

The Logic of Public Announcements, Common Knowledge, and Private Suspicions

Alexandru Baltag¹
abaltag@cwi.nl

Lawrence S. Moss²
lsm@cs.indiana.edu

Sławomir Solecki²
ssolecki@indiana.edu

Comments Welcome

Abstract

This paper presents a logical system in which various group-level epistemic actions are incorporated into the object language. That is, we consider the standard modeling of knowledge among a set of agents by multi-modal Kripke structures. One might want to consider actions that take place, such as announcements to groups privately, announcements with suspicious outsiders, etc. In our system, such actions correspond to additional modalities in the object language. That is, we do not add machinery on top of models (as in, e.g., Fagin et al [4]), but we reify aspects of the machinery in the logical language.

Special cases of our logic have been considered in Plaza [13], Gerbrandy [5, 6], and Gerbrandy and Groeneveld [7]. The latter group of papers introduce a language in which one can faithfully represent all of the reasoning in examples such as the Muddy Children scenario. In that paper we find operators for updating worlds via announcements to groups of agents who are isolated from all others. We advance this by considering many more actions, and by using a more general semantics.

Our logic contains the infinitary operators used in the standard modeling of common knowledge. We present a sound and complete logical system for the logic, and we study its expressive power.

1 Introduction: Example Scenarios and Their Representations

We introduce the issues in this paper by presenting a few *epistemic scenarios*. These are all based on the Muddy Children scenario, well-known from the literature on knowledge. The intention is to expose the problems that we wish to address. These problems are first of all to get models which are faithful to our intuitions, and then to build and study logical systems which capture some of what is going on in the scenarios.

The cast of characters consists of three children: A , B , and C . So that we can use pronouns for them in the sequel, we assume that A is male, and B and C are female. Furthermore, A and B are dirty, and C is clean. Each of the children can see all and only the others. It is known to all (say, as a result of a shout from one of the parents) that at least one child is dirty. Furthermore, each child must try to figure out his or her state only by stating “I know whether I’m dirty or not” or “I don’t know whether I’m dirty or not.” They must tell the truth, and

¹CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands.

²Mathematics Department, Indiana University, Bloomington, IN 47401 USA.

they are perfect reasoners in the sense that they know all of the semantic consequences of their knowledge. The opening situation and these rules are all assumed to be common knowledge.

SCENARIO 1. After reflection, A and B announce to everyone that at that point they do not know whether they are dirty or not. (The reason we are having A and B make this announcement rather than all three children is that it fits in better with our scenarios to follow.) Let α denote this announcement.

As in the classical Muddy Children, there are intuitions about knowledge before and after α . Here are some of those intuitions. Before α , nobody should know that he or she is dirty. However, A should think that it is possible that B knows. (For if A were clean, B would infer that she must be the dirty one.) After α , A and B should each know that they are dirty, and hence they know whether they are dirty or not. On the other hand, C should not know whether she is dirty or not.

SCENARIO 1.5. This scenario begins after α . At this point, A and B announce to all three that they *do* know whether or not they are dirty. We'll call this event α' . Our intuition is that after α' , C should know that she is not dirty. Moreover, A and B should know that C knows this. Actually, the dirty-or-not states of all the children should be *common knowledge* to all three.

SCENARIO 2. As an alternative to the first scenario, let's assume that C falls asleep for a minute. During this time, A and B got together and told each other that they didn't know whether they were dirty or not. Let β denote this event. After β , C wakes up. Part of what we mean by β is that C does not even consider it possible that β occurred, and that it's common knowledge to A and B that this is the case. Then our intuitions are that after β , C should "know" (actually: believe) that A does not know whether he is dirty (and similarly for B); and this fact about C is common knowledge for all three children. Of course, it should also be common knowledge to A and B that they are dirty.

