Pub. We take χ_{Pub} to be $\Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T})$. In addition $Pub \xrightarrow{A} Pub$ for all A, and there are no other arrows in $\mathbf{\Sigma}_{\text{pub}}$. We take α to be $Pub\ \mathsf{H}$; note that this is the only action accessible from itself in the canonical action model. To use the Action Rule, we must show that

1. $\vdash \Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \rightarrow [Pub\ \mathsf{H}]\neg\mathsf{T}$.
2. $\vdash (\Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \wedge \mathsf{H}) \rightarrow \Box_A\Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T})$ for all agents A.

From these assumptions, we may infer $\vdash \Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \rightarrow [Pub\ \mathsf{H}]\Box^*\neg\mathsf{T}$.

For the first statement,

(a)	$\mathsf{T} \leftrightarrow [Pub\ \mathsf{H}]\mathsf{T}$	Atomic Permanence
(b)	$(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \rightarrow (\mathsf{H} \rightarrow \neg[Pub\ \mathsf{H}]\mathsf{T})$	(a), propositional reasoning
(c)	$[Pub\ \mathsf{H}]\neg\mathsf{T} \leftrightarrow (\mathsf{H} \rightarrow \neg[Pub\ \mathsf{H}]\mathsf{T})$	Partial Functionality
(d)	$\Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \rightarrow (\mathsf{H} \leftrightarrow \neg\mathsf{T})$	Epistemic Mix
(e)	$\Box^*(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \rightarrow [Pub\ \mathsf{H}]\neg\mathsf{T}$	(d), (b), (c), propositional reasoning

And the second statement is an easy consequence of the Epistemic Mix Axiom.

What happens when a publically known fact is announced? One intuition about public announcements and common knowledge is that if φ is common knowledge, then announcing φ publically does not change anything. Formally, we express this by a scheme rather than a single equation:

$$\Box^*\varphi \rightarrow ([Pub\ \varphi]\psi \leftrightarrow \psi) \tag{13}$$

What we would like to say is $\Box^*\varphi \rightarrow \bigwedge_\psi([Pub\ \varphi]\Box_A\psi \leftrightarrow \psi)$, but of course this cannot be expressed in our language. So we consider only the sentences of the form (13), and we show that all of these are provable. We argue by induction on φ.

For an atomic sentence p, (13) follows from the Epistemic Mix and Atomic Permanence Axioms. The induction steps for $\wedge$ and $\neg$ are easy. Next, assume (13) for ψ. By necessitation and Epistemic Mix, we have

$$\Box^*\varphi \to (\Box_A[Pub\ \varphi]\psi \leftrightarrow \Box_A\psi)$$

Note also that by the Announcement-Knowledge Axiom

$$\Box^*\varphi \to ([Pub\ \varphi]\Box_A\psi \leftrightarrow \Box_A[Pub\ \varphi]\psi)$$

These two imply (13) for $\Box_A\psi$.

Finally, we assume (13) for ψ and prove it for $\Box^*_{\mathcal{B}}\psi$. We show first that $\Box^*\varphi \wedge \Box^*\psi \to [Pub\ \varphi]\Box^*\psi$. For this we use the Action Rule. We must show that

(1) $\Box^*\varphi \wedge \Box^*\psi \to [Pub\ \varphi]\psi$.

(2) $\Box^*\varphi \wedge \Box^*\psi \wedge \Box_A(\Box^*\varphi \wedge \Box^*\psi) \to \Box_A(\Box^*\varphi \wedge \Box^*\psi)$.

(1) is easy from our induction hypothesis, and (2) is an easy consequence of Epistemic Mix.

To conclude, we show $\Box^*\varphi \wedge [Pub\ \varphi]\Box^*\psi \to \Box^*\psi$. For this, we use the Induction Rule; that is, we show

(3) $\Box^*\varphi \wedge [Pub\ \varphi]\Box^*\psi \to \psi$.

(4) $\Box^*\varphi \wedge [Pub\ \varphi]\Box^*\psi \to \Box_A(\Box^*\varphi \wedge [Pub\ \varphi]\Box^*\psi)$ for all A.

For (3), we use Lemma 6.14, part (1) to see that $[Pub\ \varphi]\Box^*\psi \to [Pub\ \varphi]\psi$; and now (3) follows from our induction hypothesis. For (4), it will be sufficient to show that

$$\varphi \wedge [Pub\ \varphi]\Box^*\psi \to \Box_A[Pub\ \varphi]\Box^*\psi$$

This follows from Lemma 6.14, part (2).

Actions do not change common knowledge of simple sentences For yet another application, let ψ be any boolean combination of atomic sentences. Then for all actions α of any of our logics, $\vdash \psi \leftrightarrow [\alpha]\psi$. The proof is an easy induction on ψ. Even more, we have $\vdash \Box^*_{\mathcal{C}}\psi \leftrightarrow [\alpha]\Box^*_{\psi}$. In one direction, we use the Action Rule, and in the other, the Induction Rule.

5.4 Application: syntactic treatment of an earlier example

In this section, we return to the "cheating" Scenario 4. We further studied this example in Section 2.7; in that section we justified the earlier model. What we want to do here is to work out explicitly some of the points which we argued semantically. It might be noted that our entire discussion in this section is not really necessary: everything follows from the completeness of our logics and some facts on characterizing sentences. Still, it might help the reader to see how our logic works in this setting. In particular, it might be useful to see a concrete application of the Action Rule.

We consider the logic of private announcements to groups. Let α be the announcement of H to A privately, and let β be the announcement of True to both A and B. Then we have $\alpha \xrightarrow{A} \alpha$, $\alpha \xrightarrow{B} \beta$, and $\beta \xrightarrow{AB} \beta$.

We use the notation from Section 2.7. In particular, χ_{H} denotes the sentence which says that the coin lies H up but it is common knowledge that neither A nor B knows this (and it is common knowledge that $\mathsf{H} \leftrightarrow \neg\mathsf{T}$). And ψ is the sentence in equation (8) in Section 2.7. Our main claim in this section is $\vdash \chi_{\mathsf{H}} \to \langle\alpha\rangle\psi$.

We now present a list of sentences which are provable in our logic. Since a complete list would be too long, we only present the main steps. We have abbreviated NOBODY KNOWS to NK.

1. $[\alpha]\mathsf{H} \leftrightarrow (\mathsf{H} \leftrightarrow \mathsf{H})$. This is an Action-Knowledge (AK) axiom.

2. $[\alpha]\mathsf{H}$. By (1).

3. $[\alpha]\Box^*_A\mathsf{H}$. By (2) and the Action Rule, taking $\chi_\alpha = \mathsf{True}$.

4. $\neg\Box_B\neg\mathsf{H} \to (\mathsf{H} \to \neg[\beta]\Box_B\neg\mathsf{H})$. By (*) and propositional reasoning.

5. $\neg\Box_B\neg\mathsf{H} \to (\mathsf{H} \to \neg\Box_B[\beta]\neg\mathsf{H})$. By (4) and the AK Axiom $[\beta]\Box_B\neg\mathsf{H} \leftrightarrow \Box_B[\beta]\neg\mathsf{H}$.

6. $\neg\Box_B\neg\mathsf{H} \to (\mathsf{H} \to \neg[\alpha]\Box_B\neg\mathsf{H})$. By (5) and the AK Axiom $[\alpha]\Box_B\neg\mathsf{H} \leftrightarrow (\mathsf{H} \to \Box_B[\beta]\neg\mathsf{H})$.

7. $\neg\Box_B\neg\mathsf{H} \to [\alpha]\neg\Box_B\neg\mathsf{H}$. By (6) and a Partial Functionality (PF) Axiom.

8. $\Diamond_B\mathsf{T} \to [\alpha]\Diamond_B\mathsf{T}$. Similar to (4-7)

9. $\Box^*_A(\Diamond_B\mathsf{H} \wedge \Diamond_B\mathsf{T}) \to [\alpha]\Box^*_A(\Diamond_B\mathsf{H} \wedge \Diamond_B\mathsf{T})$. By the Action Rule, (7), and (8). We also use Mix and $[\alpha]$-normality.

10. $\Box^*_{A,B}\mathrm{NK} \to [\beta]\Box^*_{A,B}\mathrm{NK}$. By (*).

11. $\Box_B\Box^*_{A,B}\mathrm{NK} \to \Box_B[\beta]\Box^*_{A,B}\mathrm{NK}$. By (10) and necessitation.

12. $\Box^*_{A,B}\mathrm{NK} \to \Box_B[\beta]\Box^*_{A,B}\mathrm{NK}$. By (10) and (11).

13. $[\alpha]\Box_B\Box^*_{A,B}\mathrm{NK} \leftrightarrow (\mathsf{H} \to \Box_B[\beta]\Box^*_{A,B}\mathrm{NK})$. This is an AK axiom.

14. $\Box^*_{A,B}\mathrm{NK} \to [\alpha]\Box_B\Box^*_{A,B}\mathrm{NK}$. By (12) and (13).

15. $\Box^*_{A,B}\mathrm{NK} \to \Box_A\mathrm{NK}$. By Mix.

16. $\Box^*_{A,B}\mathrm{NK} \to [\alpha]\Box^*_A\Box_B\Box^*_{A,B}\mathrm{NK}$. From the Action Rule, (14) and (15).

17. $[\alpha]\Diamond_A\mathsf{True} \leftrightarrow \Diamond_A\langle\alpha\rangle\mathsf{True}$. Using a PF Axiom and an AK Axiom.

18. $\langle\alpha\rangle\mathsf{True} \leftrightarrow \mathsf{H}$. Using a PF Axiom and the necessitation fact $[\alpha]\mathsf{True}$.

19. $\Diamond_A\mathsf{H} \to [\alpha]\Diamond_A\mathsf{True}$. By (17) and (18).

20. $\Box^*_{A,B}(\mathsf{H} \leftrightarrow \neg\mathsf{T}) \to [\alpha]\Box^*_{A,B}(\mathsf{H} \leftrightarrow \neg\mathsf{T})$. By our work in Section 5.3.

21. $\chi_{\mathsf{H}} \to [\alpha]\psi$. Using the definition of χ_{H}, $[\alpha]$-Necessitation, (3), (16), (19), and (20).

22. $\chi_{\mathsf{H}} \to \langle\alpha\rangle\psi$. Using (21) and the equivalence of $\langle\alpha\rangle\psi$ and $\mathsf{H} \wedge [\alpha]\psi$.

5.5 Further results on the target logics

In this section, we mention some additional results on the target logics, including:

1. Interesting derivable principles in the logic of public announcements.
2. Qian's work on the question of which sentences φ in the logic of public announcements have the property that they are true after being announced.
3. Derivable principles in other logics.

More needed here!

Qian's Theorem We consider a natural question about the logic of public announcements, especially in its one-agent setting as the logic of modal relativization. Some sentences φ have the property that $\models [Pub\ \varphi]\varphi$. That is, whenever φ holds and is announced, then φ holds after the announcement. For example, the atomic sentences p have this property, as do their boolean combinations. So do sentences $\Box p$. The question here is to characterize syntactically those φ such that $[Pub\ \varphi]\varphi$ is valid. The first guess is that $[Pub\ \varphi]\varphi$ is valid iff φ is equivalent to a sentence in the *universal* fragment of modal logic, the fragment built from atomic sentences and their negations using $\Box$, $\wedge$, and $\vee$. However, this is not to be.

Suppose that φ and ψ are non-modal sentences (that is, boolean combinations of atomic sentences). Suppose that $\models \psi \rightarrow \varphi$. Consider $\varphi \vee \Diamond\psi$. This is clearly not in general equivalent to a sentence in our fragment. Yet we claim that

$$\models [Pub\ \varphi \vee \Diamond\psi](\varphi \vee \Diamond\psi).$$

To see this, fix a state model $\mathbf{S}$ and some $s \in S$. If $s \in [\![\varphi]\!]_{\mathbf{S}}$, then since φ is non-modal, $s \in [\![\varphi]\!]_{\mathbf{S}(Pub\ \varphi\vee\Diamond\psi)}$. On the other hand, suppose that $s \in [\![\Diamond\psi]\!]_{\mathbf{S}}$. Let $s \rightarrow t$ with $t \in [\![\psi]\!]_{\mathbf{S}} \subseteq [\![\varphi]\!]_{\mathbf{S}}$. Hence t belongs to the model after the announcement, $\mathbf{S}(Pub\ \varphi \vee \Diamond\psi)$. And $s \in [\![\Diamond\varphi]\!]_{\mathbf{S}(Pub\ \varphi\vee\Diamond\psi)}$ via t.

This example is due to Lei Qian. He also found a hypothesis under which the "first guess" above indeed holds. Here is his result. Let T_0 be the set of non-modal sentences. Let

$$T_1 \quad = \quad T_0 \cup \{\Box\varphi : \varphi \in T_0\} \cup \{\Diamond\varphi : \varphi \in T_0\}.$$

Finally, let T_2 be the closure of T_1 under $\wedge$ and $\vee$.

Theorem 5.9 (Qian [20]) *Let $\varphi \in T_2$ have the property that $\models [\mathrm{Pub}\ \varphi]\varphi$. Then there is some ψ in the universal fragment of modal logic such that $\models \varphi \leftrightarrow \psi$.*

We also should discuss some natural conjectures. For example, here is a very strong way to say that it is impossible to attain common knowledge by private announcements.

Conjecture Let there be just two agents, A and B. Let φ be a sentence in $\mathcal{L}_1(\mathbf{\Sigma}_{\mathrm{pri}})$, and let α be an action in $\mathcal{L}_1(\mathbf{\Sigma}_{\mathrm{pri}})$ in which all of the announcements are private. (That is, α has no occurrences of $Pri^{A,B}$, only Pri^A and Pri^B.) Suppose that $\models \varphi \rightarrow [\alpha]\Box^*_{A,B}p$. Then $\models \varphi \rightarrow \Box^*_{A,B}p$.

In other words, if φ implies that running α always leads to common knowledge, then φ itself already implies that there is common knowledge of φ (and presumably running α does not change this.)

Endnotes The logical system in this paper, like $\mathcal{L}(\mathbf{\Sigma})$ itself, is new here. However, the special cases for the target logics were known, at least in part. Plaza [18] constructs what we would call $\mathcal{L}_0(\mathbf{\Sigma}_{Pub})$, that is, the logic of public announcements without common knowledge operators or program iteration. (He worked only on models where each accessibility relation is an equivalence relation, so his system includes the S5 axioms.) Gerbrandy [7, 8], and also Gerbrandy and Groeneveld [9] went a bit further. They studied the logic of completely private announcements (generalizing public announcements) and presented a logical system which included the common knowledge operators. That is, their system included Mix. They argued that all of the reasoning in the original Muddy Children scenario can be carried out in their system. This is important because it shows that in order to get a formal treatment of that problem and related ones, one need not posit models which maintain histories. Their system was not complete since it did not have anything like the Action Rule; this first appears in a slightly different form in [2].

6 The Action Rule

Recall that our logical system is listed in Figure 9 above. We have shown the soundness of all of the axioms and rules, except for the Action Rule. But before we can even state the Action Rule, we need to discuss a structure that comes up in our work, the *canonical action model* $\mathbf{\Omega}$.

6.1 The canonical action model Ω

Recall that we defined simple action models and epistemic program models in Sections 3.1 and 3.2, respectively. At this point, we define a new simple action structure called the *canonical action model* $\mathbf{\Omega}$ of a language $\mathcal{L}_1(\mathbf{\Sigma})$.

Definition A *basic action* is a program (in $\mathcal{L}(\mathbf{\Sigma})$) of the form $\sigma\vec{\psi}$. A *simple action* is a program of $\mathcal{L}(\mathbf{\Sigma})$ in which neither the program sum operation $+$ nor the iteration operation π^* occur. We use letters like α and β to denote simple actions *only*. A *simple sentence* is a sentence of $\mathcal{L}_1(\mathbf{\Sigma})$ in which neither action sum nor action iteration occur. So all programs in simple sentences are simple actions.

The canonical action model Ω We define a program model $\mathbf{\Omega}$ in several steps. The simple actions of $\mathbf{\Omega}$ are the actions of $\mathcal{L}(\mathbf{\Sigma})$. For all A, the accessibility relation $\xrightarrow{A}$ is the smallest relation such that

1. *skip* $\xrightarrow{A}$ *skip*.
2. $\sigma_i\vec{\varphi} \to \sigma_j\vec{\psi}$ iff $\sigma_i \xrightarrow{A} \sigma_j$ in $\mathbf{\Sigma}$ and $\vec{\varphi} = \vec{\psi}$.
3. If $\alpha \xrightarrow{A} \alpha'$ and $\beta \xrightarrow{A} \beta'$, then $\alpha \cdot \beta \xrightarrow{A} \alpha' \cdot \beta'$.

Proposition 6.1 *As a frame,* $\mathbf{\Omega}$ *is* locally finite*: for each simple α, there are only finitely many β such that* $\alpha \xrightarrow{A^*} \beta$.

Proof By induction on α; we use heavily the fact that the accessibility relations on $\mathbf{\Omega}$ are the smallest family with their defining property. For the simple action expressions $\sigma\vec{\psi}$, we use the assumption that the action model $\mathbf{\Sigma}$ underlying all our definitions is finite. (Taking it to be locally finite would also be sufficient.) $\dashv$

Next, we define $\textsc{Pre} : \Omega \to \Omega$ by recursion so that

$$\begin{array}{lcl} \textsc{Pre}(\textit{skip}) & = & \mathsf{True} \\ \textsc{Pre}(\textit{crash}) & = & \mathsf{False} \\ \textsc{Pre}(\sigma_i\vec{\psi}) & = & \psi_i \\ \textsc{Pre}(\alpha \cdot \beta) & = & \langle\alpha\rangle\textsc{Pre}(\beta) \end{array}$$

Remark This function $\textsc{Pre}$ should not be confused with the function pre which is part of the structure of an epistemic action model. However, there is a connection: we are in the midst of defining the epistemic action model $\mathbf{\Omega}$, and its pre function is defined in terms of $\textsc{Pre}$. This is also perhaps a good place to remind the reader that neither $\textsc{Pre}$ nor pre is a first-class symbol in $\mathcal{L}(\mathbf{\Sigma})$; it is only a defined symbol. In Section 7 below we shall introduce another language called $\mathcal{L}_1^+(\mathbf{\Sigma})$. In that language, $\textsc{Pre}$ will be a first-class function symbol.

Completing the definition of $\mathbf{\Omega}$ we set

$$\mathrm{pre}(\sigma) \quad = \quad [\![\textsc{Pre}(\sigma)]\!].$$

This simple action model $\mathbf{\Omega}$ is the *canonical epistemic action model*; it plays the same role in our work as the canonical model in modal logic.

At this point, we need some results on the $\textsc{Pre}$ operation.

Lemma 6.2 $\vdash \langle\alpha\rangle\mathsf{True} \leftrightarrow \textsc{Pre}(\alpha)$.

Proof By induction on α. $\dashv$

Lemma 6.3 *The following monoid-type laws hold:*

1. $\vdash \textsc{Pre}(\alpha \cdot \mathrm{skip}) \leftrightarrow \textsc{Pre}(\alpha) \leftrightarrow \textsc{Pre}(\mathrm{skip} \cdot \alpha)$.
2. $\vdash \textsc{Pre}(\alpha \cdot (\beta \cdot \gamma)) \leftrightarrow \textsc{Pre}((\alpha \cdot \beta) \cdot \gamma)$.

Remark Later we shall see that these turn into *admissible rules* for the equivalence of actions. Lemma 6.3 will be used to establish this, and also several other important results.

Proof We use the definitions to calculate

$$\begin{array}{rcl} \textsc{Pre}(\alpha \cdot \mathit{skip}) & = & \langle\alpha\rangle\mathsf{True} \\ \textsc{Pre}(\mathit{skip} \cdot \alpha) & = & \langle \mathit{skip}\rangle\textsc{Pre}(\alpha). \end{array}$$

We obtain a formal proof using Lemma 6.2 and the Skip Axiom. Turning to the second law,

$$\begin{array}{rcl} \textsc{Pre}(\alpha \cdot (\beta \cdot \gamma)) & = & \langle\alpha\rangle\textsc{Pre}(\beta \cdot \gamma) \\ & = & \langle\alpha\rangle\langle\beta\rangle\textsc{Pre}(\gamma) \\ & \equiv & \langle\alpha \cdot \beta\rangle\textsc{Pre}(\gamma) \\ & = & \textsc{Pre}((\alpha \cdot \beta) \cdot \gamma). \end{array}$$

We replace $\equiv$ by provable equivalence, using the Composition Axiom. $\dashv$

This last result leads to another concept.

Definition $\approx$ is the smallest relation on $\mathbf{\Omega}$ containing the identity relation on $\mathbf{\Omega}$ and all instances of associativity,

$$\alpha \cdot (\beta \cdot \gamma) \quad \approx \quad (\alpha \cdot \beta) \cdot \gamma,$$

and closed under substitution: $\alpha \approx \beta$ and $\gamma \approx \delta$ imply $\alpha \cdot \gamma \approx \beta \cdot \delta$. We say that simple actions α and α' *differ by associativity* if $\alpha \approx \alpha'$.[5]

Proposition 6.4 *The relation $\approx$ of "differing by associativity" is a* $\textsc{Pre}$*-preserving bisimulation equivalence on $\mathbf{\Omega}$. That is, $\approx$ is an equivalence relation. And if $\alpha \approx \alpha'$, then:*

1. $\vdash \textsc{Pre}(\alpha) \leftrightarrow \textsc{Pre}(\alpha')$.
2. *Whenever $\alpha \xrightarrow{A} \beta$, there is some β' such that $\alpha' \xrightarrow{A} \beta'$ and $\alpha' \approx \beta'$.*

Proof (1) and (2) are shown by induction on the relation $\approx$. Lemma 6.3 is used in (1), and

[5] We could modify $\approx$ to include the instances of $\alpha \cdot \mathit{skip} \approx \alpha$ and $\mathit{skip} \cdot \alpha \approx \alpha$. All our results below would hold for this. However, we have no use for this larger relation.

the argument for (2) is direct. The transitivity of $\approx$ is similar. $\dashv$

We continue with some further general results on simple actions.

Proposition 6.5 *Let α be a simple action, let $\mathbf{S}$ be a state model, and let $s \in S$. For s and t' the following are equivalent:*

1. *There is some t such that s $[\![\alpha]\!]_{\mathbf{S}}$ t and $t \xrightarrow{A} t'$.*
2. *There are s' and α' such that $s \xrightarrow{A} s'$, $\alpha \xrightarrow{A} \alpha'$, and s' $[\![\alpha]\!]_{\mathbf{S}}$ t'.*

$$
\begin{tikzcd}
s \arrow[r, "{[\![\alpha]\!]_{\mathbf{S}}}"] \arrow[d, "A"'] & t \arrow[d, "A"] \\
s' \arrow[r, "{[\![\alpha']\!]_{\mathbf{S}}}"'] & t'
\end{tikzcd}
$$

Proof By induction on α. When $\alpha = skip$, the equivalence is immediate. We consider the case when $\alpha = \sigma_i\vec{\psi}$. In one direction, when s $\mathbf{r}_{\mathbf{S}}$ t, t must be (s, σ_i); and t' is of the form (s', σ_j) for some σ_j such that $\sigma_i \xrightarrow{A} \sigma_j$. We take α' to be $\sigma_j\vec{\psi}$. The other direction of the equivalence is similar.

Finally, assume the result for α and β, and consider $\alpha \cdot \beta$. First, assume that s $\mathbf{r} \cdot \mathbf{s}_{\mathbf{S}}$ t and $t \xrightarrow{A} t'$. Then there is some u such that s $\mathbf{r}_{\mathbf{S}}$ u $\mathbf{s}_{\mathbf{S}(\mathbf{r})}$ t. By induction hypothesis, there are u' and β' such that $u \xrightarrow{A} u'$, $\beta \xrightarrow{A} \beta'$, and u' $\mathbf{r}'_{\mathbf{S}(\mathbf{r})}$ t'. By induction hypothesis again, there are s' and α' such that $s \xrightarrow{A} s'$, $\alpha \xrightarrow{A} \alpha'$, and s' $\mathbf{r}'_{\mathbf{S}}$ u'. And now, $\mathbf{r} \cdot \mathbf{s} \xrightarrow{A} \mathbf{r}' \cdot \mathbf{s}'$, and s' $\mathbf{r}' \cdot \mathbf{s}'_{\mathbf{S}}$ t'. The other direction is similar. $\dashv$

Proposition 6.6 *For all simple actions α and all state models $\mathbf{S}$, $[\![\textsc{Pre}(\alpha)]\!]_{\mathbf{S}} = dom\ [\![\alpha]\!]_{\mathbf{S}}$. That is, the domain of the relation $[\![\alpha]\!]_{\mathbf{S}}$ is the interpretation in $\mathbf{S}$ of $\textsc{Pre}(\alpha)$.*

Proof By induction on α. The cases of *skip*, *crash*, and the simple action expressions are easy. We check the composition case, so assume the result for α and β. Then

$$
\begin{array}{lcl}
[\![\textsc{Pre}(\gamma)]\!]_{\mathbf{S}} & = & [\![\langle\alpha\rangle\textsc{Pre}(\beta)]\!]_{\mathbf{S}} \\
 & = & \{s \in S : \text{there is } t \text{ so that } s\ [\![\alpha]\!]_{\mathbf{S}}\ t \text{ and } t \in [\![\textsc{Pre}(\beta)]\!]_{\mathbf{S}([\![\alpha]\!])}\} \\
 & = & \{s \in S : \text{there is } t \text{ so that } s\ [\![\alpha]\!]_{\mathbf{S}}\ t \text{ and } t \in dom\ [\![\beta]\!]_{\mathbf{S}([\![\alpha]\!])}\} \\
 & = & dom\ ([\![\alpha]\!]_{\mathbf{S}} \cdot\ [\![\beta]\!]_{\mathbf{S}([\![\alpha]\!])}) \\
 & = & dom\ [\![\gamma]\!]_{\mathbf{S}}
\end{array}
$$

$\dashv$

We also need one more definition and result:

Definition We define a map $\alpha \mapsto \alpha^{\#}$, $\varphi \mapsto \varphi^{\#}$ from simple actions and simple sentences to the set N of natural numbers by the following recursion:

$$\begin{array}{lclclcl}
p^{\#} & = & 0 & \qquad ([\alpha]\varphi)^{\#} & = & 1+\max(\alpha^{\#},\varphi^{\#}) \\
(\neg\varphi)^{\#} & = & 1+\varphi^{\#} & \qquad \mathit{skip}^{\#} & = & 3 \\
(\varphi\wedge\psi)^{\#} & = & 1+\max(\varphi^{\#},\psi^{\#}) & \qquad \mathit{crash}^{\#} & = & 3 \\
(\Box_A\varphi)^{\#} & = & 1+\varphi^{\#} & \qquad (\sigma_i\vec{\psi})^{\#} & = & 1+\max_j \psi_j{}^{\#} \\
(\Box^*_{\mathcal{B}}\varphi)^{\#} & = & 1+\varphi^{\#} & \qquad (\alpha\cdot\beta)^{\#} & = & 3+\max(\alpha^{\#},\beta^{\#})
\end{array}$$

Proposition 6.7 *Concerning the operation $\#$:*

1. *For all simple actions α, $\mathrm{PRE}(\alpha)^{\#} < \alpha^{\#}$.*
2. *If φ and ψ are simple sentences and φ is a proper subsentence of ψ, then $\varphi^{\#} < \psi^{\#}$.*

Proof We prove the first part by induction on α. The cases of $\alpha = \mathit{skip}$, $\alpha = \mathit{crash}$, or α of the form $\sigma_i\vec{\psi}$ are easy. Assuming that $\mathrm{PRE}(\alpha)^{\#} < \alpha^{\#}$ and $\mathrm{PRE}(\beta)^{\#} < \beta^{\#}$, we see that

$$\begin{array}{rcl}
(\mathrm{PRE}(\alpha\cdot\beta))^{\#} & = & (\langle\alpha\rangle\mathrm{PRE}(\beta))^{\#} \\
& = & 2+\max(\alpha^{\#}, 1+(\mathrm{PRE}(\beta)^{\#}) \\
& < & 3+\max(\alpha^{\#},\beta^{\#}) \\
& = & (\alpha\cdot\beta)^{\#}
\end{array}$$

The second part is a similar induction. ⊣

At this point, we have all the general results we need. We take up the main issue of this section.

Definition For any action α, let $\hat{\alpha}$ be the update $(\mathbf{\Omega}, \{\alpha\})$ as we have defined it in Section 3.4.

Theorems 6.8 and 6.9 are the main result of this section. In reading them, recall that

$$\begin{array}{rcl}
([\hat{\alpha}][\![\varphi]\!])_{\mathbf{S}} & = & \{s\in S:\ \text{if } (s,\alpha)\in S\otimes\Omega, \text{ then } (s,\alpha)\in[\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}\} \\
& = & \{s\in S:\ \text{if } s\in[\![\mathrm{PRE}(\alpha)]\!]_{\mathbf{S}}, \text{ then } (s,\alpha)\in[\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}\}
\end{array}$$

and also that

$$[\![[\alpha]\varphi]\!]_{\mathbf{S}} \quad = \quad ([[\![\alpha]\!]][\![\varphi]\!])_{\mathbf{S}} \quad = \quad \{s\in S: \text{if } s\ [\![\alpha]\!]_{\mathbf{S}}\ t, \text{ then } t\in[\![\varphi]\!]_{\mathbf{S}([\![\alpha]\!])}\}.$$

Theorem 6.8 *For all simple actions α and simple sentences φ,*

$$[\![[\alpha]\varphi]\!] \quad = \quad [\hat{\alpha}][\![\varphi]\!].$$

That is, for all φ and all state models $\mathbf{S}$,

$$[\![[\alpha]\varphi]\!]_{\mathbf{S}} \quad = \quad \{s\in S: \textit{if } s\in[\![\mathrm{PRE}(\alpha)]\!]_{\mathbf{S}},\ \textit{then } (s,\alpha)\in[\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}\}. \tag{14}$$

Proof Write $\boldsymbol{\varphi}$ for $[\![\varphi]\!]$. We need to show that the following conditions are equivalent:

(a) If $s \; [\![\alpha]\!]_{\mathbf{S}} \; t$, then $t \in \boldsymbol{\varphi}_{\mathbf{S}([\![\alpha]\!])}$.

(b) If $s \in dom \; [\![\alpha]\!]_{\mathbf{S}}$, then $(s, \alpha) \in \boldsymbol{\varphi}_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.

We argue by induction on $\varphi^{\#}$. To strengthen the induction hypothesis, we also show that another equivalence (c)$\Longleftrightarrow$(d) holds for all β, and all $(s, \alpha) \in \mathbf{S} \otimes \boldsymbol{\Omega}$.

(c) $(s, \alpha) \in [\![[\beta]\varphi]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.

(d) $(s, \alpha) \in [\![\text{PRE}(\beta)]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$ implies $(s, \alpha \cdot \beta) \in [\![\varphi]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.

Finally, we also show at the same time the following assertion:

(e) If $\alpha \approx \alpha'$, then $(s, \alpha) \in [\![\varphi]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$ iff $(s, \alpha') \in [\![\varphi]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.

We argue all of this by induction on $\varphi^{\#}$.

$\varphi^{\#} = 0$ First, we check the equivalences when $\varphi^{\#} = 0$. In this case, φ is an atomic sentence (say) p. Fix α, $\mathbf{S}$ and $s \in S$. The following are equivalent:

1. If $s \; [\![\alpha]\!]_{\mathbf{S}} \; t$, then $t \in \boldsymbol{p}_{\mathbf{S}([\![\alpha]\!])}$.
2. If $s \in dom \; [\![\alpha]\!]_{\mathbf{S}}$, then $s \in \boldsymbol{p}_{\mathbf{S}}$.
3. If $s \in dom \; [\![\alpha]\!]_{\mathbf{S}}$, then $(s, \alpha) \in \boldsymbol{p}_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.

The equivalence (1)$\Longleftrightarrow$(2) follows from Proposition 4.3.

The equivalence (c)$\Longleftrightarrow$(d) in the atomic case is checked as follows. We assume $(s, \alpha) \in \mathbf{S} \otimes \boldsymbol{\Omega}$. We note the following equivalences:

1. $(s, \alpha) \in [\![[\beta]p]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.
2. If $(s, \alpha) \; [\![\beta]\!]_{\mathbf{S}} \; t$, then $t \in [\![p]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}([\![\beta]\!])}$.
3. If $(s, \alpha) \in dom \; [\![\beta]\!]_{\mathbf{S}}$, then $s \in [\![p]\!]_{\mathbf{S}}$.
4. If $(s, \alpha) \in [\![\text{PRE}(\beta)]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$, then $s \in [\![p]\!]_{\mathbf{S}}$.
5. If $(s, \alpha) \in [\![\text{PRE}(\beta)]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$, then $(s, \alpha \cdot \beta) \in [\![p]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.

(2)$\Longleftrightarrow$(3) here is by the semantics of p in $\mathbf{S} \otimes \boldsymbol{\Omega}$. (3)$\Longleftrightarrow$(4) is by Proposition 6.6. (4)$\Longleftrightarrow$(5) again uses the semantics of p in $\mathbf{S} \otimes \boldsymbol{\Omega}$. And if also the assumption that $(s, \alpha) \in \mathbf{S} \otimes \boldsymbol{\Omega}$. That is, this assumption and $(s, \alpha) \in [\![\text{PRE}(\beta)]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$, implies that $s \in [\![\text{PRE}(\alpha \cdot \beta)]\!]_{\mathbf{S}}$. Therefore $(s, \alpha \cdot \beta) \in \mathbf{S} \otimes \boldsymbol{\Omega}$. And the semantics of $\mathbf{S} \otimes \boldsymbol{\Omega}$ tells us that $s \in [\![p]\!]_{\mathbf{S}}$ iff $(s, \alpha \cdot \beta) \in [\![p]\!]_{\mathbf{S} \otimes \boldsymbol{\Omega}}$.

As for (e) in the atomic case, it follows from the definition of $\mathbf{S} \otimes \boldsymbol{\Omega}$, Lemma 6.3, and Proposition 6.4. This last result implies that (s, α) is an action in $\mathbf{S} \otimes \boldsymbol{\Omega}$ iff (s, α') also is one; hence our result.

The induction step For the rest of the proof, we argue by cases on the shape of φ. That is, we consider sentences $\neg\varphi$, $\varphi \wedge \psi$, $\Box_A \varphi$, $\Box^*_{\mathcal{B}} \varphi$, and $[\gamma]\varphi$. In each case, we assume that for all sentences ψ whose value under $\#$ is strictly smaller than the sentence in question, we have (a)$\Longleftrightarrow$(b), (b)$\Longleftrightarrow$(c), and (e).

$\neg\varphi$ Recall that $(\neg\varphi)^{\#} < \varphi^{\#}$. So we assume our result for φ and prove it for $\neg\varphi$. Fix α, $\mathbf{S}$ and $s \in S$. For (a)$\Longleftrightarrow$(b), we check that the following are equivalent:

1. If s $[\![\alpha]\!]_{\mathbf{S}}$ t, then $t \in (\neg\boldsymbol{\varphi})_{\mathbf{S}([\![\alpha]\!])}$.
2. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, then it is false that s $[\![\alpha]\!]_{\mathbf{S}}$ t implies $t \in \boldsymbol{\varphi}_{\mathbf{S}([\![\alpha]\!])}$.
3. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, then $(s,\alpha) \notin \boldsymbol{\varphi}_{\mathbf{S}\otimes\mathbf{\Omega}}$.
4. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, then $(s,\alpha) \in (\neg\boldsymbol{\varphi})_{\mathbf{S}\otimes\mathbf{\Omega}}$.

To see (c)$\Longleftrightarrow$(d), we use the following equivalences:

1. $(s,\alpha) \in [\![[\beta]\neg\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.
2. If $(s,\alpha) \in [\![\mathrm{P{\scriptsize RE}}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$, then $(s,\alpha) \notin [\![[\beta]\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.
3. If $(s,\alpha) \in [\![\mathrm{P{\scriptsize RE}}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$, then $(s,\alpha\cdot\beta) \notin [\![\varphi]\!]_{\mathbf{S}}$.
4. If $(s,\alpha) \in [\![\mathrm{P{\scriptsize RE}}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$, then $(s,\alpha\cdot\beta) \in [\![\neg\varphi]\!]_{\mathbf{S}}$.

The induction hypothesis is used in (1)$\Longleftrightarrow$(2) along with the overall definitions of the semantics.
(e) in this case is immediate.

$\varphi\wedge\psi$ The induction step for $\wedge$ is easy, and so we omit the details.

$\Box_A$ Next, assume (a)$\Longleftrightarrow$(b), (c)$\Longleftrightarrow$(d), and (e) for φ; we check it for $\Box_A\varphi$. Fix α, $\mathbf{S}$ and $s \in S$. First, the following establish (a)$\Longleftrightarrow$(b) for $\Box_A\varphi$.

1. If s $[\![\alpha]\!]_{\mathbf{S}}$ t, then $t \in (\Box_A\boldsymbol{\varphi})_{\mathbf{S}([\![\alpha]\!])}$.
2. If s $[\![\alpha]\!]_{\mathbf{S}}$ t and $t \xrightarrow{A} t'$, then $t' \in \boldsymbol{\varphi}_{\mathbf{S}([\![\alpha]\!])}$.
3. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, $\alpha \xrightarrow{A} \alpha'$, $s \xrightarrow{A} s'$, and s' $[\![\alpha]\!]_{\mathbf{S}}$ t', then $t' \in \boldsymbol{\varphi}_{\mathbf{S}([\![\alpha]\!])}$.
4. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, $\alpha \xrightarrow{A} \alpha'$ and $s \xrightarrow{A} s'$, then: if s' $[\![\alpha]\!]_{\mathbf{S}}$ t', then $t' \in \boldsymbol{\varphi}_{\mathbf{S}([\![\alpha]\!])}$.
5. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, $\alpha \xrightarrow{A} \alpha'$ and $s \xrightarrow{A} s'$, then: if $s' \in dom$ $[\![\alpha]\!]'_{\mathbf{S}}$, then $(s',\alpha') \in \boldsymbol{\varphi}_{\mathbf{S}\otimes\mathbf{\Omega}}$.
6. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, $s' \in dom$ $[\![\alpha]\!]'_{\mathbf{S}}$, and $(s,\alpha) \xrightarrow{A} (s',\alpha')$, then $(s',\alpha') \in \boldsymbol{\varphi}_{\mathbf{S}\otimes\mathbf{\Omega}}$.
7. If $s \in dom$ $[\![\alpha]\!]_{\mathbf{S}}$, then $(s,\alpha) \in (\Box_A\boldsymbol{\varphi})_{\mathbf{S}\otimes\mathbf{\Omega}}$.

The equivalence (2)$\Longleftrightarrow$(3) comes from Proposition 6.5. (3)$\Longleftrightarrow$(5) is by the induction hypothesis. And we check (c)$\Longleftrightarrow$(d) for $\Box_A\varphi$ as follows:

1. $(s,\alpha) \in [\![[\beta]\Box_A\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.
2. Whenever (s,α) $[\![\beta]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ x and $x \xrightarrow{A} y$, we have $y \in [\![\varphi]\!]_{(\mathbf{S}\otimes\mathbf{\Omega})([\![\beta]\!])}$.
3. Whenever $(s,\alpha) \xrightarrow{A} (s',\alpha')$, $\beta \xrightarrow{A} \beta'$, and (s',α') $[\![\beta']\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ x, $x \in [\![\varphi]\!]_{(\mathbf{S}\otimes\mathbf{\Omega})([\![\beta']\!])}$.
4. Whenever $(s,\alpha) \xrightarrow{A} (s',\alpha')$ and $\beta \xrightarrow{A} \beta'$, $(s',\alpha') \in [\![[\beta']\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

5. Whenever $(s,\alpha) \xrightarrow{A} (s',\alpha')$, $\beta \xrightarrow{A} \beta'$, and $(s,\alpha) \in [\![\mathrm{PRE}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$, $(s',\alpha'\cdot\beta') \in [\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

6. $(s,\alpha) \in [\![\mathrm{PRE}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$, $s \xrightarrow{A} s'$, $\alpha \xrightarrow{A} \alpha'$, and $\beta \xrightarrow{A} \beta'$ imply $(s,\alpha'\cdot\beta') \in [\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

7. $(s,\alpha) \in [\![\mathrm{PRE}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ implies $(s,\alpha\cdot\beta) \in [\![\Box_A\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

The crucial equivalence here is (2)$\Longleftrightarrow$(3); this is by Proposition 6.5.

To complete our discussion of $\Box_A\varphi$, we check (e) for it. We use Proposition 6.4, and the argument is straightforward.

$\Box^*_{\mathcal{B}}\varphi$ We have (a)$\Longleftrightarrow$(b), (c)$\Longleftrightarrow$(d), and (e) for φ. By what we showed above, we have these equivalences for $\Box_A\varphi$ for all A. We can use this to show by induction that (a)$\Longleftrightarrow$(b), (c)$\Longleftrightarrow$(d), and (e) hold for every sentence of the form $\Box_{A_1}\Box_{A_2}\cdots\Box_{A_k}\varphi$, with all A's from the set $\mathcal{B}$. Let W be the set of such sentences; so $\Box^*\mathcal{B}\varphi$ is semantically equivalent to $\bigwedge W$. This gives (a)$\Longleftrightarrow$(b) and (e) for $\Box^*_{\mathcal{B}}\varphi$ (for all α, $\mathbf{S}$ and s). The details are similar to the argument for (c)$\Longleftrightarrow$(d), and so we'll just prove this part. Fix α, $\mathbf{S}$, s, and β. Then the following are equivalent:

1. $(s,\alpha) \in [\![[\beta]\Box^*_{\mathcal{B}}\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

2. For all $\psi \in W$, $(s,\alpha) \in [\![[\beta]\psi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

3. $(s,\alpha) \in [\![\mathrm{PRE}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ implies that for $\psi \in W$, $(s,\alpha\cdot\beta) \in [\![[\beta]\psi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

4. $(s,\alpha) \in [\![\mathrm{PRE}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ implies $(s,\alpha\cdot\beta) \in [\![\Box^*_{\mathcal{B}}\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

(1)$\Longleftrightarrow$(2) here is by the semantics of $[\beta]\Box^*\varphi$. The other equivalences are easy from (c)$\Longleftrightarrow$(d) for $\psi \in W$.

$[\gamma]\varphi$ Assume our result for sentences with smaller value than $[\gamma]\varphi$. (In particular, we have the result for $\mathrm{PRE}(\gamma)$; this is the main point.) We first check (a)$\Longleftrightarrow$(b) for $[\gamma]\varphi$. Fix α, $\mathbf{S}$ and $s \in S$. Then the following are equivalent:

1. If $s\ [\![\alpha]\!]_{\mathbf{S}}\ t$, then $t \in ([\gamma]\boldsymbol{\varphi})_{\mathbf{S}([\![\alpha]\!])}$.

2. If $s\ [\![\alpha]\!]_{\mathbf{S}}\ t$ and $t\ [\![\gamma]\!]_{\mathbf{S}([\![\alpha]\!])}\ u$, then $u \in \boldsymbol{\varphi}_{\mathbf{S}([\![\alpha]\!])([\![\gamma]\!])}$.

3. If $s\ ([\![\alpha]\!]\cdot[\![\gamma]\!])_{\mathbf{S}}\ u$, then $u \in \boldsymbol{\varphi}_{\mathbf{S}([\![\alpha]\!]\cdot[\![\gamma]\!])}$.

4. If $s \in dom\ ([\![\alpha]\!]\cdot[\![\gamma]\!])_{\mathbf{S}}$, then $(s,\alpha\cdot\gamma) \in \boldsymbol{\varphi}_{\mathbf{S}\otimes\mathbf{\Omega}}$.

5. If $s \in dom\ [\![\alpha]\!]_{\mathbf{S}}$ and $(s,\alpha) \in [\![\mathrm{PRE}(\gamma)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$, then $(s,\alpha\cdot\gamma) \in \boldsymbol{\varphi}_{\mathbf{S}\otimes\mathbf{\Omega}}$.

6. If $s \in dom\ [\![\alpha]\!]_{\mathbf{S}}$, then $(s,\alpha) \in ([\gamma]\boldsymbol{\varphi})_{\mathbf{S}\otimes\mathbf{\Omega}}$.