SCENARIO 2.5. Following Scenario 2, we again have α' : A and B announce that they do know whether they are dirty or not. Our intuitions are not entirely clear at this point. Surely C should suspect some kind of cheating or miscalculation on the part of the others. However, we will not have much to say about the workings of this kind of real-world sensibility. Our goal will be more in the direction of modeling different alternatives.

SCENARIO 3. Now we vary Scenario 2. C merely feigned sleep and thought she heard both A and B whispering. C cannot be sure of this, however, and also entertains the possibility that nothing was communicated. (In reality, A and B did communicate.) A and B for their part, still believe that C was sleeping. We call this event γ .

One might at first glance think that A and B 's "knowledge" of C 's epistemic state is unchanged by γ . After all, the communication was not about C . However, we work with a semantic notion of knowledge, and after γ , A and B know that they are dirty, hence then know that C knows that they are dirty. A and B did not know this at the outset.

So we need to revise the initial intuition. What is correct is that if C knows some fact φ before γ , then after γ , A and B know (or rather, believe) that C knows φ . This is because after γ , A and B not only know the clean-or-dirty state of everyone, they (therefore) also know exactly which possibilities everyone is aware of, which they discard as impossible, etc. So each of them can reconstruct C 's entire epistemic state. They believe that their reconstruction is current, but of course, what they reconstruct is C 's original one, before γ .

Conversely, if after γ , A and B “know” that C knows φ , then before γ , C really did know φ . That is, the reconstruction is accurate. For example, after γ , A believes that C should not consider it possible that A knows that he is dirty. However, C thinks it is possible that A knows he is dirty.

There is a stronger statement that is true: C knows φ before γ iff after γ , it is common knowledge to A and B that each of them knows that C knows φ . Intuitively, this holds because each of A and B knows that both of them are able to carry out the reconstruction of C 's state.

Our final intuition is that after γ , C should know that if A were to subsequently announce that he knows that he is dirty, then C would know that B knows that she is dirty.

SCENARIO 3.5. Again, continue Scenario 3 by α' . At this point, C should know that her suspicions were confirmed, and hence that she is not dirty. For their part, A and B should think that C is confused by α' : they should think that C is as she was following Scenario 2.5.

SCENARIO 4. A and B are on one side of the table and C is on the other, dozing. C wakes up at what looks to her like the middle of a joint confession by A and B . The two sides stare each other down. In fact, A and B have already communicated. We call this action δ . So C suspects that δ is what happened, but can't tell if it was δ or nothing. For their part, A and B see that C suspects but does not know that δ happened.

The basic intuition is that after δ , it should be common knowledge to all three that C suspects that the communication happened. Even if C thinks that A and B did not communicate, C should not think that she is sure of this.

One related intuition is that after δ , it should be common knowledge that C suspects that A knows that he is dirty. As it happens, this intuition is wrong. Here is a detailed analysis: C thinks it possible that everyone is dirty at the outset, and if this were the case then the announcement of B 's ignorance would not help A to learn that he is dirty; from A 's point of view, he still could be clean and B would not know that she is dirty. C 's view on this does not change as a result of δ , so afterwards, C still thinks that it could be the case that A says, “It's possible that B and C are the dirty one and I am clean, Hence C would see my clean face and not suspect that I know that I am dirty.” So it certainly should not be common knowledge that C suspects that A knows he is dirty.

Notice also that C would say after δ : “I think it is possible that no announcement occurred, and yet A thinks it possible that B is the only dirty one. In that case, what A would think that I suspect that A told B that he knows that he is *not* dirty. Of course, this is not what I actually suspect.” The point is that C 's reasoning about A and B 's reasoning about her involves suspicion of a different announcement than we at first considered.

SCENARIO 4.5. Once again, we continue with α' . Our intuition is that this is tantamount to an admission of private communication by A and B . If we disregard this and only look at higher order knowledge concerning who is and is not dirty, we expect that the epistemic state after α' is the same for all three children as it is at the end of Scenario 1.5.

1.1 Models

Now that we have detailed a few scenarios and our intuitions about them, it is time to construct some Kripke models as representations for them.

THE MODELS U AND V . We begin with a representation of the situation before α . We take the Kripke model U whose worlds are u_1, \dots, u_7 and whose structure is given in the table on the left below:

World	A	B	C	$\overset{A}{\rightarrow}$	$\overset{B}{\rightarrow}$	$\overset{C}{\rightarrow}$
u_1			•	u_1, u_5	u_1, u_3	u_1
u_2		•		u_2, u_6	u_2	u_2, u_3
u_3		•	•	u_3, u_7	u_1, u_3	u_2, u_3
u_4	•			u_4	u_4, u_6	u_4, u_5
u_5	•		•	u_1, u_5	u_5, u_7	u_4, u_5
$u_6 \checkmark$	•	•		u_2, u_6	u_4, u_6	u_6, u_7
u_7	•	•	•	u_3, u_7	u_5, u_7	u_6, u_7

World	A	B	C	$\overset{A}{\rightarrow}$	$\overset{B}{\rightarrow}$	$\overset{C}{\rightarrow}$
v_1			•	v_1, v_5	v_1, v_3	v_1
v_3		•	•	v_3, v_7	v_1, v_3	v_3
v_5	•		•	v_1, v_5	v_5, v_7	v_5
$v_6 \checkmark$	•	•		v_6	v_6	v_6, v_7
v_7	•	•	•	v_3, v_7	v_5, v_7	v_6, v_7

For example, in world u_3 , A is clean, but B and C are dirty. Also, the worlds which A thinks are possible are u_3 and u_7 . That is, A sees that B and C are dirty, so A infers that the world is either u_3 or u_7 . The rest of the structure is explained similarly, except for the \checkmark mark next to u_6 . This specifies u_6 as the actual world in the model, the one which corresponds to our description of the model before α . Note that U incorporates some of the conventions stated in Scenario 1. For example, in each world, each child has a complete and correct assessment of which worlds are possible for all three reasoners.

Each of our intuitions about knowledge before α turns into a statement in the modal logic of knowledge. This logic has atomic sentences D_A , D_B , and D_C standing for “ A is dirty”, etc.; it has knowledge operators \Box_A , \Box_B , and \Box_C along with the usual boolean connectives. We are going to use the standard Kripke semantics for multi-modal logic throughout this paper. So given a model-world pair, say $\langle A, a \rangle$, and some agent, say D , we’ll write

$$\langle K, k \rangle \models \Box_D \varphi \quad \text{iff} \quad \text{whenever } k \overset{D}{\rightarrow} l \text{ in } K, \text{ we have } \langle K, l \rangle \models \varphi.$$

The boolean connectives will be interpreted classically. We can then check the following:

$$\begin{aligned} \langle U, u_6 \rangle &\models \neg \Box_A D_A \wedge \neg \Box_A \neg D_A \wedge \neg \Box_B D_B \wedge \neg \Box_B \neg D_B \wedge \neg \Box_C D_C \wedge \neg \Box_C \neg D_C \\ \langle U, u_6 \rangle &\models \Diamond_A \Box_B D_B \end{aligned}$$

The model after α is the Kripke model V , shown on the right above. The way we got V from U was to discard the worlds u_2 and u_4 of U , since in U at each of those worlds, either A or B would know if they were dirty. We also changed the u ’s to v ’s to avoid confusion, and to stress

the fact that we get a new model. Turning back to our intuitions, we can see that the following holds:

$$\langle V, v_6 \rangle \models \Box_A D_A \wedge \Box_B D_B \wedge \neg(\Box_C D_C \vee \Box \neg D_C)$$

THE MODEL W . Scenario 1.5 elaborates Scenario 1 by the event α' . So we discard the worlds where this is false in V , and we obtain a one-world model W :

World	A	B	C	$\overset{A}{\rightarrow}$	$\overset{B}{\rightarrow}$	$\overset{C}{\rightarrow}$
$w_6 \checkmark$	•	•		w_6	w_6	w_6

(We have renamed v_6 to w_6 .) This model reflects our intuition that at this point, C should know that she is not dirty.