We used (c)$\Longleftrightarrow$(d) for φ. Finally, we consider (c)$\Longleftrightarrow$(d) for $[\gamma]\varphi$. Fix β, $\mathbf{S}$, and s. The following are then equivalent:

1. $(s,\alpha) \in [\![[\beta][\gamma]\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

2. $(s,\alpha) \in [\![[\beta\cdot\gamma]\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

3. $(s, \alpha) \in [\![\text{PRE}(\beta \cdot \gamma)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ implies $(s, \alpha \cdot (\beta \cdot \gamma)) \in [\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

4. $(s, \alpha) \in [\![\text{PRE}(\beta \cdot \gamma)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ implies $(s, (\alpha \cdot \beta) \cdot \gamma) \in [\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

5. $(s, \alpha) \in [\![\langle\beta\rangle\text{PRE}(\gamma)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ implies $(s, (\alpha \cdot \beta) \cdot \gamma) \in [\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

6. $(s, \alpha) \in [\![\text{PRE}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ and $(s, \alpha \cdot \beta) \in [\![\text{PRE}(\gamma)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ imply $(s, (\alpha \cdot \beta) \cdot \gamma) \in [\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

7. $(s, \alpha) \in [\![\text{PRE}(\beta)]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$ implies $(s, \alpha \cdot \beta) \in [\![[\gamma]\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.

We have used the Composition Axiom in (1)$\Longleftrightarrow$(2). (2)$\Longleftrightarrow$(3) and (6)$\Longleftrightarrow$(7) are by the induction hypothesis. (3)$\Longleftrightarrow$(4) is by (e) for φ. (Indeed, this very step is the main place that (e) is used.) (4)$\Longleftrightarrow$(5) is by the definition of PRE. As for (5)$\Longleftrightarrow$(6), it holds by the induction hypothesis. That is, $(\text{PRE}(\gamma))^{\#} < [\gamma]\varphi^{\#}$. So we may use the equivalence of (c) and (d) for $\text{PRE}(\gamma)$. We also used some propositional reasoning at this point. This concludes the verification of (c)$\Longleftrightarrow$(d) for $[\gamma]\varphi$.

Finally, (e) for $[\gamma]\varphi$ follows from (c)$\Longleftrightarrow$(d) for φ. As in showing (e) for φ atomic, we use Lemma 6.3. ⊣

Theorem 6.9 *For all simple actions α and* all *sentences φ of $\mathcal{L}_1(\mathbf{\Sigma})$, $[\![[\alpha]\varphi]\!] = [\hat{\alpha}][\![\varphi]\!]$. That is, (14) holds even when $+$ occurs in φ.*

Proof By induction on the number of occurrences of $+$ in φ. The base case is when this number is 0, and this is Theorem 6.8. The induction step is then again by cases. The most interesting step is for sentences of the form $[\alpha + \beta]\varphi$, and for this we use the soundness of the Choice Axiom. ⊣

6.2 Soundness of the Action Rule

As we have mentioned, the Action Rule is the main rule in the system for $\mathcal{L}_1(\mathbf{\Sigma})$ that goes beyond the rules of $\mathcal{L}_0(\mathbf{\Sigma})$. It is an induction rule which allows us to prove that common knowledge obtains after various actions. Here is our statement of the Action Rule:

Let α be a simple action. Let ψ be a sentence, and let $\mathcal{C}$ be a set of agents. Let there be sentences χ_β for all β such that $\alpha \to_{\mathcal{C}}^{*} \beta$ (including α itself), and such that

1. $\vdash \chi_\beta \to [\beta]\psi$.

2. If $A \in \mathcal{C}$ and $\beta \xrightarrow{A} \gamma$, then $\vdash (\chi_\beta \wedge \text{PRE}(\beta)) \to \Box_A \chi_\gamma$.

From these assumptions, infer $\vdash \chi_\alpha \to [\alpha]\Box_{\mathcal{C}}^{*}\psi$.

Remark We use $\to_{\mathcal{C}}^{*}$ as an abbreviation for the reflexive and transitive closure of the relation $\bigcup_{A\in C} \xrightarrow{A}$. Recall from Proposition 6.1 that there are only finitely many β such that $\alpha \to_{\mathcal{C}}^{*} \beta$. So even though the Action Rule might look like it takes infinitely many premises, it really only takes finitely many.

Remark It is possible to drop the Composition Axiom in favor of a more involved version of the Action Rule. The point is we shall later introduce normal forms for sentences, and using the Composition Axiom will greatly simplify these normal forms. The Action Rule here is geared towards those normal forms. So if we were to drop the Composition Axiom, we would need a stronger, more complicated, formulation of the Action Rule, one which involved *sequences* of actions. It is not terribly difficult to formulate such a rule, and completeness can be obtained by an elaboration of the work which we shall do. We did not present this work, mostly because adding the Composition Axiom leads to shorter proofs.

The rest of this section is devoted to a proof of the soundness of the Action Rule, beginning with a lemma of independent interest. We'll turn to examples in the target logics and examples afterwards, and the reader may prefer to skip ahead to that discussion.

Lemma 6.10 *$s \in [\![\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\varphi]\!]_{\mathbf{S}}$ iff there is a sequence of states from S*

$$s \;=\; s_0 \quad \to_{A_1} \quad s_1 \quad \to_{A_2} \quad \cdots \quad \to_{A_{k-1}} \quad s_{k-1} \quad \to_{A_k} \quad s_k$$

where $k \geq 0$, and also a sequence of actions of the same length k,

$$\alpha \;=\; \alpha_0 \quad \to_{A_1} \quad \alpha_1 \quad \to_{A_2} \quad \cdots \quad \to_{A_{k-1}} \quad \alpha_{k-1} \quad \to_{A_k} \quad \alpha_k$$

such that $A_i \in \mathcal{C}$ and $s_i \in [\![\mathrm{PRE}(\alpha_i)]\!]_{\mathbf{S}}$ for all $0 \leq i < k$, and $s_k \in [\![\langle\alpha_k\rangle\varphi]\!]_{\mathbf{S}}$.

Remark The case $k = 0$ just says that $s \in [\![\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\varphi]\!]_{\mathbf{S}}$ is implied by $s \in [\![\langle\alpha\rangle\varphi]\!]_{\mathbf{S}}$.

Proof The following are equivalent:

1. $s \in [\![\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\varphi]\!]_{\mathbf{S}}$.
2. $s \in [\![\mathrm{PRE}(\alpha)]\!]_{\mathbf{S}}$, and $(s, \alpha) \in [\![\Diamond^*_{\mathcal{C}}\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.
3. $s \in [\![\mathrm{PRE}(\alpha)]\!]_{\mathbf{S}}$, and for some $k \geq 0$ there is a sequence in $\mathbf{S} \otimes \mathbf{\Omega}$,

   $$(s, \alpha) \;=\; t_0 \quad \to_{A_1} \quad t_1 \quad \to_{A_2} \quad \cdots \quad \to_{A_{k-1}} \quad t_{k-1} \quad \to_{A_k} \quad t_k$$

   such that $A_i \in \mathcal{C}$ and $t_k \in [\![\varphi]\!]_{\mathbf{S}\otimes\mathbf{\Omega}}$.
4. There are sequences of $s = s_0, \ldots, s_k$ and $\alpha_0, \ldots, \alpha_k$ as in the statement of this lemma.

The equivalences (1)$\Longleftrightarrow$(2) and (3)$\Longleftrightarrow$(4) are consequences of equation (14) in Section 6.1. (2)$\Longleftrightarrow$(3) is just the semantics of $\Diamond^*_{\mathcal{C}}$. $\dashv$

Proposition 6.11 *The Action Rule is sound.*

Proof Assume that $s \in [\![\chi_\alpha]\!]_{\mathbf{S}}$ but also $s \in [\![\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\neg\psi]\!]_{\mathbf{S}}$. According to Lemma 6.10, there is a labeled sequence of worlds from S

$$s \;=\; s_0 \quad \to_{A_1} \quad s_1 \quad \to_{A_2} \quad \cdots \quad \to_{A_{k-1}} \quad s_{k-1} \quad \to_{A_k} \quad s_k$$

where $k \geq 0$ and each $A_i \in \mathcal{C}$, and also a sequence of actions of length k, with the same labels,

$$\alpha \;=\; \alpha_0 \quad \to_{A_1} \quad \alpha_1 \quad \to_{A_2} \quad \cdots \quad \to_{A_{k-1}} \quad \alpha_{k-1} \quad \to_{A_k} \quad \alpha_k$$

such that $s_i \in [\![\mathrm{PRE}(\alpha_i)]\!]_{\mathbf{S}}$ for all $0 < i \leq k$, and $s_k \in [\![\langle\alpha_k\rangle\neg\psi]\!]_{\mathbf{S}}$. If $k = 0$, we would have $s \in [\![\langle\alpha\rangle\neg\psi]\!]_{\mathbf{S}}$. But one of the assumptions in the statement of the Action Rule is that $\vdash \chi_{\alpha_0} \to [\alpha_0]\psi$. So we would have $s \in [\![[\alpha]\psi]\!]_{\mathbf{S}}$. This would be a contradiction.

Now we argue the case $k > 0$. We show by induction on $1 \leq i \leq k$ that $s_i \in [\![\chi_{\alpha_i}]\!]_{\mathbf{S}}$. The case $i = 0$ is the opening assumptions of this proof. Assume that $s_i \in [\![\chi_{\alpha_i}]\!]_{\mathbf{S}}$. By hypothesis, $s_i \in [\![\mathrm{PRE}(\alpha_i)]\!]_{\mathbf{S}}$. In view of the second assumption in the Action Rule, $s_i \models \Box_{A_{i+1}}\chi_{\alpha_{i+1}}$. Hence $s_{i+1} \in [\![\chi_{\alpha_{i+1}}]\!]_{\mathbf{S}}$. By the first assumption, $s_i \in [\![[\alpha_{i+1}]\psi]\!]_{\mathbf{S}}$. This completes our induction.

In particular, $s_k \in [\![\chi_{\alpha_k}]\!]_{\mathbf{S}}$. Using again the first assumption in the Action Rule, we have $s_k \in [\![[\alpha_k]\psi]\!]_{\mathbf{S}}$. This is a contradiction. $\dashv$

6.3 Some syntactic results

At this point, we return to our study of the logic for $\mathcal{L}_1(\mathbf{\Sigma})$. We establish some syntactic results which will be used in the next section, and in our proof of the completeness of the logic for $\mathcal{L}_1(\mathbf{\Sigma})$.

Lemma 6.12 *Let α be a composition of actions of the form $\sigma\vec{\psi}$. Then there is some $\hat{\alpha}$ of the form*

$$\sigma_1\vec{\psi^1} \cdot (\sigma_2\vec{\psi^2} \cdot (\cdots \;\; \sigma_r\vec{\psi^r}))$$

such that for all φ, $\vdash [\alpha]\varphi \leftrightarrow [\hat{\alpha}]\varphi$, and $\vdash \mathrm{PRE}(\alpha) \leftrightarrow \mathrm{PRE}(\hat{\alpha})$. Moreover,

$$\{\beta^* : \alpha \xrightarrow{A} \beta\} \quad = \quad \{\gamma : \hat{\alpha} \xrightarrow{A} \gamma\}.$$

Proof $\hat{\alpha}$ is a re-association of the simple actions in α. We can check by induction on α that $\vdash [\alpha]\varphi \leftrightarrow [\hat{\alpha}]\varphi$ for all φ. The fact that $\vdash \mathrm{PRE}(\alpha) \leftrightarrow \mathrm{PRE}(\hat{\alpha})$ comes from Lemma 6.3. $\dashv$

Lemma 6.13 *The Action-Knowledge Axiom is provable for all simple actions α:*

$$\vdash [\alpha]\Box_A\varphi \leftrightarrow (\mathrm{PRE}(\alpha) \to \bigwedge\{\Box_A[\beta]\varphi : \alpha \xrightarrow{A} \beta \text{ in } \mathbf{\Omega}\}) \tag{15}$$

Proof By induction on the length of α. If the length of α is 0, α is *skip* or *crash*. In the first case, $\textsc{Pre}(\alpha) = \mathsf{True}$, the only β with $\alpha \xrightarrow{A} \beta$ is *skip*, and (15) reads:

$$\vdash [skip]\Box_A\varphi \leftrightarrow (\mathsf{True} \to \Box_A[skip]\varphi)$$

And this is an easy consequence of the Skip Axiom. In the second case, $\textsc{Pre}(\alpha) = \mathsf{False}$, and there are no β such that $\alpha \xrightarrow{A} \beta$. (15) then reads:

$$\vdash [crash]\Box_A\varphi \leftrightarrow (\mathsf{False} \to \mathsf{True}).$$

By the Crash Axiom and modal reasoning, $\vdash [crash]\Box_A\varphi$; and so the assertion just above holds. This takes care of the case when the length of α is 0.

If the length is 1, then we simply have the Action-Knowledge Axiom in the form we know it. So assume our lemma for actions α of length n. We may replace α by $\hat{\alpha}$ in Lemma 6.12 if necessary, to assume that composition in α associates to the right. (That is, (15) for $\hat{\alpha}$ implies (15) for α, as implied by the various assertions in Lemma 6.12.) We therefore show that for all actions $\sigma\vec{\psi}$,

$$\vdash [\sigma\vec{\psi}\cdot\alpha]\Box_A\varphi \leftrightarrow (\textsc{Pre}(\sigma\vec{\psi}\cdot\alpha) \to \bigwedge\{\Box_A[\tau\vec{\psi}\cdot\beta]\varphi : \alpha \xrightarrow{A} \beta \text{ in } \mathbf{\Omega} \text{ and } \sigma \xrightarrow{A} \tau \text{ in } \mathbf{\Sigma}\}) \qquad (16)$$

We start with (15), use $[\sigma\vec{\psi}]$-necessitation and normality, and get

$$\vdash [\sigma\vec{\psi}][\alpha]\Box_A\varphi \leftrightarrow ([\sigma\vec{\psi}]\textsc{Pre}(\alpha) \to \bigwedge\{[\sigma\vec{\psi}]\Box_A[\beta]\varphi : \alpha \xrightarrow{A} \beta \text{ in } \mathbf{\Omega}\}) \qquad (17)$$

The Composition Axiom tells us that $\vdash [\sigma\vec{\psi}][\alpha]\Box_A\varphi \leftrightarrow [\sigma\vec{\psi}\cdot\alpha]\Box_A\varphi$. And the Action-Knowledge Axiom that we already know tells us that for each β

$$[\sigma\vec{\psi}]\Box_A[\beta]\varphi \leftrightarrow (\textsc{Pre}(\sigma\vec{\psi}) \to \bigwedge\{\Box_A[\tau\vec{\psi}][\beta]\varphi : \sigma \xrightarrow{A} \tau\})$$

So (17) leads to

$$\vdash [\sigma\vec{\psi}\cdot\alpha]\Box_A\varphi \leftrightarrow ([\sigma\vec{\psi}]\textsc{Pre}(\alpha) \to (\textsc{Pre}(\sigma\vec{\psi}) \to \bigwedge\{\Box_A[\tau\vec{\psi}\cdot\beta]\varphi : \alpha \xrightarrow{A} \beta \text{ in } \mathbf{\Omega} \text{ and } \sigma \xrightarrow{A} \tau \text{ in } \mathbf{\Sigma}\}))$$

But $\textsc{Pre}(\sigma\vec{\psi}\cdot\alpha) = \textsc{Pre}(\sigma\vec{\psi})\cdot[\sigma\vec{\psi}]\textsc{Pre}(\alpha)$. Using this and some propositional reasoning, we get (16), as desired. $\dashv$

Remark The Action-Knowledge Axiom is not sound if the actions contain $+$. Here is an example. Let $\mathcal{A}$ be a singleton, hence omitted from the notation. $\mathbf{\Sigma} = \{\alpha, \beta, \gamma, \delta\}$ with $\alpha \to \beta$, $\gamma \to \delta$, and no other arrows.

Lemma 6.14 *For all $A \in \mathcal{C}$ and all β such that $\alpha \to_A \beta$,*

1. $\vdash [\alpha]\Box^*_{\mathcal{C}}\psi \to [\alpha]\psi$.
2. $\vdash [\alpha]\Box^*_{\mathcal{C}}\psi \wedge \textsc{Pre}(\alpha) \to \Box_A[\beta]\Box^*_{\mathcal{C}}\psi$.

Proof Part (1) follows easily from the Epistemic Mix Axiom and modal reasoning. For part (2), we start with a consequence of the Epistemic Mix Axiom: $\vdash \Box^*_{\mathcal{C}}\psi \to \Box_A\Box^*_{\mathcal{C}}\psi$. Then by modal reasoning, $\vdash [\alpha]\Box^*_{\mathcal{C}}\psi \to [\alpha]\Box_A\Box^*_{\mathcal{C}}\psi$. By the Action-Knowledge Axiom generalized as we have it in Lemma 6.13, we have $\vdash [\alpha]\Box^*_{\mathcal{C}}\psi \wedge \textsc{Pre}(\alpha) \to \Box_A[\beta]\Box^*_{\mathcal{C}}\psi$. $\dashv$

6.4 Equivalence of updates, programs, and actions

In this section, we formulate syntactic and semantic notions of equivalence for simple actions. As it happens, there are many open questions about these notions. But we are able to prove a few results for later use.

Equivalence of updates A bisimulation preserving update may be regarded as a relation on the quotient of states by $\equiv$. It makes sense to say that bisimulation preserving $\mathbf{r}$ and $\mathbf{s}$ are *equivalent* if they induce the same relation on the quotient. More concretely,

1. If $s\ \mathbf{r}_{\mathbf{S}}\ s'$, then there is some t' such that $t\ r^{\mathbf{S}}_{\mathbf{T}}\ t'$ and $(\mathbf{S}(\mathbf{r}), s') \equiv (\mathbf{T}(\mathbf{s}), t')$.

2. If $t\ \mathbf{r}_{\mathbf{T}}\ t'$, then there is some s' such that $s\ \mathbf{r}_{\mathbf{S}}\ s'$ and $(\mathbf{S}(\mathbf{r}), s') \equiv (\mathbf{T}(\mathbf{s}), t')$.

So it makes sense to say that two programs π and ρ are equivalent if their semantics $[\![\pi]\!]$ and $[\![\rho]\!]$ are equivalent updates. We shall study this notion primarily when π and ρ are simple actions.

The notion of equivalence turns out to be too semantic for us to use in connection with our proof system. So we need a more syntactic concept.

Definition A $\textsc{Pre}$-*preserving bisimulation* is a relation R on simple actions such that if $\alpha R \beta$, then

1. $\vdash \textsc{Pre}(\alpha) \leftrightarrow \textsc{Pre}(\beta)$.

2. For all α' and A such that $\alpha \xrightarrow{A} \alpha'$, there is some β' such that $\beta \xrightarrow{A} \beta'$ and $\alpha' R \beta'$.

3. For all β' and A such that $\beta \xrightarrow{A} \beta'$, there is some α' such that $\alpha \xrightarrow{A} \alpha'$ and $\alpha' R \beta'$.

We say that α and β are *equivalent* if there is some $\textsc{Pre}$-preserving bisimulation relating them. We write $\alpha \equiv \beta$ in this case.

For sentences, we write $\varphi \equiv \psi$ and say that φ and ψ are *equivalent* if $\vdash \varphi \leftrightarrow \psi$.

Lemma 6.15 *Concerning equivalence of actions:*

1. *The relation $\equiv$ of action equivalence is an equivalence relation.*

2. *If $\models \varphi_1 \leftrightarrow \psi_1$, …, $\models \varphi_n \leftrightarrow \psi_n$, then $\sigma_i\vec{\varphi} \equiv \sigma_i\vec{\psi}$.*

3. *$\alpha \cdot skip \equiv \alpha \equiv skip \cdot \alpha$.*

4. *$\alpha \cdot (\beta \cdot \gamma) \equiv (\alpha \cdot \beta) \cdot \gamma$.*

5. *If $\alpha \equiv \alpha'$ and $\beta \equiv \beta'$, then $\alpha \cdot \beta \equiv \alpha' \cdot \beta'$.*

Proof The first part is routine. In the second, we take R to be the set of pairs $(\tau\vec{\varphi}, \tau\vec{\psi})$ as τ

ranges over the simple action structure Σ. Lemma 6.3 is used in parts (3) and (4). The rest of arguments for part(3) is easy and we omit it. For part (4), we take R to be the set of pairs

$$(\alpha' \cdot (\beta' \cdot \gamma'), (\alpha' \cdot \beta') \cdot \gamma')$$

such that for some word w in $\mathcal{A}^*$, α' is reachable from α via w, and similarly for β' and γ' (via the same w). Part (5) is similar. ⊣

Lemma 6.16 *If α and β are equivalent actions, then $[\![\alpha]\!]$ and $[\![\beta]\!]$ are equivalent updates.*

Proof Fix a PRE-preserving bisimulation R_0 relating α and β. Write $\mathbf{r}$ for $[\![\alpha]\!]$ and $\mathbf{s}$ for $[\![\beta]\!]$. Fix $\mathbf{S}$. We check that the following is a bisimulation between the state models $\mathbf{S}(\mathbf{r})$ and $\mathbf{S}(\mathbf{s})$: $x\ R\ y$ iff there is some $s \in S$ such that $s\ \mathbf{r}_\mathbf{S}\ x$ and $s\ \mathbf{s}_\mathbf{S}\ y$. The "atomic harmony" part comes from Proposition 4.3. Suppose that $x\ R\ y$ via $s \in S$, and also $x \xrightarrow{A} x'$. By Proposition 6.5 there are some s' and α' such that $s \xrightarrow{A} s'$, $\alpha \xrightarrow{A} \alpha'$, and $s'\ \mathbf{r}'_\mathbf{S}\ x'$. (Here $\mathbf{r}' = [\![\alpha']\!]$.) ⊣

add a few more details in the last proof

We present next our main result on the syntactic formulation of equivalence of actions.

Lemma 6.17 *If α and β are equivalent, then for all φ, $\vdash [\alpha]\varphi \leftrightarrow \beta[\varphi]$.*

Proof By induction on φ. Fix a PRE-preserving bisimulation R relating α and β. For φ atomic, our result is easy. The induction steps for $\neg$ and $\wedge$ are trivial. The step for $\Box_A$ is not hard, and so we omit it.

We next check the result for sentences $[\gamma]\varphi$. We need to see that

$$\vdash [\alpha][\gamma]\varphi \leftrightarrow [\alpha'][\gamma]\varphi.$$

For this, it is sufficient by the Composition Axiom to show that $\vdash [\alpha \cdot \gamma]\varphi \leftrightarrow [\alpha' \cdot \gamma]\varphi$. By Lemma 6.15, parts (1) and (5), $\alpha \cdot \gamma \equiv \alpha' \cdot \gamma$. So we are done by the induction hypothesis.