THE MODEL X . This corresponds to Scenario 2. We start with U and see the effect of the private announcement β . The resulting model X is:

World	A	B	C	$\overset{A}{\rightarrow}$	$\overset{B}{\rightarrow}$	$\overset{C}{\rightarrow}$
u_1, \dots, u_7						
x_1			•	x_1, x_5	x_1, x_3	u_1
x_3		•	•	x_3, x_7	x_1, x_3	u_2, u_3
x_5	•		•	x_1, x_5	x_5, x_7	u_4, u_5
$x_6 \checkmark$	•	•		x_6	x_6	u_6, u_7
x_7	•	•	•	x_3, x_7	x_5, x_7	u_6, u_7

Notice that the worlds u_1, \dots, u_7 are also worlds in X . We did not put any information in the chart above for those worlds since it should be exactly the same as in U above. The reason for having these “old worlds” in X is that since C was asleep, the worlds that C considers possible after β should be just the ones that were possible before β . We can check that

$$\langle X, x_6 \rangle \models \neg \Box_C (\Box_A D_A \vee \Box_A \neg D_A).$$

Let φ be the sentence above. Then also, $\langle X, x_6 \rangle \models \Box_{\{A, B, C\}}^* \varphi$. This is our formal statement that it is common knowledge in the group of three children that φ holds. The semantics of this is that for all sequences $D_1, \dots, D_m \in \{A, B, C\}^*$, $\langle X, x_6 \rangle \models \Box_{D_1} \dots \Box_{D_m} \varphi$. Note that we have no way of saying in the modal language that C suspects that an announcement happened; the best we can do is (roughly) to say that C thinks that some sentence ψ is possible in the sense that ψ holds in some possible world. Of course, we have no way to say that A and B know that C was asleep, either.

Note as well that in X , we do not have $x_6 \overset{C}{\rightarrow} x_6$. In other words, the real world would not be possible for C . This is some indication that something strange is going on in this model. Further, we consider the model of what happens after A and B 's announcement. Then in this model, *no worlds* would be accessible for C from the actual world. These anomalies should justify our interest in the more complicated scenarios and models involving suspicions of announcements.

THE MODEL OBTAINED BY ANNOUNCING α' IN X . This would be the one-world model below:

World	A	B	C	\xrightarrow{A}	\xrightarrow{B}	\xrightarrow{C}
$x_6^* \checkmark$	•	•		x_6^*	x_6^*	

We have not only deleted the worlds where either A or B does not know that they are dirty in X , but we also discarded all worlds not reachable from the new version x_6^* of x_6 . The anomaly here is that C thinks no worlds are possible.

THE MODEL Y . We consider γ from Scenario 3, in which C thought she might have heard A and B , while A and B think that C is unaware of γ . We get the model Y displayed in Figure 1 below. Y has 24 worlds, and so we won't justify all of them individually. We will give a more principled construction of Y from W and γ , once we have settle on a mathematical model of γ . For now, the ideas are that the y worlds are those where the announcement happened, and the y' worlds are those in which it did not. Note that some of the y worlds are missing, since the truthful announcement by A and B presupposes that they don't know whether they are dirty in U at the corresponding world. The x 's and u 's are from above, and they inherit the