This leaves the step for sentences of the form $\Box^*_\mathcal{B}\varphi$, assuming the result for φ. We use the Action Rule to show that $\vdash [\alpha]\Box^*_\mathcal{C}\varphi \to [\beta]\Box^*_\mathcal{C}\varphi$. We need a functional witness to R; that is, a map $\beta' \mapsto \alpha'$ for all β' such that $\beta \xrightarrow{\mathcal{A}^*} \beta'$, such that $\alpha' R \beta'$ for all these actions. Further, let $\chi_{\beta'}$ be $[\alpha']\Box^*_\mathcal{C}\varphi$. We need to show that for all $A \in C$ and all $\beta' \xrightarrow{A} \beta''$,

a. $\vdash [\alpha']\Box^*_\mathcal{C}\varphi \to [\beta']\varphi$; and

b. If $\beta' \xrightarrow{A} \beta''$, then $\vdash [\alpha']\Box^*_\mathcal{C}\varphi \wedge \text{PRE}(\beta') \to \Box_A[\alpha'']\Box^*_\mathcal{C}\varphi$.

For (a), we know from Lemma 6.14, that $\vdash [\alpha']\Box^*_\mathcal{C}\varphi \to [\alpha']\varphi$. By induction hypothesis on φ, $\vdash [\alpha']\varphi \leftrightarrow [\beta']\varphi$. And this implies (a). For (b), Lemma 6.14 tells us that under the assumptions,

$$\vdash [\alpha']\Box^*_\mathcal{C}\varphi \wedge \text{PRE}(\alpha') \to \Box_A[\alpha'']\Box^*_\mathcal{C}\varphi.$$

The fact that R is a PRE-preserving bisimulation tells us that $\vdash \text{PRE}(\alpha') \leftrightarrow \text{PRE}(\beta')$. This implies (b).

This completes the induction on φ. ⊣

Lemma 6.17 is the key step in the proof of the Substitution Lemma below. In order to formulate substitution, it is convenient to formulate an extension of $\mathcal{L}_1(\mathbf{\Sigma})$ which has additional *sentence variables* and *action variables*. We also need the notion of substitution, of course.

Lemma 6.18 (Substitution Lemma for $\mathcal{L}_1(\mathbf{\Sigma})$.) *Equivalence is preserved under substitutions: Let v and w be maps from sentence variables to sentences (without variables) and action variables to actions (without variables). Assume that $v(x)$ and $w(x)$ are equivalent for all x. Let v^* and w^* be the extensions of v and w to $\mathcal{L}_1(\mathbf{\Sigma})$. Then $v^*(t)$ and $w^*(t)$ are equivalent for all sentences or actions t.*

Proof By induction on t. The cases of atomic sentences and actions are trivial, and the cases for the sentence variables and action variables are immediate. The induction steps for the boolean connectives are easy calculations of propositional equivalence. The induction steps for sentences $\Box_A\varphi$ and $\Box^*_{\mathcal{B}}$ are also easy using the necessitation rules of $\mathcal{L}_1(\mathbf{\Sigma})$.

The key step is for sentences $[\alpha]\varphi$. By induction hypothesis, $v^*(\alpha) \equiv w^*(\alpha)$ and also $v^*(\varphi) \equiv w^*(\varphi)$. So

$$\begin{array}{rcll} v^*([\alpha]\varphi) & = & [v^*(\alpha)]v^*(\varphi) & \\ & = & [w^*(\alpha)]v^*(\varphi) & \text{by Lemma 6.17} \\ & = & [w^*(\alpha)]w^*(\varphi) & \\ & = & w^*([\alpha]\varphi) & \end{array}$$

We also must consider actions. Actions $\sigma\vec{\psi}$: see Lemma 6.15, part (2). The induction step for actions $\alpha \cdot \beta$ follows from part (5) of Lemma 6.15. The steps for $\alpha + \beta$ and *skip* are easy. $\dashv$

Equivalence of basic actions in the fragments We have noted in Lemma 6.15 a few general facts about basic action equivalence. It is natural to ask whether action equivalence is the smallest relation $\equiv$ satisfying all parts of Lemma 6.15. One can ask this in a general way, for all logics irrespective of the underlying action signature. Or one can restrict to the smallest action signatures. Both matters are open.

For example, in the logic of public announcements, we have the iterated announcement law

$$Pub\ \varphi \cdot Pub\ \psi \quad \equiv \quad Pub\ (\varphi \wedge [Pub\ \varphi]\psi) \tag{18}$$

This leads to consequences such as the fact that for atomic p and q, $Pub\ p \cdot Pub\ q \equiv Pub\ q \cdot Pub\ p \equiv Pub(p \wedge q)$, and also that, e.g., $Pub\ \Box_A p \cdot Pub\ \Box_A p \equiv Pub\ \Box_A p$. But we do not know if we obtain a complete proof system for $\equiv$ by taking our logical system for sentential validities and adding the laws of Lemma 6.15 and also (18).

Equivalence in bigger fragments The general notion of equivalence of programs extends to our larger languages. We clearly have the following facts:

1. $\alpha + crash \equiv \alpha$.

2. $\alpha + \beta \equiv \beta + \alpha$.

3. $\alpha + (\beta + \gamma) \equiv (\alpha + \beta) + \gamma$.

4. $\alpha + \alpha \equiv \alpha$.

5. If $\pi \equiv \rho$, then $\pi^* \equiv \rho^*$.

7 Completeness Theorems

In this section, we prove the completeness of our logical systems for $\mathcal{L}_0(\mathbf{\Sigma})$ and $\mathcal{L}_1(\mathbf{\Sigma})$. Recall that the difference between the two languages is that the second has the common-knowledge propositional operators $\Box^*_{\mathcal{B}}$ while the first does not. This difference makes the bigger system much more expressive, and more difficult to study. As it happens, the extension of $\mathcal{L}_1(\mathbf{\Sigma})$ to the full logic $\mathcal{L}(\mathbf{\Sigma})$, the extension via the action iteration operation π^*, leads to a system which is not complete. And due to the common knowledge operators $\Box *_{\mathcal{B}}$, the logic is not *strongly* complete; that is, completeness fails for the notion of validity under hypotheses $T \models \varphi$. The best one can hope for is weak completeness: $\vdash \varphi$ if and only if $\models \varphi$. We prove this in Theorem 7.20. As a result of some preliminary results aimed toward that result, we establish the completeness of our axiomatization of the weaker logic $\mathcal{L}_0(\mathbf{\Sigma})$. The work there is easier because it relies on a translation into modal logic.

In a later section, we show that in contrast to our translation results for $\mathcal{L}_0(\mathbf{\Sigma})$, the larger language $\mathcal{L}_1(\mathbf{\Sigma})$ cannot be translated into $\mathcal{L}$ or even to $\mathcal{L}(\Box^*)$ (modal logic with extra modalities $\Box^*_{\mathcal{B}}$). So completeness results for $\mathcal{L}_1(\mathbf{\Sigma})$ cannot simply be based on translation.

The ideas Our completeness proofs are somewhat involved, and it might help the reader to have a preview of some of the ideas before we get started. For, $\mathcal{L}_0(\mathbf{\Sigma})$, the leading idea is that we can translate the logic back to ordinary modal logic. Then we get completeness by taking any logical system for modal logic and adding whatever principles are needed in order to make the translation. This idea is simple enough, and after looking at a few examples one can see how the translation of $\mathcal{L}_0(\mathbf{\Sigma})$ to modal logic should go. But the formal definition of the translation is complicated. Our definition goes via a *term rewriting system* related to the axioms of the logic, and in particular we'll need to do prove the *termination* of our system. It would have been nice to use some off-the-shelf results of term rewriting theory to get the termination of our system, but this does not seem to be possible. In any case, we prove termination by establishing a *decreasing interpretation* of the system; that is, rewriting a sentence leads to a decrease in an order $<$ that we study at length. Our interpretation is exponential rather than polynomial, and it was found by hand.

When we turn to $\mathcal{L}_1(\mathbf{\Sigma})$, our model is the filtration proof of the completeness of PDL due to Kozen and Parikh [12]. We need to use the Action Rule rather than an Induction Rule, but modulo this difference, the work is similar. We also need to use some of the details on the ordering $<$ mentioned above. That is, even if one were interested in the completeness theorem of $\mathcal{L}_1(\mathbf{\Sigma})$, our proof would still involve the general rewriting apparatus.

We admit that our proofs are complicated and circuitous, so there might be shorter arguments. We don't know of any arguments that avoid term rewriting theory. For this reason, if not for the overall interest of our system itself, we are interested in questions of term rewriting theory related to our logical system. We discuss some of these matters in Section 7.5.

For the record, here are some of the reasons why we feel that the study of our system is complicated:

1. The original statement of axioms such as the Action-Knowledge Axiom is in terms of actions of the form $[\sigma_i \vec{\psi}]$.

2. On the other hand, the Action Rule is best stated in terms of actions which are compo-

sentences φ	p_i	$\neg\varphi$	$\varphi \wedge \psi$	$\varphi \to \psi$	$\Box_A \varphi$	$\Box^*_{\mathcal{B}} \varphi$	$[\alpha]\varphi$	$\textsc{Pre}(\alpha)$
actions α	$\sigma\psi_1, \ldots, \psi_n$	$\alpha \cdot \beta$						

Figure 11: The language $\mathcal{L}_1^+(\mathbf{\Sigma})$

sitions of the actions $[\sigma_i \vec{\psi}]$. In a term-rewriting setting, this point and the previous one work against each other. Our work will be to work with the versions of the axioms that are generalized to the case of all simple actions.

3. Again mentioning term rewriting, we'll need to pick an orientation for the Composition Axiom. This will either be $[\alpha][\beta]\varphi \rightsquigarrow [\alpha \cdot \beta]\varphi$, or $[\alpha \cdot \beta]\varphi \rightsquigarrow [\alpha][\beta]\varphi$.[6] Both alternatives lead to difficulties at various points. We chose the first alternative, and for this reason, we'll need a formulation of the Action-Knowledge Axiom as an infinite scheme.

4. Our language has $+$, but because the axioms are not in general sound for sums, we need to reformulate things to avoid $+$. This is not difficult, but it would have been nice to avoid it.

It is convenient to replace $\mathcal{L}_1(\mathbf{\Sigma})$ by a slightly different language which we call $\mathcal{L}_1^+(\mathbf{\Sigma})$. This new language is shown in Figure 11. For the purposes of this section, we stress the formulation of this as an algebra for a two-sorted signature. Starting with some fixed action signature $\mathbf{\Sigma}$, we construct a two-sorted signature $\Delta = \Delta(\mathbf{\Sigma})$ of terms, obtained in the following way. Let n be the number of simple actions in Σ.

1. Δ has two sorts: s (for sentences) and a (for actions).
2. Each $p \in \mathsf{AtSen}$ is a constant symbol of sort s.
3. $\neg$, $\Box_A$, and $\Box^*_{\mathcal{B}}$ are function symbols of type $s \to s$.
4. $\wedge$ and $\to$ are binary function symbols of type $s \times s \to s$.
5. Each $\sigma \in \Sigma$ is a function symbol of sort $s^n \to a$.
6. *app* is an binary symbol of type $a \times s \to s$.
7. $\textsc{Pre}$ is a function symbol of sort $a \to s$.

The most important addition here is that we have $\textsc{Pre}$ as a first-class part of the syntax; previously it had been an abbreviation. We also add the implication symbol $\to$, but this is only for convenience. Obviously, $\to$ may be dropped from the system. On the other hand, we dropped $+$ (as we mentioned, some axioms are not sound as equations in general if we have $+$). We might as well drop *skip* as well, since it, too, can be translated away.

Incidentally, the fact that $\textsc{Pre}$ is not a symbol of $\mathcal{L}_1(\mathbf{\Sigma})$ makes the issue of translating between $\mathcal{L}_1(\mathbf{\Sigma})$ and $\mathcal{L}_1^+(\mathbf{\Sigma})$ delicate. The reader might wish to formulate a careful translation

[6] Another possibility which we did not explore is to consider rewriting modulo the identity $[\alpha \cdot \beta]\varphi = [\alpha][\beta]\varphi$.

in order to appreciate some of the features of our ordering $<$ which we shall introduce in due course.

We adopt the usual notational conventions that $\wedge$, $\leftrightarrow$, and $\cdot$, are used as infix symbols. We continue our practice of writing $\sigma_i\vec{\psi}$ for what technically would be $\sigma_1(\psi_1,\ldots,\psi_n)$. Also, we write $[\alpha]\psi$ instead of $app(\alpha,\psi)$.

When dealing with $\mathcal{L}_1^+(\mathbf{\Sigma})$, we let α, β, etc., range over terms of sort action, and φ, ψ, χ, etc., range over terms of sort sentence. Finally, we use letters like t and u for terms of either sort (so as to shorten many of our statements). We also will need to adjoin new variables to our signature in order to formulate the a term rewrite system. For this, we let X_a and X_s be sets of new symbols. In order to simplify our notation, we'll use letters like x, y, and z to range over *both* of these sets. That is, we will not notationally distinguish the two sorts of variables. The context will always make it clear what the sort of any given variable is.

Let $\mathcal{L}_1^+(\mathbf{\Sigma},X)$ be the terms built from our signature which now may contain the variables from X. Examples of such terms may be found in Figure 12.

We remind the reader that *ground* terms are terms without variables. For example, a sentence of actions of $\mathcal{L}_1(\mathbf{\Sigma})$ which does not contain $+$ or *skip* may be regarded as a ground term of $\mathcal{L}_1^+(\mathbf{\Sigma})$.

7.1 The rewriting system $\mathcal{R}$ and its interpretation

We recall here the general notion of term rewriting as it applies to $\mathcal{L}_1^+(\mathbf{\Sigma},X)$. Consider a *rewrite rule* (r) of the form $l \rightsquigarrow r$, where l and r are elements of $\mathcal{L}_1^+(\mathbf{\Sigma},X)$. This (r) generates a relation of *immediate rewriting* on $\mathcal{L}_1^+(\mathbf{\Sigma},X)$: we say that t_1 *rewrites to* t_2 *via (r)* if there is a term u with exactly one occurrence of a variable x (of either sort) and a substitution σ such that $u(x \leftarrow l^\sigma) = t_1$ and $v(x \leftarrow r^\sigma) = t_2$. That is, t_1 and t_2 result from u by substituting l and r in for x; however, we need not use l and r literally, but we could as well take a substitution instance of them. We would write $t_1 \xrightarrow{r} t_2$ for this.

Given a set $\mathcal{R}$ of rewrite rules, we write $t_1 \xrightarrow{\mathcal{R}} t_2$ if for some $r \in \mathcal{R}$, $t_1 \xrightarrow{r} t_2$. We naturally consider the transitive closure $\xrightarrow{\mathcal{R}^*}$ of $\xrightarrow{\mathcal{R}}$. We say that t_1 *rewrites to* t_2 *via* $\mathcal{R}$ if they stand in this relation $\xrightarrow{\mathcal{R}^*}$.

We are mostly interested in the rewriting of ground terms. But we need the notion of variables to formulate this, and this is the only reason for introducing variables.

The rewriting system $\mathcal{R}$. Now that we have made this preliminary discussion, we consider the term rewriting system $\mathcal{R}$ shown in Figure 12. We have numbered the rules, and we use this numbering in the sequel.

Remark (r7_α) is an infinite scheme, and in it, α is a term of action sort. But α need not be an action variable. (Indeed, the scheme is only interesting when α is a composition of actions of the form $\sigma_i\vec{\psi}$.) The right side of (r7_α) depends on α, and this is why we use a scheme rather than a single rule. We need to explain what it means. First, the action terms of $\mathcal{L}_1^+(\mathbf{\Sigma},X)$ carry the structure of a frame by taking for each A the smallest relation $\xrightarrow{A}$ such that the following two conditions hold:

1. If $i \xrightarrow{A} j$ in $\mathbf{\Sigma}$, then $\sigma_i\vec{\psi} \xrightarrow{A} \sigma_j\vec{\psi}$.

(r1)	$x \to y$	$\rightsquigarrow$	$\neg(x \wedge \neg y)$
(r2)	$\mathrm{PRE}(\sigma_i y_1 \cdots y_n)$	$\rightsquigarrow$	y_i
(r3)	$\mathrm{PRE}(x \cdot y)$	$\rightsquigarrow$	$\mathrm{PRE}(x) \wedge [x]\mathrm{PRE}(y)$
(r4)	$[x]p$	$\rightsquigarrow$	$\mathrm{PRE}(x) \to p$
(r5)	$[x]\neg y$	$\rightsquigarrow$	$\mathrm{PRE}(x) \to \neg[x]y$
(r6)	$[x](y \wedge z)$	$\rightsquigarrow$	$[x]y \wedge [x]z$
($\mathrm{r}7_\alpha$)	$[\alpha]\Box_A x$	$\rightsquigarrow$	$\mathrm{PRE}(\alpha) \to \bigwedge\{\Box_A[\beta]x : \alpha \xrightarrow{A} \beta\}$
(r8)	$[x][y]z$	$\rightsquigarrow$	$[x \cdot y]z$
(r9)	$x \cdot (y \cdot z)$	$\rightsquigarrow$	$(x \cdot y) \cdot z$

Figure 12: The rewrite system $\mathcal{R}$.

2. if $\alpha \xrightarrow{A} \alpha'$ and $\beta \xrightarrow{A} \beta'$, then $\alpha \cdot \alpha' \xrightarrow{A} \beta \cdot \beta'$.

Moreover, we use the notation $\bigwedge\{\Box_A[\beta]\varphi : \alpha \xrightarrow{A} \beta\}$ to stand for some fixed, but arbitrary, rendering of the conjunction as a nested list of binary conjuncts. The order will not matter, but we shall need an estimate of how many conjuncts there are.

We have a length function ℓ on action terms: $\ell(x) = 0$ for variables x, $\ell(\sigma_i\vec{\psi}) = 1$, and $\ell(\alpha \cdot \beta) = \ell(\alpha) + \ell(\beta)$. For each α, the number of β such that $\alpha \xrightarrow{A} \beta$ is at most $n^{\ell(\alpha)}$, where n is the size of Σ. This is easy to check by induction on $\ell(\alpha)$.

Also on ($\mathrm{r}7_\alpha$): one is tempted to replace this infinite scheme by the finite one

$$\text{(r10)} \quad [\sigma_i\vec{\psi}]\Box_A\varphi \quad \rightsquigarrow \quad \psi_i \to \bigwedge\{\Box_A[\sigma_j\vec{\psi}]\varphi : i \xrightarrow{A} j\}$$

If one did this, one would have to reverse (r8) to read

$$\text{(r11)} \quad [\alpha \cdot \beta]\varphi \quad \rightsquigarrow \quad [\alpha][\beta]\varphi$$

The reason is that (r10) above only allows us to reduce expressions $[\alpha]\Box_A\varphi$ when α is of the form $\sigma_i\vec{\psi}$. The problem is that the normal forms of (r1)–(r6) + (r9) + (r10) + (r11) would include terms like

$$[\sigma_1\vec{\psi^1}][\sigma_2\vec{\psi^2}]\cdots[\sigma_k\vec{\psi^k}]\Box^*_{\mathcal{C}}\varphi,$$

and these are not suitable for use in connection with the Action Rule. We will show about our system (r1) – (r9) that its normal forms includes terms like

$$[(\cdots(\sigma_1\vec{\psi^1} \cdot \sigma_2\vec{\psi^2})\cdots\sigma_k\vec{\psi^k})]\Box^*_{\mathcal{C}}\varphi.$$

Again, this is mainly due to our choice in the direction of (r8).

Our interpretation An *interpretation* of a signature Δ is a Δ-algebra. This is a carrier set for the sentences, a carrier set for the actions, and for each n-ary function symbol f of Δ a function of the appropriate sort. In our setting, both carrier sets will be $N_{\geq 3}$, the set of natural numbers which are at least 3. We shall use a, b, etc., to range over $N_{\geq 3}$ in this section. Our interpretation shown in Figure 13. By recursion on terms $t(x_1, \ldots, x_k)$, we build an interpretation

$$[\![t]\!](x_1, \ldots, x_k) : (N_{\geq 3})^k \to N_{\geq 3}.$$

$[\![p_i]\!]$	$=$	3	$[\![\Box^*_{\mathcal{B}}]\!](a)$	$=$	$a+1$
$[\![\neg]\!](a)$	$=$	$a+1$	$[\![app]\!](a,b)$	$=$	a^b
$[\![\wedge]\!](a,b)$	$=$	$a+b$	$[\![\text{PRE}]\!](a)$	$=$	a
$[\![\rightarrow]\!](a,b)$	$=$	$a+b+3$	$[\![\sigma_i]\!](a_1,\ldots,a_n)$	$=$	$a_1+\cdots+a_n+1$
$[\![\Box_A]\!](a)$	$=$	$a+2$	$[\![\cdot]\!](a,b)$	$=$	a^{b+1}

Figure 13: The interpretation

The interpretation of each function symbol is strictly monotone in each argument. (For example, we check this for $[\![\cdot]\!]$. For fixed a and b, we consider the functions $\lambda x.a^{x+2}$ and $\lambda x.x^{b+2}$. Then if $x < y$, $a^{x+2} < a^{y+2}$ and $x^{b+2} < y^{b+2}$.) So by induction, each $[\![t]\!](x_1,\ldots,x_k)$ is strictly monotone as a function in each argument.

Lemma 7.1 *Let $n = |\Sigma|$. For all action terms α, and all maps ι of variables to $N_{\geq 3}$. $[\![\alpha]\!](\iota(x_1),\ldots,\iota(x_n)) > n^{l(\alpha)}$.*

Proof By induction on α. If α is a variable x, then $\ell(\alpha) = 0$. So $[\![\alpha]\!](\iota(x)) = \iota(x) \geq 3 > 1 = n^{\ell(\alpha)}$. If α is a term of the form $\sigma_i\vec{\psi}$, then $\ell(\alpha) = 1$, and

$$\begin{array}{rcl} [\![\sigma_i\psi_1\cdots\psi_n]\!](\iota(x_1),\ldots,\iota(x_n)) & = & [\![\psi_1]\!](\iota(x_1),\ldots,\iota(x_n)) + \cdots [\![\psi_n]\!](\iota(x_1),\ldots,\iota(x_n)) \\ & > & 3n \\ & > & n^{\ell(\alpha)} \end{array}$$

Now assume our lemma for α and β.

$$\begin{array}{rcll} [\![\alpha\cdot\beta]\!](\iota(x_1),\ldots,\iota(x_n)) & = & ([\![\alpha]\!](\iota(x_1),\ldots,\iota(x_n)))^{[\![\beta]\!](\iota(x_1),\ldots,\iota(x_n))+1} & \\ & > & n^{\ell(\alpha)\cdot(n^{\ell(\beta)+1})} & \\ & > & n^{\ell(\alpha)+\ell(\beta)} & \text{see below} \\ & = & n^{\ell(\alpha\cdot\beta)} & \end{array}$$

In asserting that $\ell(\alpha)\cdot n^{\ell(\beta)} > \ell(\alpha)+\ell(\beta)$, we assume that $n > 1$. If $n = 1$, then our lemma is trivial. ⊣

Proposition 7.2 *If α is an action term of $\mathcal{L}_1^+(\mathbf{\Sigma})$ and $\alpha \xrightarrow{A} \beta$, then $[\![\alpha]\!]$ and $[\![\beta]\!]$ are the same function. In particular, if α and β are ground action terms, then $[\![\alpha]\!] = [\![\beta]\!]$.*

Proof By induction on α. The point is that the interpretation of the actions $\sigma_i\vec{\psi}$ does not use the number i in any way. The rest follows by an easy induction. ⊣

The following proposition is a technical but elementary, result on exponentiation that will be used in Lemma 7.4 below. One should check as well that if we replace 3 by 2 in the statement, then some of the parts are no longer true. This is the reason why our overall interpretation uses $N_{\geq 3}$ rather than N or $N_{\geq 2}$.