World	A	B	C	\xrightarrow{A}	\xrightarrow{B}	\xrightarrow{C}
u_1, \dots, u_7						
x_1, x_3, x_5, x_6, x_7						
y_1			•	x_1, x_5	x_1, x_3	y_1, y'_1
y_3		•	•	x_3, x_7	x_1, x_3	y_3, y'_2, y'_3
y_5	•		•	x_1, x_5	x_5, x_7	y_5, y'_4, y'_5
$y_6 \checkmark$	•	•		x_6	x_6	y_6, y_7, y'_6, y'_7
y_7	•	•	•	x_3, x_7	x_5, x_7	y_6, y_7, y'_6, y'_7
y'_1			•	u_1, u_5	u_1, u_3	y_1, y'_1
y'_2		•		u_2, u_6	u_2	y_3, y'_2, y'_3
y'_3		•	•	u_3, u_7	y_1, y_3	y_3, y'_2, y'_3
y'_4	•			u_4	u_4, u_6	y_5, y_4, y'_5
y'_5	•		•	u_1, u_5	u_5, u_7	y_5, y'_4, y'_5
y'_6	•	•		u_2, u_6	u_4, u_6	y_6, y_7, y'_6, y'_7
y'_7	•	•	•	u_3, u_7	u_5, u_7	y_6, y_7, y'_6, y'_7

Figure 1: The model Y

accessibility relations which we have seen.

Now our main intuition here is that $\langle U, u_6 \rangle \models \Box_C \varphi$ iff $\langle Y, y_6 \rangle \models \Box_{\{A, B\}}^+ \Box_C \varphi$. (The sentence $\Box_{\{A, B\}}^+ \chi$ means that A knows φ , A knows B knows χ , etc. It differs from $\Box_{\{A, B\}}^* \chi$ in that it does not entail that χ is true.) To see this, note that $u_6 \xrightarrow{C} u_6, u_7$ and no other worlds. And the only worlds reachable from y_6 using one or more \xrightarrow{A} or \xrightarrow{B} transitions followed by a \xrightarrow{C} transition are again u_6 and u_7 .

Another intuition is that in $\langle Y, y_6 \rangle$, C should think that it is possible that A knows that he is dirty. This is justified since $y_6 \xrightarrow{C} x_6$, and $\langle Y, x_6 \rangle \cong \langle X, x_6 \rangle$ (that is, the submodels of X and Y generated by x_6 are isomorphic), and $\langle X, x_6 \rangle \models \Box_A \text{DA}$.

Our final intuition is that in $\langle Y, y_6 \rangle$, C should know that if A were to subsequently announce that he knows that he is dirty, then C would know that B knows that she is dirty. To check this, we need to modify Y by deleting the worlds where A does not know that he is dirty. These include y_7 , y'_6 and y'_7 . In the updated model, the only world accessible for C from (the new version of) y_6 is y_6 itself, and at y_6 in the new structure, B correctly knows she is dirty.

THE MODEL OBTAINED BY ANNOUNCING α' IN Y . As when α' is announced in X , we only keep the worlds of Y worlds where both A or B do know they are dirty. So we drop y_7 , y'_6 , and y'_7 .

World	A	B	C	$\overset{A}{\rightarrow}$	$\overset{B}{\rightarrow}$	$\overset{C}{\rightarrow}$
$y_6^* \checkmark$	•	•		$x_6^\#$	$x_6^\#$	y_6^*
$x_6^\#$	•	•		$x_6^\#$	$x_6^\#$	

We also only keep the worlds accessible from y_6 (this change is harmless). C knows she is not dirty. Technically, A and B “know” this, but this is for the nonsensical reason that they “know” that C knows everything.

THE MODEL Z . Corresponding to Scenario 4, we get the model Z shown below.