Proposition 7.3 *Let $a, b, c \geq 3$.*

1. $2a^b > a + 4$.

2. $a^{b+c} > a^b + a^c$.

3. $a^2 > a + 5$.

4. $a^{b+1} > 3a + 3$.

5. $a^b > (a+1)(b+1)$.

Lemma 7.4 *For each of (r1) - (r9), the interpretation of the left side of the rule is strictly larger than the right side on all tuples of arguments in $N_{\geq 3}$.*

Proof We remind the reader that letters a, b, c, etc., denote elements of $N_{\geq 3}$. We use Proposition 7.3 without explicit mention.

(r1) $[\![x \rightarrow y]\!](a,b) = a + b + 3 > a + b + 2 = [\![\neg(x \wedge \neg y)]\!](a,b)$.

(r2) $[\![\mathrm{PRE}(\sigma_i \vec{y})]\!](a_1,\ldots,a_n) = [\![\sigma_i \vec{y}]\!](a_1,\ldots,a_n) \geq a_i + 1 > [\![y_i]\!](a_1,\ldots,a_n)$.

(r3)

$$\begin{array}{rcl} [\![\mathrm{PRE}(x \cdot y)]\!](a,b) & = & a^{b+1} \\ & > & a + a^b \\ & = & [\![\mathrm{PRE}(x)]\!](a,b) + [\![[x]\mathrm{PRE}(y)]\!](a,b) \\ & = & [\![\mathrm{PRE}(x) \wedge [x]\mathrm{PRE}(y)]\!](a,b) \end{array}$$

(r4) $[\![[x]p]\!](a) = a^2 > a + 5 = [\![\mathrm{PRE}(x) \rightarrow p]\!](a)$.

(r5)

$$\begin{array}{rcl} [\![[x]\neg y]\!](a,b) & = & a^{b+1} \\ & \geq & a^b + 2a^b \\ & > & a + a^b + 4 \\ & = & a + [\![[x]y]\!](a,b) + 4 \\ & = & [\![\mathrm{PRE}(x) \rightarrow \neg[x]y]\!](a,b) \end{array}$$

(r6) $[\![[x](y \wedge z)]\!](a,b,c) = a^{b+c} > a^b + a^c = [\![[x]y \wedge [x]z]\!](a,b,c)$.

($\mathrm{r}7_\alpha$) Let $y_1,\ldots,y_n$ be the free variables of α. Fix $c_1,\ldots,c_n \in N_{\geq 3}$. Let $a = [\![\alpha]\!](c_1,\ldots,c_n)$, and let $b \in N_{\geq 3}$ be arbitrary. Then

$$\begin{array}{rcl} [\![[\alpha]\Box_A x]\!](c_1,\ldots,c_n,b) & = & a^{b+2} \\ & > & a^{b+1} + a^{b+1} \\ & > & a^{b+1} + 3a + 3 \\ & = & a + a(2 + a^b) + 3 \\ & > & a + n^{\ell(a)}(2 + a^b) + 3 \\ & \geq & [\![\mathrm{PRE}(\alpha) \rightarrow \bigwedge\{\Box_A[\beta]x : \alpha \xrightarrow{A} \beta\}]\!](c_1,\ldots,c_n,b) \end{array}$$

Notice that we used Proposition 7.2 and our observation that the size of the conjunction on the right side of ($\mathrm{r}7_\alpha$) is at most $n^{\ell(\alpha)}$.

(r8) $[\![[x][y]z]\!](a,b,c) = a^{b^c} > a^{(b+1)c} = (a^{b+1})^c = [\![[x \cdot y]z]\!](a,b,c)$.

(r9) $[\![x \cdot (y \cdot z)]\!](a,b,c) = a^{b^{c+1}+1} > a^{(b+1)(c+1)} = (a^{b+1})^{c+1} = [\![(x \cdot y) \cdot z]\!](a,b,c)$. ⊣

Definition Let φ and ψ be sentences in $\mathcal{L}_1(\mathbf{\Sigma})$. We regard each as ground terms in $\mathcal{L}_1^+(\mathbf{\Sigma})$.

We write $\varphi > \psi$ if $[\![\varphi]\!] > [\![\psi]\!]$, where our interpretation is the one in Figure 13. We write $\varphi < \psi$ if $\psi > \varphi$, and $\varphi \leq \psi$ if either φ and ψ are identical, or $\varphi < \psi$.

Theorem 7.5 *$\mathcal{R}$ is terminating: there are no infinite sequences*

$$t_1 \quad \xrightarrow{\mathcal{R}} \quad t_2 \quad \xrightarrow{\mathcal{R}} \quad \cdots \quad \xrightarrow{\mathcal{R}} \quad t_n \quad \xrightarrow{\mathcal{R}} \quad t_{n+1} \quad \cdots$$

Proof We recall a standard argument. Assume we had a counterexample, an infinite sequence of rewrites. We may assume without loss of generality that each t_i is a ground term. Recall Lemma 7.4 and our observation that each $[\![t]\!](x_1, \ldots, x_k)$ is strictly monotone as a function in each argument. From this it follows that $[\![t_1]\!] > [\![t_2]\!] > \cdots$, and these are all numbers. So we have a contradiction. ⊣

7.2 Normal forms

We remind the reader that we have formulated the notions of equivalences of simple actions and also equivalence of sentences of $\mathcal{L}_1(\mathbf{\Sigma})$ in Section 6.4. We consider now a translation $\mathsf{trans} : \mathcal{L}_1^+(\mathbf{\Sigma}) \to \mathcal{L}_1(\mathbf{\Sigma})$. The definition is by recursion on the well-order $<$. We take trans to be a homomorphism for all symbols except PRE. For this, we require the following:

$$\begin{array}{rcl} \mathsf{trans}(\textsc{Pre}(\sigma\psi_1, \ldots, \psi_n)) & = & \mathsf{trans}(\psi_i) \\ \mathsf{trans}(\textsc{Pre}(\alpha \cdot \beta)) & = & \langle \mathsf{trans}(\alpha) \rangle \mathsf{trans}(\textsc{Pre}(\beta)) \\ & = & \neg[\mathsf{trans}(\alpha)]\neg\mathsf{trans}(\textsc{Pre}(\beta)) \end{array}$$

The reason we need recursion on $<$ is that $\mathsf{trans}(\textsc{Pre}(\beta))$ figures in $\mathsf{trans}(\textsc{Pre}(\alpha \cdot \beta))$. So we need to know that $\textsc{Pre}(\beta) < \textsc{Pre}(\alpha \cdot \beta)$. Of course, we don't really need $<$ for this; we might as well define $\varphi < \psi$ iff φ has fewer symbols than ψ.

Lemma 7.6 *For all action terms α, $\mathsf{trans}(\textsc{Pre}(\alpha)) = \textsc{Pre}(\mathsf{trans}(\alpha))$.*

Proof By induction on $<$. The important thing again is that when we write $\textsc{Pre}(\mathsf{trans}(\alpha))$ we are using PRE here as a defined symbol. Its recursion equations match the definition of trans. ⊣

Lemma 7.7 *Let α be a ground action term of $\mathcal{L}_1^+(\mathbf{\Sigma})$. Then*

$$\{\mathsf{trans}(\beta) : \alpha \xrightarrow{A} \beta \textit{ in } \mathcal{L}_1^+(\mathbf{\Sigma})\} \quad = \quad \{\beta : \mathsf{trans}(\alpha) \xrightarrow{A} \beta \textit{ in } \mathcal{L}_1(\mathbf{\Sigma})\}.$$

Proof By induction on α. ⊣

Lemma 7.8 *If t_1 and t_2 are ground actions terms of $\mathcal{L}_1^+(\mathbf{\Sigma})$ such that $t_1 \xrightarrow{\mathcal{R}} t_2$, then $\mathsf{trans}(t_1)$ and $\mathsf{trans}(t_2)$ are equivalent actions of $\mathcal{L}_1(\mathbf{\Sigma})$. If t_1 and t_2 are ground sentences, then $\vdash \mathsf{trans}(t_1) \leftrightarrow \mathsf{trans}(t_2)$.*

Proof Let u be a term with exactly one free variable x, and let σ be such that for some i,

$t_1 = u(x \leftarrow l^\sigma)$ and $t_2 = u(x \leftarrow r^\sigma)$. We argue by induction on u. If u is just x, then we examine the rules of the system. All of them pertain to sentences except for (r9), and in the case of (r9) we use Lemma 6.15, part (4). The arguments for the other rules use Lemma 7.6, and we'll give the details for the infinite scheme (r7_α), repeated below:

$$(\mathrm{r}7_\alpha) \qquad [\alpha]\Box_A x \quad \rightsquigarrow \quad \textsc{Pre}(\alpha) \to \bigwedge\{\Box_A[\beta]x : \alpha \xrightarrow{A} \beta\}.$$

Recall that α might have variables. What we need to check here is that ground instances of (r7_α) have equivalent translations. So take a ground sentence of the form $[\alpha]\Box_A\varphi$. Its translation is $[\mathsf{trans}(\alpha)]\Box_A(\mathsf{trans}(\varphi))$. Using the Action-Knowledge Axiom and Lemmas 7.7 and 7.6:

$$\begin{array}{lcl} [\mathsf{trans}(\alpha)]\Box_A(\mathsf{trans}(\varphi)) & \equiv & \textsc{Pre}(\mathsf{trans}(\alpha)) \to \Box_A \bigwedge\{[\beta]\mathsf{trans}(\varphi) : \mathsf{trans}(\alpha) \xrightarrow{A} \beta\} \\ & \equiv & \mathsf{trans}(\textsc{Pre}(\alpha)) \to \bigwedge\{\Box_A[\mathsf{trans}(\beta)]\mathsf{trans}(\varphi) : \alpha \xrightarrow{A} \beta\}) \\ & \equiv & \mathsf{trans}\Big(\textsc{Pre}(\alpha) \to \bigwedge\{\Box_A[\beta]\varphi : \alpha \xrightarrow{A} \beta\}\Big) \end{array}$$

This settles all of the cases in this lemma when u is a variable by itself. In the general case, we need a fact about substitutions. Let t be any term of $\mathcal{L}^+(\mathbf{\Sigma}, X)$. Let σ and τ be two ground substitutions such that $\mathsf{trans}(\sigma(x))$ and $\mathsf{trans}(\tau(x))$ equivalent (actions or sentences) for all variables x. Then $\mathsf{trans}(t^\sigma)$ and $\mathsf{trans}(t^\tau)$ are again equivalent. The argument here is an easy induction on terms of $\mathcal{L}^+(\mathbf{\Sigma}, X)$, and the only step that is not immediate is the one for Pre. ⊣

A *normal form* in a rewriting system is a term which cannot be rewritten in the system. In a terminating system such as our $\mathcal{R}$, one may define a normal form for a term by rewriting as much as possible in some fixed way (for example, by always rewriting in the leftmost possible way). We shall need some information on the normal forms of $\mathcal{R}$, and we shall turn to this shortly.

Definition Let $\mathcal{NF}$ be the smallest set of ground sentences and actions in $\mathcal{L}_1^+(\mathbf{\Sigma})$ with the following properties:

1. Each atomic p belongs to $\mathcal{NF}$.

2. If $\varphi, \psi \in \mathcal{NF}$, then also $\neg\varphi$, $\varphi \wedge \psi$, $\Box_A\varphi$, and $\Box^*_{\mathcal{C}}\varphi$ belong to $\mathcal{NF}$.

3. If $\alpha \in \mathcal{NF}$, $\varphi \in \mathcal{NF}$, and $\mathcal{C} \subseteq \mathcal{A}$, then $[\alpha]\Box^*_{\mathcal{C}}\varphi$ belongs to $\mathcal{NF}$.

4. If $k \geq 0$, if $\vec{\psi}^1, \ldots, \vec{\psi}^k$ is a sequence of sequences of length $n(\mathbf{\Sigma})$ of elements of $\mathcal{NF}$, and if $\sigma_1, \ldots, \sigma_k \in \Sigma$,
$$(\cdots((\sigma_1\vec{\psi}^1 \cdot \sigma_2\vec{\psi}^2) \cdots \sigma_k\vec{\psi}^k) \tag{19}$$
is an action term in $\mathcal{NF}$.

Lemma 7.9 *A term $t \in \mathcal{L}_1^+(\mathbf{\Sigma})$ is a normal form of $\mathcal{R}$ iff $t \in \mathcal{NF}$. In particular,* Pre *and* $\to$ *do not occur in normal forms.*

Proof An induction on $\mathcal{NF}$ shows that all sentences in $\mathcal{NF}$ are normal forms: no rules of $\mathcal{R}$

can apply at any point. Going the other way, we show by induction on φ that if φ is a normal form, then $\varphi \in \mathcal{NF}$. The base case and the induction steps for $\neg$, $\wedge$, $\Box_A$, and $\Box^*_{\mathcal{B}}$ are all easy. Suppose our result holds for φ, and consider a normal form $[\alpha]\varphi$. Then α must be an action of the form $[\sigma_i\vec{\psi}]$ with all ψ_j in normal form; if not, some rule would apply to α and hence to $[\alpha]\varphi$. As for φ, by induction hypothesis it, too, is a normal form. We claim that φ must be of the form (19). This is because all of the other possibilities are reducible in the system. $\dashv$

Lemma 7.10 *A sentence $\varphi \in \mathcal{L}_0(\mathbf{\Sigma})$ is a normal form of $\mathcal{R}$ iff φ is a modal sentence (that is, if φ contains no actions).*

Proof By the easy part of Lemma 7.9, the modal sentences are normal forms, In the more significant direction, one first checks by induction on rewrite sequences that if $\varphi \in \mathcal{L}_0(\mathbf{\Sigma})$ and φ' is reachable from φ by a finite number of rewrites in $\mathcal{R}$, then $\Box^*$ does not appear in φ'. So every normal form of φ is a purely modal sentence. It follows that if φ were a normal form to begin with, then φ would be a modal sentence. $\dashv$

Corollary 7.11 *For every $\mathcal{L}_1(\mathbf{\Sigma})$ there is a normal form $nf(\varphi)$ such that $\vdash \varphi \leftrightarrow nf(\varphi)$. Moreover, $nf(\varphi) \leq \varphi$.*

Proof Regard φ as a term in $\mathcal{L}_1^+(\mathbf{\Sigma})$. Let

$$\varphi = \varphi_1 \quad \rightsquigarrow \quad \varphi_2 \qquad \rightsquigarrow \quad \cdots \rightsquigarrow \varphi_n = nf(\varphi)$$

be the sequence of rewrites which leads to a normal form of φ. By Lemma 7.8, $\vdash \mathsf{trans}(\varphi) \leftrightarrow \mathsf{trans}(nf(\varphi))$. But neither φ nor $nf(\varphi)$ contain $\textsc{Pre}$ or $\rightarrow$, so they are literally equal to their translations. Thus $\vdash \varphi \leftrightarrow nf(\varphi)$, just as desired. $\dashv$

Lemma 7.12 *For all φ and α of $\mathcal{L}_1^+(\mathbf{\Sigma})$:*

1. $\textsc{Pre}(\alpha) < [\alpha]\varphi$.
2. *If $\alpha \xrightarrow{A} \beta$, then $[\beta]\varphi < \neg[\beta]\varphi < [\alpha]\Box^*_{\mathcal{C}}\varphi$.*
3. *If $[\alpha]\Box^*_{\mathcal{C}}\psi$ is a normal form sentence and $\alpha \xrightarrow{A} \beta$, then $\Box_A[\beta]\Box^*_{\mathcal{C}}\psi$ and $[\beta]\Box^*_{\mathcal{C}}\psi$ are again normal forms.*
4. *Every subterm of a normal form is a normal form.*
5. *If α is a normal form action and $\alpha \xrightarrow{A} \beta$, then β is a normal form action as well.*

Proof Part (1) is an easy calculation. In part (2) we use the fact that $[\![\alpha]\!]$ and $[\![\beta]\!]$ are the same number, by Proposition 7.2. We use the fact here (and here only) that $[\![\Box^*]\!](a) = a + 1$ rather than $[\![\Box^*]\!](a) = a$. The remaining parts are also easy using the characterization of normal forms of Lemma 7.9. $\dashv$

7.3 Strong completeness for $\mathcal{L}_0(\mathbf{\Sigma})$

Recall that in the languages $\mathcal{L}_0(\mathbf{\Sigma})$, we do not have the common knowledge operators $\Box^*_{\mathcal{B}}$ or the action iterations π^*. At this point, we can put together several results from our previous work to obtain a completeness theorem for languages of the form $\mathcal{L}_0(\mathbf{\Sigma})$. The overall ideas are: (1) we need a logical system which is strong enough to translate each sentence of $\mathcal{L}_0(\mathbf{\Sigma})$ to a normal form; (2) this normal form will be a purely modal sentence; and (3) the system should be at least as strong as multimodal K.

Proposition 7.13 *Every sentence φ of $\mathcal{L}_0(\mathbf{\Sigma})$ is provably equivalent to a sentence φ^* in which there are no occurrences of $+$, $\cdot$, or* skip.

Theorem 7.14 *The logical system for $\mathcal{L}_0(\mathbf{\Sigma})$ is strongly complete: For all sets $T \subseteq \mathcal{L}_0(\mathbf{\Sigma})$, $T \vdash \varphi$ iff $T \models \varphi$.*

Proof The soundness half being easy, we only need to show that if $T \models \varphi$, then $T \vdash \varphi$.

First, we may assume that the symbols $+$, *crash*, and *skip* do not occur in T or φ. Thus, we may work with $\mathcal{L}_0(\mathbf{\Sigma}) \cap \mathcal{L}_1^+(\mathbf{\Sigma})$. In particular, we have normal forms.

Next, for each χ of $\mathcal{L}_0(\mathbf{\Sigma})$, $\vdash \chi \leftrightarrow nf(\chi)$. As a result, $T \vdash nf(\chi)$ for all $\chi \in T$.

Finally, write $nf(T)$ for $\{nf(\chi) : \chi \in T\}$. By soundness, $nf(T) \models nf(\varphi)$. Since our system extends the standard complete proof system of modal logic, $nf(T) \vdash nf(\varphi)$. So $T \vdash nf(\varphi)$. As we know $\vdash \varphi \leftrightarrow nf(\varphi)$. So we have our desired conclusion: $T \vdash \varphi$. $\dashv$

7.4 Completeness for $\mathcal{L}_1(\mathbf{\Sigma})$

The proof of completeness and decidability of $\mathcal{L}_1(\mathbf{\Sigma})$ is based on the filtration argument for completeness of PDL due to Kozen and Parikh [12]. We show that every consistent φ has a finite model, and that the size of the model is recursive in φ. We shall need to use some results concerning the rewriting system $\mathcal{R}^*$ from Section 6.

Definition For t an action or a sentence, let $s(t)$ be the set of subsentences of t, including t itself (if t is a sentence). This includes all sentences occurring in actions which occur in φ and their subsentences. We define a function $f : \mathcal{NF} \to \mathcal{P}(\mathcal{NF})$ by recursion on the wellfounded relation $<$ as follows:

$$
\begin{array}{lcll}
f(p) & = & & \{p\} \\
f(\neg\varphi) & = & & f(\varphi) \cup \{\neg\varphi\} \\
f(\varphi \wedge \psi) & = & & f(\varphi) \cup f(\psi) \cup \{\varphi \wedge \psi\} \\
f(\Box_A \varphi) & = & & f(\varphi) \cup \{\Box_A \varphi\} \\
f(\Box^*_{\mathcal{B}} \varphi) & = & & f(\varphi) \cup \{\Box^*_{\mathcal{B}} \varphi\} \cup \{\Box_A \Box^*_{\mathcal{B}} \varphi : A \in \mathcal{B}\} \\
f([\alpha]\Box^*_{\mathcal{C}} \varphi) & = & \bigcup & \{s(\Box_A[\beta]\Box^*_{\mathcal{C}}\varphi) \ : \ \alpha \to^*_{\mathcal{C}} \beta \ \& \ A \in \mathcal{C}\} \\
 & & \cup & \bigcup\{f(\psi) \ : \ (\exists\beta)\alpha \to^*_{\mathcal{C}} \beta \ \& \ \psi \in s(\beta)\} \\
 & & \cup & \bigcup\{f(nf(\textsc{Pre}(\beta))) \ : \ \alpha \to^*_{\mathcal{C}} \beta\} \\
 & & \cup & f(\Box^*_{\mathcal{C}}\varphi) \\
 & & \cup & \bigcup\{f(nf([\beta]\varphi)) \ : \ \alpha \to^*_{\mathcal{C}} \beta\}
\end{array}
$$

The definition makes sense because the calls to f on the right-hand sides are all $<$ the arguments on the left-hand sides; see Lemma 7.12.)

Lemma 7.15 *For all $\varphi \in \mathcal{NF}$:*

1. $\varphi \in f(\varphi)$.
2. $f(\varphi)$ *is a finite set of normal form sentences.*
3. *If* $\psi \in f(\varphi)$, *then* $s(\psi) \subseteq f(\varphi)$.
4. *If* $\psi \in f(\varphi)$, *then* $f(\psi) \subseteq f(\varphi)$.
5. *If* $[\gamma]\Box^*_{\mathcal{C}}\chi \in f(\varphi)$, $\gamma \to^*_{\mathcal{C}} \delta$, *and* $A \in \mathcal{C}$, *then* $f(\varphi)$ *also contains* $\Box_A[\delta]\Box^*_{\mathcal{C}}\chi$, $[\delta]\Box^*_{\mathcal{C}}\chi$, $nf(\textsc{Pre}(\delta))$, *and* $nf([\delta]\chi)$.

Proof Part (1) is by cases on φ. The only interesting case is for $[\alpha]\Box^*_{\mathcal{C}}\varphi$, and this is a subsentence of $\Box_A[\alpha]\Box^*_{\mathcal{C}}\varphi$ for any $A \in \mathcal{C}$.

All of other the parts are by induction on φ in the well-order $<$. For part (2), we use Lemma 7.12.

In part (3), we only need to consider φ in normal form. The result is immediate when φ is an atomic sentence p, and the induction steps for $\neg$, $\wedge$, and $\Box_A$ are easy. For $\Box^*_{\mathcal{B}}\varphi$, note that since $\varphi < \Box^*_{\mathcal{B}}\varphi$, our induction hypothesis implies the result for φ; we verify it for $\Box^*_{\mathcal{B}}\varphi$. When ψ is $\Box^*_{\mathcal{B}}\varphi$, then

$$s(\psi) \quad = \quad s(\varphi) \cup \{\Box^*_{\mathcal{B}}\varphi\} \quad \subseteq \quad f(\varphi) \cup \{\Box^*_{\mathcal{B}}\varphi\} \quad \subseteq \quad f(\Box^*_{\mathcal{B}}\varphi).$$

And then, when ψ is $\Box_A\Box^*_{\mathcal{B}}\varphi$ for some $A \in \mathcal{B}$, we have

$$s(\psi) \quad = \quad s(\Box^*_{\mathcal{B}}\varphi) \cup \{\Box_A\Box^*_{\mathcal{B}}\varphi\} \quad \subseteq \quad f(\Box^*_{\mathcal{B}}\varphi).$$

To complete part (3), we consider $[\alpha]\Box^*_{\mathcal{C}}\varphi$. If there is some $\chi < [\alpha]\Box^*_{\mathcal{C}}\varphi$ such that $\psi \in f(\chi)$ and $f(\chi) \subseteq f([\alpha]\Box^*_{\mathcal{C}}\varphi)$, then we argue as follows: by induction hypothesis, $s(\psi) \subseteq f(\chi)$; and by hypothesis $f(\chi) \subseteq f([\alpha]\Box^*_{\mathcal{C}}\varphi)$. This covers all of the cases except for ψ a subsentence of $\Box_A[\beta]\Box^*_{\mathcal{C}}\varphi$. And here $s(\psi) \subseteq s(\Box_A[\beta]\Box^*_{\mathcal{C}}\varphi) \subseteq f([\alpha]\Box^*_{\mathcal{C}}\varphi)$.

In part (4), we only need to consider φ in normal form. The result is immediate when φ is an atomic sentence p. The induction steps for $\neg$, $\wedge$, and $\Box_A$ are easy. For $\Box^*_{\mathcal{B}}\varphi$, note that since $\varphi < \Box^*_{\mathcal{B}}\varphi$, our induction hypothesis implies the result for φ; we verify it for $\Box^*_{\mathcal{B}}\varphi$. The only interesting case is when ψ is $\Box_A\Box^*_{\mathcal{B}}\varphi$ for some $A \in \mathcal{B}$. And in this case

$$f(\psi) \quad = \quad f(\Box^*_{\mathcal{B}}\varphi) \cup \{\Box_A\Box^*_{\mathcal{B}}\varphi\} \quad \subseteq \quad f(\Box^*_{\mathcal{B}}\varphi).$$

To complete part (4), we consider $[\alpha]\Box^*_{\mathcal{C}}\varphi$. If there is some $\chi < [\alpha]\Box^*_{\mathcal{C}}\varphi$ such that $\psi \in f(\chi)$ and $f(\chi) \subseteq f([\alpha]\Box^*_{\mathcal{C}}\varphi)$, then we are easily done by the induction hypothesis. This covers all of the cases except for ψ of the form $[\beta]\Box^*_{\mathcal{C}}\varphi$ or of the form $\Box_A[\beta]\Box^*_{\mathcal{C}}\varphi$. (If ψ is a subsentence of some β, then $f(\psi) \subseteq f([\alpha]\Box^*_{\mathcal{C}}\varphi)$ directly. If ψ is a subsentence of $\Box^*_{\mathcal{C}}\varphi$, then by induction hypothesis, $f(\psi) \subseteq f(\Box^*_{\mathcal{C}}\varphi)$; and directly, $f(\Box^*_{\mathcal{C}}\varphi) \subseteq f([\alpha]\Box^*_{\mathcal{C}}\varphi)$.) For $[\beta]\Box^*_{\mathcal{C}}\varphi$, we use the transitivity of $\to^*_{\mathcal{C}}$ to check that $f([\beta]\Box^*_{\mathcal{C}}\varphi) \subseteq f([\alpha]\Box^*_{\mathcal{C}}\varphi)$. And now the case of $\Box_A[\beta]\Box^*_{\mathcal{C}}\varphi$ follows:

$$f(\Box_A[\beta]\Box^*_{\mathcal{C}}\varphi) \quad = \quad f([\beta]\Box^*_{\mathcal{C}}\varphi) \cup \{\Box_A[\beta]\Box^*_{\mathcal{C}}\varphi\} \quad \subseteq \quad f([\alpha]\Box^*_{\mathcal{C}}\varphi).$$

For part (5), assume that $[\gamma]\Box^*_{\mathcal{C}}\chi \in f(\varphi)$. By part (2), $[\gamma]\Box^*_{\mathcal{C}}\chi$ is a normal form. The definition of f implies that $\Box_A[\delta]\Box^*_{\mathcal{C}}\chi$, $[\delta]\Box^*_{\mathcal{C}}\chi$, $nf(\textsc{Pre}(\delta))$, and $nf([\delta]\chi)$ all belong to $f([\gamma]\Box^*_{\mathcal{C}}\chi)$, and then part (4) tells us that $f([\gamma]\Box^*_{\mathcal{C}}\chi) \subseteq f(\varphi)$. $\dashv$

The set $\Delta = \Delta(\varphi)$ Fix a sentence φ. We set $\Delta = f(\varphi)$ (i.e., we drop φ from the notation). This set Δ is the version for our logic of the Fischer-Ladner closure of φ. Let $\Delta = \{\psi_1, \ldots, \psi_n\}$. Given a maximal consistent set U of $\mathcal{L}_1(\mathbf{\Sigma})$, let

$$U^* \quad = \quad \pm\psi_1 \wedge \cdots \wedge \pm\psi_n,$$

where the signs are taken in accordance with membership in U. That is, if $\psi_i \in U$, then ψ is a conjunct of U^*; but if $\psi_i \notin U$, then $\neg\psi_i$ is a conjunct.

Two (standard) observations are in order. Notice that if $U^* \neq V^*$, then $U^* \wedge V^*$ is inconsistent. Also, for all $\psi \in \Delta$,

$$\vdash \psi \leftrightarrow \bigvee\{W^* : W \text{ is maximal consistent and } \psi \in W\}. \tag{20}$$

and

$$\vdash \neg\psi \leftrightarrow \bigvee\{W^* : W \text{ is maximal consistent and } \neg\psi \in W\}. \tag{21}$$

(The reason is that ψ is equivalent to the disjunction of *all* complete conjunctions which contain it. However, some of those complete conjunctions are inconsistent and these can be dropped from the big disjunction. The others are consistent and hence can be extended to maximal consistent sets.)

Definition The *filtration* $\mathcal{F}$ is the model whose worlds are the equivalence classes $[U]$, where U is a maximal consistent set in the logic for $\mathcal{L}_1(\mathbf{\Sigma})$, and the equivalence relation is $U \equiv V$ iff $U^* = V^*$ (iff $U \cap \Delta = V \cap \Delta$). We set

$$\|p\|_{\mathcal{F}} \quad = \quad \{[U] : p \in U \cap \Delta\}. \tag{22}$$

Furthermore,

$$[U] \xrightarrow{A} [V] \text{ in } \mathcal{F} \quad \text{iff} \quad \text{whenever } \Box_A\psi \in U \cap \Delta, \text{ then also } \psi \in V. \tag{23}$$

Both (22) and (23) are independent of the choice of representatives: we use part (3) of Lemma 7.15 to see that if $\Box_A\chi \in \Delta$, then also $\chi \in \Delta$.

Proposition 7.16 *If $U^* \wedge \Diamond_A V^*$ is consistent, then $[U] \to_A [V]$.*

Proof Assume $\Box_A\psi \in U \cap \Delta$ and toward a contradiction that $\psi \notin V$. Then $\neg\psi \in V$. And since $\psi \in \Delta$, $\vdash V^* \to \neg\psi$. Thus, $\vdash \Diamond_A V^* \to \Diamond_A\neg\psi$, and so $\vdash U^* \wedge \Diamond_A V^* \to \Box_A\psi \wedge \Diamond_A\neg\psi$. Hence $U^* \wedge \Diamond_A V^*$ is inconsistent. $\dashv$

Definition Let $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\psi$ be a normal form. A *good path from $[V_0]$ for $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\psi$* is a path in $\mathcal{F}$

$$[V_0] \quad \to_{A_1} \quad [V_1] \quad \to_{A_2} \quad \cdots \quad \to_{A_{k-1}} \quad [V_{k-1}] \quad \to_{A_k} \quad [V_k]$$

such that $k \geq 0$, each $A_i \in \mathcal{C}$, and such that there exist actions

$$\alpha \;=\; \alpha_0 \quad \to_{A_1} \quad \alpha_1 \quad \to_{A_2} \quad \cdots \quad \to_{A_{k-1}} \quad \alpha_{k-1} \quad \to_{A_k} \quad \alpha_k$$

such that $\textsc{Pre}(\alpha_i) \in V_i$ for all $0 \leq i \leq k$, and $\langle\alpha_k\rangle\psi \in V_k$.

The idea behind a good path comes from considering Lemma 6.10 in $\mathcal{F}$. Of course, the special case of that result would require that $\langle\mathcal{F}, [V_i]\rangle \models \textsc{Pre}(\alpha_i)$ rather than $\textsc{Pre}(\alpha_i) \in V_i$, and similarly for $\langle\alpha_k\rangle\psi$ and V_k. The exact formulation above was made in order that the Truth Lemma will go through for sentences of the form $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\psi$ (see the final paragraphs of the proof of Lemma 7.19).

Lemma 7.17 *Let $[\alpha]\Box^*_{\mathcal{C}}\psi \in \Delta$. If there is a good path from $[V_0]$ for $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\neg\psi$, then $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\neg\psi \in V_0$.*

Proof By induction on the length k of the path. If $k = 0$, then $\langle\alpha\rangle\neg\psi \in V_0$. If $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\neg\psi \notin V_0$, then $[\alpha]\Box^*_{\mathcal{C}}\psi \in V_0$. By Lemma 7.15, part 5, we have $nf([\alpha]\psi) \in V_0$. This is a contradiction.

Assume the result for k, and suppose that there is a good path from $[V_0]$ for $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\neg\psi$ of length $k+1$. Then there is a good path of length k from $[V_1]$ for $\langle\alpha_1\rangle\Diamond^*_{\mathcal{C}}\neg\psi$. Also, $[\alpha_1]\Box^*_{\mathcal{C}}\psi \in \Delta$, by Lemma 7.15, part (5). By induction hypothesis, $\langle\alpha_1\rangle\Diamond^*_{\mathcal{C}}\neg\psi \in V_1$.

If $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\neg\psi \notin V_0$, then $[\alpha]\Box^*_{\mathcal{C}}\psi \in V_0$. V_0 contains $[\alpha]\Box^*_{\mathcal{C}}\psi \wedge nf(\textsc{Pre}(\alpha)) \rightarrow \Box_A[\alpha_1]\Box^*_{\mathcal{C}}\psi$. (That is, this sentence is valid by Lemma 6.14, part 2. Hence it belongs to every maximal consistent set.) So V_0 contains $\Box_A[\alpha_1]\Box^*_{\mathcal{C}}\psi$. This sentence belongs to Δ by Lemma 7.15, part (5). Now by definition of $\xrightarrow{A}$ in $\mathcal{F}$, we see that $[\alpha_1]\Box^*_{\mathcal{C}}\psi \in V_1$. This is a contradiction to our observation at the end of the previous paragraph. ⊣

Lemma 7.18 *If $V_0{}^* \wedge \langle\alpha\rangle\Diamond^*_{\mathcal{C}}\psi$ is consistent, then there is a good path from $[V_0]$ for $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}\psi$.*

Proof For each β such that $\alpha \rightarrow^*_{\mathcal{C}} \beta$, let S_β be the (finite) set of all $[W] \in \mathcal{F}$ such that there is *no* good path from $[W]$ for $\langle\beta\rangle\Diamond^*_{\mathcal{C}}\psi$. We need to see that $[V_0] \notin S_\alpha$; suppose toward a contradiction that $[V_0] \in S_\alpha$. Let

$$\chi_\beta \quad = \quad \bigvee\{W^* : W \in S_\beta\}.$$

Note that $\neg\chi_\beta$ is logically equivalent to $\bigvee\{X^* : [X] \in \mathcal{F} \text{ and } X \notin S_\beta\}$. Since we assumed $[V_0] \in S_\alpha$, we have $\vdash V_0{}^* \rightarrow \chi_\alpha$.

We first claim for β such that $\alpha \rightarrow^*_{\mathcal{C}} \beta$, $\chi_\beta \wedge \langle\beta\rangle\psi$ is inconsistent. Otherwise, there would be $[W] \in S_\beta$ such that $\chi_\beta \wedge \langle\beta\rangle\psi \in W$. Note that by the Partial Functionality Axiom, $\vdash \langle\beta\rangle\psi \rightarrow \textsc{Pre}(\beta)$. But then the one-point path $[W]$ is a good path from $[W]$ for $\langle\beta\rangle\Diamond^*_{\mathcal{C}}\psi$. Thus $[W] \notin S_\beta$, and this is a contradiction. So indeed, $\chi_\beta \wedge \langle\beta\rangle\psi$ is inconsistent. Therefore, $\vdash \chi_\beta \rightarrow [\beta]\neg\psi$.

We next show that for all $A \in \mathcal{C}$ and all β such that $\beta \rightarrow_A \gamma$, $\chi_\beta \wedge \textsc{Pre}(\beta) \wedge \Diamond_A\neg\chi_\gamma$ is inconsistent. Otherwise, there would be $[W] \in S_\beta$ with χ_β, $\textsc{Pre}(\beta)$, and $\Diamond_A\neg\chi_\gamma$ in it. Then $\bigvee\{\Diamond_A X^* : X \notin S_\gamma\}$, being equivalent to $\Diamond_A\neg\chi_\gamma$, would belong to W. It follows that $\Diamond_A X^* \in W$ for some $[X] \notin S_\gamma$. By Proposition 7.16, $[W] \rightarrow_A [X]$. Since $[X] \notin S_\gamma$, there is a good path from $[X]$ for $\langle\gamma\rangle\Diamond^*_{\mathcal{C}}\psi$. But since $\beta \rightarrow_A \gamma$ and W contains $\textsc{Pre}(\beta)$, we also have a good path from $[W]$ for $\langle\beta\rangle\Diamond^*_{\mathcal{C}}\psi$. This again contradicts $[W] \in S_\beta$. As a result, for all relevant A, β, and γ, $\vdash \chi_\beta \wedge \textsc{Pre}(\beta) \rightarrow \Box_A\chi_\gamma$.

By the Action Rule, $\vdash \chi_\alpha \rightarrow [\alpha]\Box^*_{\mathcal{C}}\neg\psi$. Now $\vdash V_0{}^* \rightarrow \chi_\alpha$. So $\vdash V_0{}^* \rightarrow [\alpha]\Box^*_{\mathcal{C}}\neg\psi$. This contradicts the assumption with which we began this proof. ⊣

Lemma 7.19 (Truth Lemma) *Consider a sentence φ, and also the set $\Delta = f(\varphi)$. For all $\chi \in \Delta$ and $[U] \in \mathcal{F}$: $\chi \in U$ iff $\langle \mathcal{F}, [U] \rangle \models \chi$.*

Proof We argue by induction on the wellfounded $<$ that if $\chi \in \Delta$, then: $\chi \in U$ iff $\langle \mathcal{F}, [U] \rangle \models \chi$. The case of χ atomic is trivial. Now assume this Truth Lemma for sentences $< \chi$. Recall that $\Delta \subseteq \mathcal{NF}$, our set of normal forms (see Section 7.2). We argue by cases on χ.

The cases that χ is either a negation or conjunction are trivial.

Suppose next that $\chi \equiv \Box_A \psi$. Suppose $\Box_A \psi \in U$; we show $\langle \mathcal{F}, [U] \rangle \models \Box_A \psi$. Let $[V]$ be such that $[U] \xrightarrow{A} [V]$. Then by definition of $\xrightarrow{A}$, $\psi \in V$. The induction hypothesis applies to ψ, since $\psi < \Box_A \psi$, and since $\psi \in \Delta$ by Lemma 7.15, part (3). So by induction hypothesis, $\langle \mathcal{F}, [V] \rangle \models \psi$. This gives half of our equivalence. Conversely, suppose that $\langle \mathcal{F}, [U] \rangle \models \Box_A \psi$. Suppose towards a contradiction that $\Diamond_A \neg\psi \in U$. So $U^* \wedge \Diamond_A \neg\psi$ is consistent. We use equation (21) and the fact that $\Diamond_A$ distributes over disjunctions to see that $U^* \wedge \Diamond_A \neg\psi$ is logically equivalent to $\bigvee (U^* \wedge \Diamond_A V^*)$, where the disjunction is taken over all V which contain $\neg\psi$. Since $U^* \wedge \Diamond_A \neg\psi$ is consistent, one of the disjuncts $U^* \wedge \Diamond_A V^*$ must be consistent. The induction hypothesis again applies, and we use it to see that $\langle \mathcal{F}, [V] \rangle \models \neg\psi$. By Proposition 7.16, $[U] \xrightarrow{A} [V]$. We conclude that $\langle \mathcal{F}, [U] \rangle \models \Diamond_A \neg\psi$, and this is a contradiction.

For χ of the form $\Box^*_{\mathcal{C}} \psi$, we use the standard argument for PDL (see Kozen and Parikh [12]). This is based on lemmas that parallel Lemmas 7.17 and 7.18. The work is somewhat easier than what we do below for sentences of the form $[\alpha]\Box^*_{\mathcal{C}} \psi$, and so we omit these details.

We conclude with the case when χ is a normal form sentence of the form $[\alpha]\Box^*_{\mathcal{C}} \psi \in \Delta$. Assume that $[\alpha]\Box^*_{\mathcal{C}} \psi \in \Delta$. First, suppose that $[\alpha]\Box^*_{\mathcal{C}} \psi \notin U$. Then by Lemma 7.18, there is a good path from $[U]$ for $\langle \alpha \rangle \Diamond^*_{\mathcal{C}} \neg\psi$. We want to apply Lemma 6.10 in $\mathcal{F}$ to assert that $\langle \mathcal{F}, [U] \rangle \models \langle \alpha \rangle \Diamond^*_{\mathcal{C}} \neg\psi$. Let k be the length of the good path. For $i \leq k$, $\textsc{Pre}(\alpha_i) \in U_i$. Now each $\mathit{nf}(\textsc{Pre}(\alpha_i))$ belongs to Δ by Lemma 7.15, part (5), and is $< [\alpha]\Box^*_{\mathcal{C}} \psi$. So by induction hypothesis, $\langle \mathcal{F}, [U_i] \rangle \models \mathit{nf}(\textsc{Pre}(\alpha_i))$. By soundness, $\langle \mathcal{F}, [U_i] \rangle \models \textsc{Pre}(\alpha_i)$. We also need to check that $\langle \mathcal{F}, [U_k] \rangle \models \langle \alpha_k \rangle \neg\psi$. For this, recall from Lemma 7.12, part (2) that Δ contains $\mathit{nf}(\neg[\alpha_k]\psi) \leq \neg[\alpha_k]\psi < [\alpha]\Box^*_{\mathcal{C}} \psi$. Since the path is good, U_k contains $\langle \alpha_k \rangle \neg\psi$; thus it contains $\neg[\alpha_k]\psi$; and finally it contains $\mathit{nf}(\neg[\alpha_k]\psi)$. By induction hypothesis, $\langle \mathcal{F}, [U_k] \rangle \models \mathit{nf}(\neg[\alpha_k]\psi)$. By soundness, $\langle \mathcal{F}, [U_k] \rangle \models \neg[\alpha_k]\psi$. Thus $\langle \mathcal{F}, [U_k] \rangle \models \langle \alpha_k \rangle \neg\psi$. Now it does follow from Lemma 6.10 that $\langle \mathcal{F}, [U] \rangle \models \langle \alpha \rangle \Diamond^*_{\mathcal{C}} \neg\psi$.

Going the other way, suppose that $\langle \mathcal{F}, [U] \rangle \models \langle \alpha \rangle \Diamond^*_{\mathcal{C}} \neg\psi$. By Lemma 6.10, we get a path in $\mathcal{F}$ witnessing this. The argument of the previous paragraph shows that this path is a good path from $[U]$ for $\langle \alpha \rangle \Diamond^*_{\mathcal{C}} \neg\psi$. By Lemma 7.17, U contains $\langle \alpha \rangle \Diamond^*_{\mathcal{C}} \neg\psi$. This completes the proof. $\dashv$

Theorem 7.20 (Completeness) *For all φ, $\vdash \varphi$ iff $\models \varphi$. Moreover, this relation is decidable.*

Proof By Corollary 7.11, $\vdash \varphi \leftrightarrow \mathit{nf}(\varphi)$. Let φ be consistent. By the Truth Lemma, $\mathit{nf}(\varphi)$ holds at some world in the filtration $\mathcal{F}$. So $\mathit{nf}(\varphi)$ has a model; thus φ has one, too. This establishes completeness. For decidability, note that the size of the filtration is computable in the size of the original φ. (Another proof of decidability: we show in Section 8.1 that $\mathcal{L}_1(\mathbf{\Sigma})$ can be translated into the modal mu-calculus fairly directly, and that logic is decidable. Neither argument gives a good estimate of the complexity.) $\dashv$

7.5 Some other rewriting systems related to $\mathcal{L}_1(\mathbf{\Sigma})$

In this section, we mention a few other results concerning rewriting systems for $\mathcal{L}_1(\mathbf{\Sigma})$.