w	A	B	C	$\overset{A}{\rightarrow}$	$\overset{B}{\rightarrow}$	$\overset{C}{\rightarrow}$	w	A	B	C	$\overset{A}{\rightarrow}$	$\overset{B}{\rightarrow}$	$\overset{C}{\rightarrow}$
z_1			•	z_1, z_5	z_1, z_3	z_1, z'_1	z'_1			•	z'_1, z'_5	z'_1, z'_3	z_1, z'_1
z_2		•		z_2	z_2	z_2, z_3, z'_2, z'_3	z'_2		•		z'_2, z'_6	z_2, z'_2	z_2, z_3, z'_2, z'_3
z_3		•	•	z_3, z_7	z_1, z_3	z_2, z_3, z'_2, z'_3	z'_3		•	•	z'_3, z'_7	z'_1, z'_3	z_2, z_3, z'_2, z'_3
z_4	•			z_4	z_4	z_4, z_5, z'_4, z'_5	z'_4	•			z'_4, z'_6	z'_4, z'_6	z_4, z_5, z'_4, z'_5
z_5	•		•	z_1, z_5	z_5, z_7	z_4, z_5, z'_4, z'_5	z'_5	•		•	z'_1, z'_5	z'_5, z'_7	z_4, z_5, z'_4, z'_5
$z_6 \checkmark$	•	•		z_6	z_6	z_6, z_7, z'_6, z'_7	z'_6	•	•		z'_2, z'_6	z'_4, z'_6	z_6, z_7, z'_6, z'_7
z_7	•	•	•	z_3, z_7	z_5, z_7	z_6, z_7, z'_6, z'_7	z'_7	•	•	•	z'_3, z'_7	z'_5, z'_7	z_6, z_7, z'_6, z'_7

Recall our last point in Scenario 4, that we need to consider a few possible announcements for C to suspect. This is reflected in the fact that the z worlds are of three types. In z_2 , B announced that she knows whether she is dirty, and A announced that he doesn't. Similar remarks apply to z_4 . In all other z worlds, both announced that they do not know. The worlds accessible from each of these is based on the relevant announcement. For example, in z_2 , neither A nor B thinks any other world is possible. (One might think that $z_2 \overset{A}{\rightarrow} z_6$. But in z_6 , B could not announce that she knows she is dirty. So if the world were z_2 and the relevant announcement made, then A would not think z_6 is possible.) The z' worlds are those in which no announcement actually happened.

Our key intuition was that it is common knowledge that C suspects that δ happened. This will not correspond to anything in the formal language $\mathcal{L}([\alpha], \square^*)$ introduced later in this paper. (However, it will be representable in an auxiliary language about actions; see Example 2.3.) Informally, the intuition is valid for Z because for every z_i (or z'_i) there is some z_j (unprimed) such that $z_i \overset{C}{\rightarrow} z_j$ (or $z'_i \overset{C}{\rightarrow} z_j$). In addition, in this particular model there is a sentence in our

formal language which happens to hold only at the worlds where an announcement occurred. Here is one:

$$\chi \equiv \Box_A D_A \vee \Box_A \neg D_A \vee \Diamond_{\{A,B\}}^* \Diamond_C \Box_A D_A$$

So $\langle Z, z_6 \rangle \models \Box_{\{A,B,C\}}^* \Diamond_C \chi$.

The explanation of the mistaken intuition in Scenario 4 is that $z_6 \xrightarrow{C} z_7 \xrightarrow{A} z_3$, and $\Diamond_A \Box_A D_A$ fails in z_2, z_3, z'_2 , and z'_3 . Overall, $\langle Z, z_6 \rangle \models \neg \Box_C \Box_A \Diamond_A \Box_A D_A$.

The point that C 's suspicion varies corresponds to the fact that $\Diamond_C \Diamond_A \Diamond_C \Box_A \neg D_A$ holds at $\langle Z, z_6 \rangle$. Indeed $z_6 \xrightarrow{C} z'_6 \xrightarrow{A} z'_2 \xrightarrow{C} z_2$, and $\langle Z, z_2 \rangle \models \Box_A \neg D_A$.

A few more involved statements are true in Z . For example, $\Box_{\{A,B,C\}}^* (\Box_A D_A \rightarrow \Diamond_C \Box_A D_A)$. It is common knowledge to all three that *if* A knows he is dirty, then C thinks it possible that A knows this.

THE MODEL OBTAINED BY ANNOUNCING α' IN Z . This model is W from above. (Actually, it is bisimilar to W ; see Section 2.2.) This corresponds to the intuition that Scenarios 2.5 and 4.5 lead to the same model.

1.2 Epistemic Actions

We will formalize a language in Section 2 along with the notions of (epistemic) *action structure* and *actions*. Before we do that, it makes sense to present the idea informally based on the examples which we have already dealt with.