A rewriting system for the logic of public announcements In the logic of public announcements, the Action-Knowledge Axiom may be stated for simple announcements in a way which avoids the exponential blowup on the right: Rather than (15), we have

$$\vdash [\alpha]\Box_A\varphi \leftrightarrow (\textsc{Pre}(\alpha) \to \Box_A[\alpha]\varphi) \tag{24}$$

This leads to a rewriting system for the logic, obtained by taking (r1) - (r6) + (r8) in Figure 12 and a variant of (r7):

$$\text{(r7')} \quad [x]\Box_A y \quad \rightsquigarrow \quad \textsc{Pre}(x) \to \Box_A[x]y$$

It turns out that there is no reason to adopt (r9) or its converse. The resulting system may be shown to terminate by a *polynomial* interpretation:

$$\begin{array}{lcl@{\qquad}lcl}
[\![p_i]\!] & = & 1 & [\![\Box^*_{\mathcal{B}}]\!](a) & = & a+1 \\
[\![\neg]\!](a) & = & a+1 & [\![app]\!](a,b) & = & 6ab+a+6b+1 \\
[\![\wedge]\!](a,b) & = & a+b+1 & [\![\textsc{Pre}]\!](a) & = & a \\
[\![\to]\!](a,b) & = & a+b+4 & [\![Pub]\!](a) & = & a+1 \\
[\![\Box_A]\!](a) & = & a+1 & [\![\cdot]\!](a,b) & = & 6ab+6a+6b+5
\end{array}$$

Indeed, this interpretation also has the property that $\wedge$ and $\cdot$ are interpreted by functions which are both commutative and associative.

An LPO For background on the lexicographic partial order (LPO), see, e.g., Dershowitz [4] or Plaisted [17]. We construct an LPO using a wellorder $<$ on the symbols in Δ. We first fix a wellfounded relation $<$ on the function symbols. We set *app* to be greater in $<$ than all other symbols. We also set $\sigma_i \equiv \sigma_j$ for all i, j. In all other cases, distinct function symbols are unordered.

Let $<$ be the LPO obtained from these choices.

We remind the reader that our rewriting system $\mathcal{R}$ is given by

$$\begin{array}{lcl}
app(\sigma_i\vec{\psi}, p) & \rightsquigarrow & \psi_i \to p \\
app(\sigma_i\vec{\psi}, \neg\varphi) & \rightsquigarrow & \psi_i \to \neg app(\sigma_i\vec{\psi}, \varphi) \\
app(\sigma_i\vec{\psi}, \varphi \wedge \chi) & \rightsquigarrow & app(\sigma_i\vec{\psi}, \varphi) \wedge app(\sigma_i\vec{\psi}, \chi) \\
app(\sigma_i\vec{\psi}, \Box_A\varphi) & \rightsquigarrow & \psi_i \to \bigwedge\{\Box_A app(\sigma_j\vec{\psi}, \varphi) : \sigma_i \xrightarrow{A} \sigma_j\} \\
app(\alpha \cdot \beta, \varphi) & \rightsquigarrow & app(\alpha, app(\beta, \varphi))
\end{array}$$

Lemma 7.21 *For all rules $\varphi \rightsquigarrow \psi$ of $\mathcal{R}$, $\psi < \varphi$.*

Proof For the first part, recall that $\to < app$. So we use (LPO3): we must check that both ψ_i and p are $< app(\sigma_i\vec{\psi}, p)$. This is by the subterm property.

In the second part, we argue similarly. We must check this time that $app(\sigma_i\vec{\psi}, \varphi) < app(\sigma_i\vec{\psi}, \neg\varphi)$. This is by (LPO1) and the fact that $\varphi < app(\sigma_i\vec{\psi}, \neg\varphi)$; this last is by the subterm property.

Turning to the third part, we check that $app(\sigma_i\vec{\psi},\varphi) < app(\sigma_i\vec{\psi},\varphi\wedge\chi)$ This uses (LPO1) and the fact that $\varphi < \varphi\wedge\psi$, and also $\varphi < app(\sigma_i\vec{\psi},\varphi\wedge\chi)$.

Here is the proof for the fourth part. We first check that for all j, $\Box_A app(\sigma_j\vec{\psi},\varphi) < app(\sigma_i\vec{\psi},\Box_A\varphi)$. Since $\Box_A < app$, we need only show that $app(\sigma_j\vec{\psi},\varphi) < app(\sigma_i\vec{\psi},\Box_A\varphi)$ and use (LPO3). For this, we use (LPO1).

Here is the proof for the last part. Note first that $\beta < \alpha\cdot\beta$, and $\varphi < app(\alpha\cdot\beta,\varphi)$. So by (LPO1), $app(\beta,\varphi) < app(\alpha\cdot\beta,\varphi)$. This with $\alpha < \alpha\cdot\beta$ allows us to conclude again by (LPO1) that $app(\alpha, app(\beta,\varphi)) < app(\alpha\cdot\beta,\varphi)$. ⊣

7.6 Extension to the completeness theorem

We briefly mention extensions of the Completeness Theorem 7.20.

First, consider the case of S5 (or K45) actions. We change our logical system by restricting to these S5 actions, and we add the S5 axioms to our logical system. We interpret this new system on S5 models. It is easy to check that applying an S5 action to an S5 model gives another S5 model. Further, the S5 actions are closed under composition. Finally, if α is an S5 action and $\alpha \to_A \beta$, then β also is an S5 action. These easily imply the soundness of the new axioms. For completeness, we need only check that if we assume the S5 axioms, then the filtration $\mathcal{F}$ from the previous section has the property that each $\xrightarrow{A}$ is an equivalence relation. This is a standard exercise in modal logic (see, e.g., Fagin et al [5], Theorem 3.3.1).

Second, we also have completeness not only for the languages $\mathcal{L}_1(\mathbf{\Sigma})$, but also for each fragment. This is important, because it gives completeness for all of our formalizations of the target logics. To see that we have completeness for fragments takes a certain amount of checking, so we only outline the details. First, our notion of a fragment insures that if φ belongs to a fragment F then so does its normal form. Also if $\varphi\in F$, then $f(\varphi)\subseteq F$; here f is the function giving the analog of the Fischer-Ladner closure of φ. These are the points that take a lot of checking. But once we work through those details, then the completeness theorem which we have extends to all fragments.

Our final extension concerns the move from actions as we have been working them to actions which change the truth values of atomic sentences. If we make this move, then the axiom of Atomic Permanence is no longer sound. However, it is easy to formulate the relevant axioms. For example, if we have an action α which effects the change $p := p\wedge\neg q$, then we would take an axiom $[\alpha]p \leftrightarrow (\textsc{Pre}(\alpha) \to p\wedge\neg q)$. Having made these changes, all of the rest of the work we have done goes through. In this way, we get a completeness theorem for this logic.

Endnotes The special case of Theorem 7.14 when $\mathbf{\Sigma}$ is the signature $\mathbf{\Sigma}_{Pub}$ of public announcements, is due to Plaza [18] and Gerbrandy [7, 8]. The proofs there also go via translation to modal logic. **I need to look up Gerbrandy's papers to make sure that I'm quoting the right one, and also to see if he got completeness for the logic of private announcements, too.**

8 Results on expressive power

In this section, we study a number of expressive power issues related to our logics. The four subsections are for the most part independent.

8.1 Translation of $\mathcal{L}_1(\mathbf{\Sigma})$ into PDL

In this section $\mathbf{\Sigma}$ is an arbitrary action signature. In Section 7.2, we saw normal forms for $\mathcal{L}_0(\mathbf{\Sigma})$ and $\mathcal{L}_1(\mathbf{\Sigma})$. We showed in that section that every sentence in $\mathcal{L}_1(\mathbf{\Sigma})$ is provably equivalent to its normal form, and the normal forms of sentences in $\mathcal{L}_0(\mathbf{\Sigma})$ are exactly the purely modal sentences. This proves that in terms of expressive power, $\mathcal{L}_0(\mathbf{\Sigma})$ is equivalent to $\mathcal{L}_0$, ordinary modal logic.

Further, we can show that $\mathcal{L}_1(\mathbf{\Sigma})$ and indeed the full language $\mathcal{L}(\mathbf{\Sigma})$ is a sublogic of $\mathcal{L}_0^\omega$, the extension of modal logic with countable boolean conjunction and disjunction. The idea is contained in the following clauses:

$$\begin{array}{lcl}
(\Box^*_{\mathcal{B}}\varphi)^t & = & \bigwedge_{\langle A_1,\ldots,A_n\rangle\in\mathcal{B}^*}(\Box_{A_1}\cdots\Box_{A_n}\varphi)^t \\
([\alpha]\Box^*_{\mathcal{B}}\psi)^t & = & \bigwedge_{\langle A_1,\ldots,A_n\rangle\in\mathcal{B}^*}([\alpha]\Box_{A_1}\cdots\Box_{A_n}\psi)^t \\
([\pi]^*\varphi)^t & = & \bigwedge_n([\pi]^n\varphi)^t
\end{array}$$

This observation is improved in our next result.

Theorem 8.1 *Every sentence of $\mathcal{L}_1(\mathbf{\Sigma})$ is equivalent to a sentence of PDL.*

Proof (Sketch) We argue by induction on the wellorder $<$ introduced and studied in Section 7.1. It is sufficient to show that each sentence in the set $\mathcal{NF}$ of normal forms of $\mathcal{L}_1(\mathbf{\Sigma})$ is equivalent to a sentence of PDL. (The normal forms were introduced in Section 7.2, and the reader may wish to look back at Lemma 7.9.) We just give the main induction step.

Suppose that $[\alpha]\Box^*_{\mathcal{C}}\varphi$ is a normal form sentence. Our induction hypothesis implies that each $\psi < [\alpha]\Box^*_{\mathcal{C}}\varphi$ is equivalent to some PDL sentence ψ'. We shall show that $\langle\alpha\rangle\Diamond^*_{\mathcal{C}}$ is itself equivalent to some PDL sentence; hence also $[\alpha]\Box^*_{\mathcal{C}}\varphi$ has this property. For this, we use the semantic equivalent given in Lemma 6.10. Recall first that there are only finitely many β such that $\alpha \xrightarrow{*} \beta$. Let X be the (finite) set of all such β such that $\alpha \xrightarrow{\mathcal{C}} \beta$. That is, the actions β reachable from α by a path labeled by agents in the set $\mathcal{C}$. We consider X as a sub-action structure of $\mathbf{\Omega}$.

We consider $\mathcal{C}\cup X$, and we assume that this union is disjoint. We consider the set $(\mathcal{C}\cup X)^*$ of all finite words on this set. We are interested in finite $\mathcal{C}$-labeled paths through X beginning at α and ending at an arbitrary element of X. At this point, we shall develop our proof only by example.

Suppose that $\mathcal{C} = \{A, B\}$ and that $X = \{\alpha,\beta\}$, with $\alpha \xrightarrow{A} \alpha$, $\alpha \xrightarrow{B} \beta$, and $\beta \xrightarrow{A} \alpha$. Then one of the paths of interest would be

$$\alpha \quad A \quad \alpha \quad B \quad \beta \quad A \quad \alpha$$

(Note that this corresponds to $\alpha \xrightarrow{A} \alpha \xrightarrow{B} \beta \xrightarrow{A} \alpha$.) We are only interested in paths that respect the structure of $\mathbf{\Omega}$. By Kleene's Theorem, the set P of finite paths of this type beginning at our fixed action α is a regular language on $\mathcal{C} \cup X$. In our example, P is given by the regular expression

$$((\alpha(A\alpha)^*(B\beta A))^*(\epsilon + B\beta).$$

With each such regular expression we associate a PDL program that represents it. In our example, we would have

$$(?\text{PRE}'(\alpha);((\langle A\rangle;?\text{PRE}'(\alpha))^*;(\langle B\rangle;?\text{PRE}'(\beta);\langle A\rangle))^*;(?\mathsf{True} + \langle B\rangle;?\text{PRE}'(\beta)).$$

(The notation $\text{PRE}'(\gamma)$ means the PDL translation of $\text{PRE}(\gamma)$; such a PDL sentence exists since $\text{PRE}(\gamma) < [\alpha]\Box^*_{\mathcal{C}}\varphi$.) Again, this is a PDL program $\pi = \pi(X,\alpha)$ whose denotation $[\![\pi]\!]_{\mathbf{S}}$ in a state model $\mathbf{S}$ is the set of pairs (s,t) of states such that there is a path

$$s = s_0 \quad \rightarrow_{A_1} \quad s_1 \quad \rightarrow_{A_2} \quad \cdots \quad \rightarrow_{A_{k-1}} \quad s_{k-1} \quad \rightarrow_{A_k} \quad s_k$$

and also a sequence of actions of the same length k,

$$\alpha \;=\; \alpha_0 \quad \rightarrow_{A_1} \quad \alpha_1 \quad \rightarrow_{A_2} \quad \cdots \quad \rightarrow_{A_{k-1}} \quad \alpha_{k-1} \quad \rightarrow_{A_k} \quad \alpha_k$$

such that each $A_i \in \mathcal{C}$, each $s_i \in [\![\text{PRE}(\alpha_i)]\!]_{\mathbf{S}}$ for all $0 \leq i \leq k$, and finally such that $s_k = t$. Consider now the PDL sentence $\langle\pi\rangle(\langle\alpha_k\rangle\varphi)'$. (As above, $(\langle\alpha_k\rangle\varphi)'$ means the PDL translation of $\langle\alpha_k\rangle\varphi$. This exists because $[\alpha_k]\varphi < [\alpha]\Box^*_{\mathcal{C}}\varphi$; see Lemma 7.9.) A state s satisfies $\langle\pi\rangle(\langle\alpha_k\rangle\varphi)'$ in $\mathbf{S}$ iff there is some $(s,t) \in [\![\pi]\!]_{\mathbf{S}}$ such that $t \in [\![\langle\alpha_k\rangle\varphi]\!]_{\mathbf{S}}$. Putting together our description of $[\![\pi]\!]_{\mathbf{S}}$ with this, we see that $(s,t) \in [\![\pi]\!]_{\mathbf{S}}$ iff $s \in [\![\langle\alpha\rangle\Diamond^*_{\mathcal{C}}]\!]_{\mathbf{S}}$; see Lemma 6.10. ⊣

At this point, we know that $\mathcal{L}_1(\mathbf{\Sigma})$ is a sublogic of PDL. So it is interesting to ask whether this extends to the full logic $\mathcal{L}(\mathbf{\Sigma})$; recall that this last language has the operation of *program* iteration π^*. It turns out that $\mathcal{L}(\mathbf{\Sigma})$ lacks the finite model property even when $\mathbf{\Sigma}$ is as simple as the signature of public announcements and indeed when the only sentenced announced publically and repeatedly is $\Diamond\mathsf{True}$ (see Section 8.4 below). Since PDL has the finite model property, we see that $\mathcal{L}(\mathbf{\Sigma})$ is not a sublogic of PDL. In fact, it also shows that $\mathcal{L}(\mathbf{\Sigma})$ is not even a sublogic of the modal mu-calculus.

8.2 $\mathcal{L}_1(\mathbf{\Sigma}_{\mathbf{pub}})$ is more expressive than $\mathcal{L}_1$

Recall that $\mathcal{L}_1$ in this paper is multi-agent modal logic together with the common-knowledge operators $\Box^*_{\mathcal{B}}$ for sets of agents. Our main result here is that $\mathcal{L}_1$ is strictly weaker than $\mathcal{L}_1(\mathbf{\Sigma}_{\text{pub}})$, the logic obtained by adding public announcements to $\mathcal{L}_1$.

We define a *rank* $|\varphi|$ on sentences from $\mathcal{L}_1(\mathbf{\Sigma}_{\text{pri}})$. Let $|p| = 0$ for p atomic, $|\neg\varphi| = |\varphi|$, $|\varphi \wedge \psi| = \max(|\varphi|, |\psi|)$, $|\Diamond_A\varphi| = 1 + |\varphi|$, for all $A \in \mathcal{A}$, and $|\Diamond^*_{\mathcal{B}}\varphi| = 1 + |\varphi|$ for all $\mathcal{B} \subseteq \mathcal{A}$.

Games for $\mathcal{L}_1$ The main technique in the proof is an adaptation of Fraisse-Ehrenfeucht games to the setting of modal logic. Let $(\mathbf{S}, s)$ and $(\mathbf{T}, t)$ be states; i.e., model-world pairs. By recursion on the natural number n we define a game $G_n((\mathbf{S}, s), (\mathbf{T}, t))$. For $n = 0$, II immediately wins if the following holds: for all $p \in \mathsf{AtSen}$, $(\mathbf{S}, s) \models p$ iff $(\mathbf{T}, t) \models p$. And if s and t differ on some atomic sentence, I immediately wins. Continuing, here is how we define $G_{n+1}((\mathbf{S}, s), (\mathbf{T}, t))$. As in the case of G_0, we first check if s and t differ on some atomic sentence. If they do, then I immediately wins. Otherwise, the play continues. Now I can make two types of moves.

1. A $\Diamond_A$-move: I has a choice of playing from $\mathbf{S}$ or from $\mathbf{T}$, and also some agent A. If I chooses $\mathbf{S}$, then I continues by choosing some s' such that $s \xrightarrow{A} s'$ in $\mathbf{S}$. Then II replies with some $t' \in T$ such that $t \xrightarrow{A} t'$. Of course, if I had chosen in $\mathbf{T}$, then II would have chosen in S. Either way, points s' and t' are determined, and the two players then play $G_n((\mathbf{S}, s'), (\mathbf{T}, t'))$.

2. A $\Diamond^*_{\mathcal{B}}$-move: I plays by selecting $\mathbf{S}$ (or $\mathbf{T}$, but we ignore this symmetric case below), and some set $\mathcal{B}$ of agents, and then I continues by playing some s' (say) reachable from s in the reflexive-transitive closure $\xrightarrow{\mathcal{B}^*}$ of $\xrightarrow{\mathcal{B}}$; II responds with a point t' in the other model, $\mathbf{T}$, which is similarly related to t.

We write $(\mathbf{S}, s) \sim_n (\mathbf{T}, t)$ if II has a winning strategy in the game $G_n((\mathbf{S}, s), (\mathbf{T}, t))$. It is easy to check that by induction on m that if $(\mathbf{S}, s) \sim_n (\mathbf{T}, t)$ and $m < n$, then $(\mathbf{S}, s) \sim_m (\mathbf{T}, t)$.

Proposition 8.2 *If $(\mathbf{S}, s) \sim_n (\mathbf{T}, t)$, then for all φ with $|\varphi| \leq n$, $(\mathbf{S}, s) \models \varphi$ iff $(\mathbf{T}, t) \models \varphi$.*

The proof is standard, except perhaps for the easy extra step for $\Diamond^*$ moves.

We have two results that show that public announcements add expressive power to $\mathcal{L}_1$. The first is for one agent on arbitrary models, and the second is for two agents on equivalence relations.

In the first result, let $\mathcal{A}$ be a singleton $\{A\}$. We drop the A from the notation.

Theorem 8.3 *The $\mathcal{L}_1(\mathbf{\Sigma}_{pub})$ sentence $\langle \mathrm{Pub}\ p\rangle \Diamond^* q$ is not expressible in $\mathcal{L}_1$, even by a set of sentences.*

Proof Fix a number N. We first show that $\langle Pub\ p\rangle \Diamond^* q$ is not expressible by a single sentence of $\mathcal{L}_1$ of rank N. Let A_N be the cycle

$$a_0 \to a_1 \to \cdots \to a_{n+2} \to a_{n+3} \to \cdots a_{2n+4} = a_0.$$

We set p true everywhere except a_{n+2} and q true only at a_0.

Announcing p means that we delete a_{n+2}. So $A_N(Pub\ p)$ splits into two disjoint pieces. this means that in $A_N(Pub\ p)$, a_1 does not satisfy $\Diamond^* q$. But a_{n+3} does satisfy it.

We show that a_1 and a_{n+3} agree on all sentences in $\mathcal{L}_1$ of rank $\leq n$. For this, we show that II has a winning strategy in the N-round game on between (A_N, a_{n+1}) and (A_N, a_{3n+1}). Then we appeal to Proposition 8.2. II's strategy is as follows: if I ever makes a $\Diamond^*$ move, II should make a move on the other side to the exact same point. (Recall that A_N is a cycle.) Thereafter, II should mimic I's moves exactly. Since the play will end with the same point in the two structures, II wins. But if I never makes a $\Diamond^*$ move, the play will consist of n $\Diamond$-moves. II should simply make the same moves in the appropriate structures. Since the original points a_1 and a_{n+3} are $n+1$ steps from a_0 and a_{n+2}, II will win the play in this case.

So at this point we conclude that $\langle Pub\ p\rangle \Diamond^* q$ is not expressible by a single sentence of $\mathcal{L}_1$ of rank N. We conclude by extending this to show that $\langle Pub\ p\rangle \Diamond^* q$ is not expressible by any set of sentences of $\mathcal{L}_1$. Suppose towards a contradiction that that $\langle Pub\ p\rangle \Diamond^* q$ were equivalent to the set $T \subseteq \mathcal{L}_1$. Consider the following models $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} = \oplus(A_n, a_1^n)$; i.e., the disjoint union of the models $\oplus A_n$, where we identify all points a_1^n. We also rename the common a_1^n point to be

a, and we take this as the distinguished point of the state $(\mathcal{A}, a)$. Similarly, $\mathcal{B} = \oplus(A_n, a^n_{n+3})$. For clarity we'll rename each point a^m_j $\mathcal{B}$ to be b^m_j. So now the distinguished point of $\mathcal{B}$ is called b. Clearly $(\mathcal{A}, a) \models \neg\langle Pub\ p\rangle\Diamond^* q$ and $(\mathcal{B}, b) \models \langle Pub\ p\rangle\Diamond^* q$.

By definition of T, $(\mathcal{B}, b) \models \varphi$ for all $\varphi \in T$. Let $\varphi \in T$ be such that $(\mathcal{A}, a) \models \neg\varphi$. Let $m = |\varphi|$. Let

$$\mathcal{C} \quad = \quad \oplus(\{(A_n, a^n_1) : n \neq m\} \cup \{(A_m, a^m_{m+3})\}).$$

Notice that we switched exactly one of the identified points. Just as before, we'll rename the points of $\mathcal{C}$ and call the overall distinguished point c. It is easy to check that $(\mathcal{C}, c) \models \langle Pub\, p\rangle\Diamond^* q$. And since this sentence is equivalent to T, we also have $(\mathcal{C}, c) \models \varphi$. But we now show that $(\mathcal{A}, a)$ and $(\mathcal{C}, c)$ agree on all sentences of rank $\leq m$. This implies that $(\mathcal{A}, a) \models \varphi$, giving the needed contradiction.

Here is a winning strategy for *II* in the m-round game between $(\mathcal{A}, a)$ and $(\mathcal{C}, c)$. If *I* opens with anything besides c^m_{m+3} or a^m_1, *II* should play the corresponding point on the other side and thereafter play in the obvious way. If *I* opens with c^m_{m+3}, *II* should play with a^m_1 and thereafter play with basically the same strategy as in the first part of this theorem; similarly, if *I* opens with a^m_1, *II* should reply c^m_{m+3} and thereafter play via the strategy above. $\dashv$

Partition models Often in epistemic logic one is concerned with models in which every accessibility relation $\xrightarrow{A}$ is an equivalence relation. We can obtain a version of Theorem 8.3 which shows that even on this smaller class of models, public announcements add expressive power to $\mathcal{L}_1$. However, we must use two agents: with one agent and an equivalence relation $\Box^*\varphi$ is equivalent to $\Box\varphi$. Thus every sentence in $\mathcal{L}_1(\mathbf{\Sigma}_{\text{pub}})$ on one-agent equivalence relations is equivalent to a purely modal sentence.