α AND α' : ANNOUNCEMENTS TO EVERYONE. We first consider α of Scenario 1. Let ψ be given by

$$\psi := \neg(\Box_A D_A \vee \Box_A \neg D_A) \wedge \neg(\Box_B D_B \vee \Box_B \neg D_B) \quad (1)$$

So ψ says that neither A nor B know whether or not they are dirty. This is the *precondition* of the announcement, but it is not the *structure*. The structure of this announcement is quite simple (so much so that the reader will need to read further to get an idea for what we mean by structure). It is the following Kripke structure K : we take one point, call it k , and we set $k \xrightarrow{D} k$ for all $D \in \{A, B, C\}$. We call $\langle K, k \rangle$ an *action structure*. Along with K , we also have a *precondition*; this will be ψ from (1). To deal with action structures with more than one point, the precondition will be a function PRE from worlds to sentences. In this case, the function PRE is just $\{\langle k, \psi \rangle\}$. The tuple $\langle K, k, \text{PRE} \rangle$ will be an example of what we call an *action*. This particular action is our model of the announcement α . Henceforth we use the symbol α to refer ambiguously to the pretheoretic notion of the announcement event and to our mathematical model of it.

Another example of an announcement to everyone is α' . Here we just change ψ from (1) to the sentence ψ' which says that both A and B know whether or not they are dirty. Yet another example is the *null announcement*. This models the announcement of a tautology *true* to everyone. We'll write this as τ .

β : A SECURE ANNOUNCEMENT TO A SET OF AGENTS. Next, suppose we have an announcement made to some possibly proper subset $\mathcal{B} \subseteq \mathcal{A}$ in the manner of Scenario 2. So there is some dispute as to what happened: the agents in \mathcal{B} think that there was an announcement, while

those out of \mathcal{B} are sure that nothing happened. We formalize this with a Kripke structure of two points, l and t . We set $l \xrightarrow{D} l$ for all $D \in \mathcal{B}$, $l \xrightarrow{D} t$ for $D \notin \mathcal{B}$, and $t \xrightarrow{D} t$ for all D . The point is that l here is the actual announcement, and the agents in \mathcal{B} know that this is the announcement. The agents not in \mathcal{B} think that t is for sure the only possible action, and t in this model will behave just like the null announcement. The precondition function will be called PRE in all of our examples. Here PRE is given by $\text{PRE}(l) = \psi$ and $\text{PRE}(t) = \text{true}$, where ψ is from (1). The action overall is $\langle L, l, \text{PRE} \rangle$, where $L = \{l, t\}$. We call this action β .

γ : AN ANNOUNCEMENT WITH A SUSPICIOUS OUTSIDER. This is based on Scenario 3. The associated structure has four points, as follows:

World	$A \rightarrow$	$B \rightarrow$	$C \rightarrow$	PRE
$m\checkmark$	l	l	m, n	ψ
n	t	t	m, n	true
l	l	l	t	ψ
t	t	t	t	true

The idea is that m is the (private) announcement that C suspects, and n is other announcement that C thinks is possible (where nothing was communicated by A and B). Then if m happened, A and B were sure that what happened was l ; similarly, if n happened, A and B would think that t was what happened. We call this action γ ; technically it is $\langle \{m, n, l, t\}, m, \text{PRE} \rangle$. We get a different action, say γ' if we use the same model as above but change the designated (“real”) world from m to n .

δ : AN ANNOUNCEMENT WITH COMMON KNOWLEDGE OF SUSPICION. Corresponding to Scenario 4, we have the following model. In it ψ_A denotes the sentence saying that A knows whether he is dirty but B does not, ψ_B the sentence saying that B knows whether she is dirty but A does not, and ψ_\emptyset the sentence stating that neither knows.