Theorem 8.4 *The $\mathcal{L}_1(\mathbf{\Sigma}_{pub})$ sentence $\langle \text{Pub}\ p\rangle\Diamond^*_{A,B} q$ is not expressible by any set of sentences of $\mathcal{L}_1$, even on the class of models in which $\xrightarrow{A}$ and $\xrightarrow{B}$ are equivalence relations.*

Proof Again we fix a number N and first show that $\langle Pub\ p\rangle\Diamond^* q$ is not expressible by a single sentence of $\mathcal{L}_1$ of rank N. Let n be the smallest even number strictly larger than N. We consider a model C_n which is a cycle of $5n$ points $a_1, \ldots, a_{5n}$ arranged as follows:

$$a_1 \xleftrightarrow{A} a_2 \xleftrightarrow{B} a_3 \qquad \cdots \qquad a_{5n-1} \xleftrightarrow{A} a_{5n} \xleftrightarrow{B} a_1$$

Since n is even, for $1 \leq i \leq 5$, the connection is $a_{in-1} \xrightarrow{A} a_{in} \xrightarrow{B} a_{in+1}$. (We are taking subscripts modulo $5n$ here.)

We also specify that p is true at all points except a_1 and a_{2n+1}, and q is true *only* at a_{4n+1}.

Note that this model C_n does have the property that $\xrightarrow{A}$ and $\xrightarrow{B}$ are equivalence relations. In addition, $(\xrightarrow{A})^*$ is the same relation as $\xrightarrow{A}$. So $\Box^*_A\varphi$ is equivalent to $\Box_A\varphi$ (and similarly for B). Finally, $(\xrightarrow{A} \cup \xrightarrow{B})^*$ is the universal relation. So if any point whatsoever satisfies a sentence $\Box^*_{A,B}\varphi$, then all points satisfy it.

When p is publically announced, a_1 and a_{2n+1} disappear. The cycle breaks into two disconnected components, and in the resulting model $C_n(Pub\ p)$, $a_{n+1} \models \neg\Diamond^*_{A,B} q$. On the other hand, a_{3n+1} and a_{4n+1} are in the same connected component after the update, so in $C_n(Pub\ p)$, $a_{3n+1} \models \Diamond^*_{A,B} q$. Returning to C_n, the points a_{n+1} and a_{3n+1} differ on our sentence $\langle Pub\ p\rangle\Diamond^*_{A,B} q$.

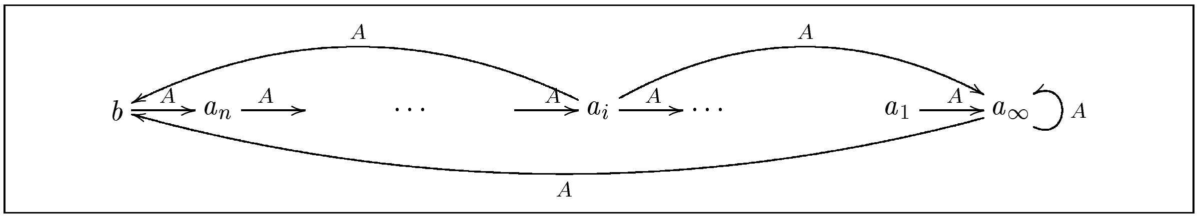

We continue by showing that II has a winning strategy in the N-round game on C_N from a_{n+1} and a_{3n+1}. If I ever makes a $\Diamond^*$ move, then II should move to the same point and thereafter mimic I perfectly. Thus we may assume that I never makes any $\Diamond^*$ moves. In this case, I will never move either point to a_1, a_{2n+1}, or a_{4n+1}. In other words, all points in the play will satisfy $p \wedge \neg q$. So as long as I plays $\Diamond$-moves, II can follow arbitrarily. Once again, since the game goes for N rounds, this will be a winning strategy.

At this point we know that $\langle Pub\ p\rangle\Diamond^* q$ is not expressible by a single sentence of $\mathcal{L}_1$ of rank N. We use the same idea as in Theorem 8.3 to show that $\langle Pub\ p\rangle\Diamond^* q$ is not expressible by any set of sentences of $\mathcal{L}_1$. In fact, virtually the same proof goes through. $\dashv$

8.3 $\mathcal{L}_1(\mathbf{\Sigma_{pri}})$ is more expressive than $\mathcal{L}_1(\mathbf{\Sigma_{pub}})$

It is natural to assume that *privacy* adds expressive power to logics of communication. The next result is the first result we know that establishes this.

Theorem 8.5 *The $\mathcal{L}_1(\mathbf{\Sigma}_{pri})$-sentence $\chi \equiv [\mathrm{Pri}^A\ p]\Diamond_A^*\Diamond_B\neg p$ cannot be expressed by any set of sentences in $\mathcal{L}_1(\mathbf{\Sigma}_{pub})$.*

Proof Let G_n be the model described as follows: We begin with a cycle in $\xrightarrow{A}$:

$$b \xrightarrow{A} a_n \xrightarrow{A} a_{n-1} \xrightarrow{A} \cdots \xrightarrow{A} a_2 \xrightarrow{A} a_1 \xrightarrow{A} a_\infty \xrightarrow{A} b \tag{25}$$

We add edges $a_i \xrightarrow{A} b$ for all i (including $i = \infty$), and also $x \xrightarrow{A} a_\infty$ for all x (again including $x = a_\infty$). This gives the picture above. The only $\xrightarrow{B}$ edge is $a_1 \xrightarrow{B} b$. The atomic sentence p is true at all points except b.

The first thing to note is that in $G_n(Pri^A\ p)$, $a_i \models \chi$ for all $i < \infty$. The relevant path is $a_i \xrightarrow{A} \cdots \xrightarrow{A} a_1 \xrightarrow{B} b$; the important point is that since the announcement was private, the edge $a_1 \xrightarrow{B} b$ survives the update. On the other hand, in this same model $G_n(Pri^A\ p)$, $a_\infty \models \neg\chi$. This is because the only path in the original model from a_∞ to b is

$$a_\infty \xrightarrow{A} b \xrightarrow{A} a_n \xrightarrow{A} \cdots \xrightarrow{A} a_1 \xrightarrow{B} b,$$

and the edge $a_\infty \xrightarrow{A} b$ is lost in the update.

Suppose towards a contradiction that χ were equivalent to $\varphi \in \mathcal{L}_1(\mathbf{\Sigma}_{\mathrm{pub}})$. Let $i = |\varphi|$, and let $n = i+1$. As we know from our discussion of χ, $\langle G_n, a_n\rangle \models \chi$ and $\langle G_n, a_\infty\rangle \models \neg\chi$. However, this contradicts the claim below.

Claim Assume that $1 < j \leq n$, $\varphi \in \mathcal{L}_1(\mathbf{\Sigma}_{\mathrm{pub}})$ and $|\varphi| < j$. Then in G_n, $a_j \models \varphi$ iff $a_\infty \models \varphi$.

The proof is by induction on φ. For $\varphi = p$, the result is clear, as are the induction steps for $\neg$ and $\wedge$. For $\Diamond_A\varphi$, suppose that $a_j \models \Diamond_A\varphi$. Either $a_\infty \models \varphi$, in which case $a_\infty \models \Diamond_A\varphi$, or else $a_{j-1} \models \varphi$. In the latter case, by induction hypothesis, $a_\infty \models \varphi$; whence $a_\infty \models \Diamond_A\varphi$. The converse is similar.

The case of $\Diamond_B\varphi$ is trivial: $a_j \models \neg\Diamond_B\varphi$ and $a_\infty \models \neg\Diamond_B\varphi$.

For $\Diamond^*_A\varphi$, note that since we have a cycle (25) containing all points, the truth value of $\Diamond^*_A\varphi$ does not depend on the point. The cases of $\Diamond^*_B\varphi$ and $\Diamond^*_{AB}\varphi$ are similar.

For $[Pub\ \varphi]\psi$, assume the result for φ and ψ, and let $|[Pub\ \varphi]\psi| < j$. Then also $|\varphi| < j$ and $|\psi| < j$. Let $H = \{x : x \models \varphi\}$ be the updated model, and recall that $\langle G_n, x\rangle \models [Pub\ \varphi]\psi$ iff $x \in H$ and $\langle H, x\rangle \models \psi$. We have two cases: First, $H = G_n$. Then $\langle G_n, x\rangle \models [Pub\ \varphi]\psi$ iff $\langle G_n, x\rangle \models \psi$. So we are done by the induction hypothesis.

The other case is when there is some $x \notin H$. If $a_k \notin H$ for some $k \geq j$ or for $k = \infty$, then *all* these a_k do not belong to H. In particular, neither a_j nor a_∞ belong. And so both a_j and a_∞ satisfy $[Pub\ \varphi]\psi$. If $b \notin H$, then H is bisimilar to a one-point model. This is because every $a_i \in H$ would have some $\xrightarrow{A}$ -successor in H (e.g., a_∞), and there would be no $\xrightarrow{B}$ edges. So we assume $b \in H$. Thus $a_i \notin H$ for some $i < j$. Let k be least so that for $k \leq l \leq \infty$, $a_l \models \varphi$. Then $1 < k \leq j$. Let $A_{\geq k} = \{a_l : k \leq l \leq \infty\}$. The submodels generated by a_j and a_∞ contain the same worlds: all worlds in $A_{\geq k}$ and b. We claim that $(A_{\geq k} \times A_{\geq k}) \cup \{\langle b, b\rangle\}$ is a bisimulation on H. The verification here is easy.

So in H, a_j and a_∞ agree on the sentences of $\mathcal{L}_1(\mathbf{\Sigma}_{\text{pub}})$. (They agree on all sentences in all our languages, by Proposition 4.2; indeed they agree on all sentences of infinitary modal logic). In particular, $\langle H, a_j\rangle \models \psi$ iff $\langle H, a_\infty\rangle \models \psi$. This concludes the claim.

So at this point we conclude that $[Pri^A\ p]\Diamond^*_A\Diamond_B\neg p$ is not expressible by a single sentence of $\mathcal{L}_1(\mathbf{\Sigma}_{\text{pub}})$. of rank N. We present a lemma which allows us, in certain circumstances, to do the following: from the existence of a sentence in a language $\mathcal{L}_1$ which is not equivalent to any sentence in a language $\mathcal{L}_0$ infer that there exists a sentence in $\mathcal{L}_1$ not even equivalent to any *theory* in $\mathcal{L}_0$.

Lemma 8.6 *Let $\mathcal{L}_0$ be a language included in $\mathcal{L}_1(\mathbf{\Sigma}_{pri})$, and let ψ be a sentence of $\mathcal{L}_1(\mathbf{\Sigma}_{pri})$. Assume that for each n we have models F_n and G_n with some worlds $f_n \in F_n$ and $g_n \in G_n$ such that $\langle F_n, f_n\rangle$ satisfies $\neg\psi$, $\langle G_n, g_n\rangle$ satisfies ψ, and $\langle F_n, f_n\rangle$ and $\langle G_n, g_n\rangle$ agree on all sentences in $\mathcal{L}_0$ of rank $\leq n$. Then $\Diamond_A\psi$ is not equivalent with any theory in $\mathcal{L}_0$.*

Proof For a sequence of model-world pairs $\langle H_n, h_n\rangle$, $n \in D \subseteq \omega$, we let $\bigoplus_{n\in D}(H_n, h_n)$ be a model-world pair defined as follows. Let h be a new world. Take disjoint copies of the H_n's and add an A-arrow from h to each h_n. All other arrows are within the H_n's and stay the same as in H_n. No atomic sentences are true at h. Atomic sentences true in the worlds belonging to the copy of H_n in $\bigoplus_{n\in D}(H_n, h_n)$ are precisely those true at the corresponding worlds of H_n.

Let F be $\bigoplus_{n\in\omega}(F_n, f_n)$ with the new world denoted by f. Define also F^m, for $m \in \omega$, to be $\bigoplus_{n\in\omega}(H_n, h_n)$ with the new world f^m where $H_m = G_m$, $h_m = g_m$ and for all $n \neq m$, $H_n = F_n$ and $h_n = f_n$

Now assume towards a contradiction that $\Diamond_A\psi$ is equivalent with a theory Φ in $\mathcal{L}_0$. Clearly $\Diamond_A\psi$ fails in $\langle F, f\rangle$. Thus some sentence $\varphi \in \Phi$ fails in $\langle F, f\rangle$. On the other hand, each $\langle F^m, f^m\rangle$ satisfies $\Diamond_A\psi$, whence $\langle F^m, f^m\rangle$ satisfies φ. Let $m_0 = |\varphi|$. The following claim shows that both $\langle F, f\rangle$ and $\langle F^{m_0}, f^{m_0}\rangle$ make φ true or both of them make it false, which leads to a contradiction.

Claim Let φ be a sentence in Φ of rank $\leq m$. Let H_n, K_n, $n \in D$, with $h_n \in H_n$ and $k_n \in K_n$ be models such that $\langle H_n, h_n\rangle$ and $\langle K_n, k_n\rangle$ agree on sentences in Φ of rank $\leq m$. Then $\langle \bigoplus_n (H_n, h_n), h\rangle$ and $\langle \bigoplus_n (K_n, k_n), k\rangle$ agree on φ.

This claim is proved by induction on complexity of φ. It is clear for atomic sentences. The induction steps for boolean connectives are trivial. A moment of thought gives the induction step for $\Diamond$ and $\Diamond^*$ with various subscripts. It remains to consider the case when $\varphi = [\varphi_1]_A\varphi_2$. (The cases when $\varphi = [\varphi_1]_B\varphi_2$ and $\varphi = [\varphi_1]_{AB}\varphi_2$ are similar.) Fix H_n, K_n, $h_n \in H_n$, $k_n \in K_n$, with $n \in D$, such that $\langle H_n, h_n\rangle$ and $\langle K_n, k_n\rangle$ agree on sentences in Φ of rank $\leq m$. Note that, for each $n \in D$, $\langle H_n, h_n\rangle \models \varphi_1$ if and only if $\langle K_n, k_n\rangle \models \varphi_1$. Let D_1 be the set of all $n \in D$ for which $\langle H_n, h_n\rangle \models \varphi_1$. Let H'_n and K'_n be models obtained by updating H_n and K_n by $[\varphi_1]_A$. By the definition of rank and the fact that $|\varphi_1| \leq m$, we have that $\langle H'_n, h_n\rangle$ and $\langle K'_n, k_n\rangle$ agree on sentences from Φ of rank $\leq m$. Therefore, by our inductive hypothesis

$$\langle \bigoplus_{n\in D_1} H'_n, h\rangle \models \varphi_2 \quad \text{iff} \quad \langle \bigoplus_{n\in D_1} K'_n, k\rangle \models \varphi_2.$$

However,

$$\langle \bigoplus_n H_n, h\rangle \models \varphi \quad \text{iff} \quad \langle \bigoplus_{n\in D_1} H'_n, h\rangle \models \varphi_2$$

and

$$\langle \bigoplus_n K_n, k\rangle \models \varphi \quad \text{iff} \quad \langle \bigoplus_{n\in D_1} K'_n, k\rangle \models \varphi_2,$$

and we are done. $\dashv$

Returning to our theorem We get Theorem 8.5 directly from the claim, the observation that $\langle G_n, a_n\rangle \models \chi$ and $\langle G_n, a_\infty\rangle \models \neg\chi$, and Lemma 8.6. $\dashv$

8.4 $\mathcal{L}(\mathbf{\Sigma}_{\mathbf{pub}})$ lacks the finite model property

We conclude with a result on *iterating* epistemic actions. We gave a semantics of sentences of the form $[\alpha^*]\varphi$ earlier in the paper. The ability to iterate actions is useful in algorithms and even in informal presentations of scenarios. As it happens, the iteration operation on actions makes our systems quite expressive, so much so that the finite model property fails for $\mathcal{L}_1(\mathbf{\Sigma})$ as soon as $\mathbf{\Sigma}$ contains public announcements.

Proposition 8.7 $\langle \text{Pub } \Diamond\mathsf{True}\rangle^*\Diamond\Box\mathsf{False}$ *is satisfiable, but not in any finite model.*

Proof A state s in a model $\mathbf{S}$ is called an *end state* if s has no successors. For each model $\mathbf{S}$, let $\mathbf{S}'$ be the same model, except with the end states removed. $\mathbf{S}'$ is isomorphic to what we have written earlier as $\mathbf{S} \otimes (\mathit{Pub}\ \Diamond\mathsf{True})$, that is, the result of publically announcing that some world is possible. So we see that our sentence $\langle \mathit{Pub}\ \Diamond\mathsf{True}\rangle^*\Diamond\Box\mathsf{False}$ holds of s just in case the following holds for all n:

1. $s \in \mathbf{S}^{(n)}$, the n-fold application of the derivative operation to $\mathbf{S}$.

2. s has some child which is an end node (and hence t would not belong to $\mathbf{S}^{(n+1)}$).

It is clear that any model of $\langle Pub\ \Diamond\mathsf{True}\rangle^{*}\Diamond\Box\mathsf{False}$ must be infinite, since the sets $S^{(n+1)} \setminus S^{(n)}$ are pairwise disjoint and nonempty.

There is well-known models of $\langle Pub\ \Diamond\mathsf{True}\rangle^{*}\Diamond\Box\mathsf{False}$. One would be the set of decreasing sequences of natural numbers, with $s \to t$ iff t is a one-point extension of s. The end nodes are the sequences that end in 0. For each n, $\mathbf{S}^{(n)}$ is the submodel consisting of the sequences which end in n. $\dashv$

This has several dramatic consequences for this work. First, it means that the logics cannot be translated into the modal mu-calculus, since that logic is known to have the finite model property. More importantly, we have the following extension of this result:

Theorem 8.8 (Miller and Moss [14]) *Concerning $\mathcal{L}(\mathbf{\Sigma}_{pub})$:*

1. *$\{\varphi : \varphi$ is satisfiable$\}$ is Σ^1_1-complete.*
2. *$\{\varphi : \varphi$ is satisfiable on a finite model$\}$ is Σ^0_1-complete.*

Indeed, these even hold when $\mathcal{L}(\mathbf{\Sigma}_{\text{pub}})$ is replaced by small fragments, such as the fragment built from $[Pub\ \Diamond\mathsf{True}]$, $\Diamond^{*}$, $\Diamond$, and boolean connectives (yet without atomic sentences). These negative results go via reduction from domino problems.

The upshot is that logics which iterate epistemic actions are not going to be axiomatizable.

We feel that our results on expressive power are just a sample of what could be done in this area. We did not investigate the next natural questions: Do announcements with suspicious outsiders extend the expressive power of modal logic with all secure private announcements and common knowledge operators? And then do announcements with common knowledge of suspicion add further expressive power?

References

[1] Alexandru Baltag, A logic of epistemic actions, ms. CWI, Amsterdam, 1999.

[2] Alexandru Baltag, Lawrence S. Moss, and Sławomir Solecki, The logic of common knowledge, public announcements, and private suspicions, *Proceedings of TARK-VII (Theoretical Aspects of Rationality and Knowledge)*, 1998.

[3] Jon Barwise and Lawrence S. Moss, **Vicious Circles: On the Mathematics of Non-Wellfounded Phenomena**, CSLI Lecture Notes Number 60, CSLI Publications, Stanford University, 1996.

[4] Nachum Dershowitz, Orderings for term-rewriting systems, *Theoretical Computer Science* **17** (1982) 279–301.

[5] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi, **Reasoning About Knowledge**, MIT Press, 1996.

[6] Michael J. Fischer and Richard E. Ladner, Propositional modal logic of programs. *J. Comput. System Sci.* 18 (1979), no. 2, 194–211.

[7] Jelle Gerbrandy, Dynamic epistemic logic, in Lawrence S. Moss, et al (eds.) **Logic, Language, and Information, vol. 2**, CSLI Publications, Stanford University 1999.

[8] Jelle Gerbrandy, *Bisimulations on Planet Kripke*, Ph.D. dissertation, University of Amsterdam, 1999.

[9] Jelle Gerbrandy and Willem Groeneveld, Reasoning about information change, *Journal of Logic, Language, and Information* **6** (1997) 147–169.

[10] P. Gochet and P. Gribomont, **Epistemic Logic**, ms., 2002.

[11] Joseph Y. Halpern and Moshe Y. Vardi, The complexity of reasoning about knowledge and time. I. Lower bounds. *Journal of Computer and System Sciences* **38** (1989), no. 1, 196–237.

[12] D. Kozen and R. Parikh, An elementary proof of the completeness of PDL, *Theoretical Computer Science* (1981) 113–118.

[13] J.- J. Mayer and W. van der Hoek, **Epistemic Logic for AI and Computer Science**, Cambridge University Press, Cambridge, 1995.

[14] Joseph Miller and Lawrence S. Moss, The Undecidability of Iterated Modal Relativization, ms., Indiana University, 2003.

[15] Lawrence S. Moss, From hypersets to Kripke models in logics of announcements, in J. Gerbrandy et al (eds.) **JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 60th Birthday**, Vossiuspers, Amsterdam University Press, 1999.

[16] Lawrence S. Moss and Sławomir Solecki, Dynamic logics of epistemic actions, ms, Indiana University, 2001.

[17] David A. Plaisted, Termination orderings, in D. Gabbay (et al) eds., **Handbook of Logic in Artificial Intelligence and Logic Programming**, vol. I, 273–364.

[18] Jan Plaza, Logics of public communications, *Proceedings,* 4^{th} *International Symposium on Methodologies for Intelligent Systems*, 1989.

[19] Vaughn R. Pratt, Semantical considerations on Floyd-Hoare logic, in *17th Annual Symposium on Foundations of Computer Science*, IEEE Comput. Soc., Long Beach, Calif., 1976, 109–121.

[20] Lei Qian, Sentences true after being announced, in *Proceedings, Student Conference of the 2002 North American Summer School in Logic, Language, and Information*, Stanford University 2002.

[21] Johan van Benthem, Update delights, ms., 2000.

[22] Hans P. van Ditmarsch, *Knowledge Games*, Ph.D. dissertation, University of Groningen, 2000.

[23] Frank Veltman, Defaults in update semantics, *Journal of Philosophical Logic* **26** (1996) 221–261.