World	$A \rightarrow$	$B \rightarrow$	$C \rightarrow$	PRE
$o\checkmark$	o	o	o, s	ψ
p	p	p	p, s	ψ_A
q	q	q	q, s	ψ_B
r	r	r	r, s	ψ_\emptyset
s	s	s	o, p, q, r, s	true

We call this action δ . There are five possible actions here, depending on whether it was ψ , ψ_A , ψ_B , ψ_\emptyset or nothing which was announced. In each case, A and B are sure of what happened. Even if nothing actually happened (s), C would suspect one of the other four possibilities. In those, C still considers it possible that nothing happened.

Still to come. The reader is perhaps wondering what the actual connection is between the (formal) actions just introduced and the concrete models of the previous section. The connection is that there is a way of taking a model and an action and producing another model. When applied to the specific model U and the actions of this section, we get the models V, \dots, Z .

We delay this connection until Section 2.2 below, since it is high time that we introduce our language of epistemic actions and its semantics. The point is that there is a principled reason behind the models.

The question also arises as to whether there are any principles behind the particular actions which we put down in this section. As it happens, there is more which can be said on this matter. We postpone that discussion until Section 2.3, after we have formally defined the syntax and semantics of our logical languages.

1.3 The Issues

The main issue we address in this paper is to formally represent *epistemic updates*, i.e., changes in the information states of agents in a distributed system. We think of these changes as being induced by specific information-updating actions, which can be of various types: (1) information-gathering and processing (e.g., realizing the possibility of other agents' hidden actions, and more generally, learning of any kind of new possibility via experiment, computation, or introspection); (2) information-exchange and communication (learning by sending/receiving messages, public announcements, secret interception of messages, etc.); (3) information-hiding (lying or other forms of deceiving actions, such as communication over secret channels, sending encrypted messages, holding secret suspicions); (4) information-loss and misinformation (being lied to, starting to have gratuitous suspicions, non-introspective learning, wrong computations or faulty observations, paranoia); (5) and more generally sequential or synchronous combinations of all of the above.

Special cases of our logic, dealing only with public or semi-public *announcements to mutually isolated groups*, have been considered in Plaza [13], Gerbrandy [5, 6], and Gerbrandy and Groeneveld [7]. These deal with actions such as α and β in our Introduction. Our examples γ and δ go beyond what is possible in the setting of these papers. 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 formal system, we capture only the *epistemic aspect* of these real actions, disregarding other (intentional) aspects. In particular, for simplicity reasons, 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 the end of Section 2.3 and also Section 5.3).

On the semantical side, the main original technical contribution of our paper lies in our decision to represent not only the epistemic states, but also the epistemic actions, by Kripke structures. 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-structures 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 for the consistency of the action with the state.

On the syntactical side, we use a mixture of dynamic and epistemic logic, with dynamic modalities associated to each action-structure, 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.

1.4 Further Contents of This Paper

Section 2 gives our basic logic $\mathcal{L}([\alpha])$ of epistemic actions and knowledge. The idea is to define the logic together with the action structures which we have just looked at informally. So in $\mathcal{L}([\alpha])$ we finally will present the promised formal versions of the announcements of Section 1.2. In Section 3 we present a sound and complete axiomatization of $\mathcal{L}([\alpha])$. We add the common knowledge operators to get $\mathcal{L}([\alpha], \Box^*)$ in Section 4. Completeness for this logic is proved in Section 5. Two results on the expressive power are presented in Section 6. An Appendix contains some technical results which, while needed for our work, seem to interrupt the flow of the paper.

2 A Logical Language with Epistemic Actions

2.1 Syntax

We begin with a set AtSen of atomic sentences, and we define two sets simultaneously: the language $\mathcal{L}([\alpha])$, and a set of *actions* (over $\mathcal{L}([\alpha])$).

$\mathcal{L}([\alpha])$ is the smallest collection which includes AtSen and which is closed under \neg , \wedge , \Box_A for $A \in \mathcal{A}$, and $[\alpha]\varphi$, where α is an action over $\mathcal{L}([\alpha])$, and $\varphi \in \mathcal{L}([\alpha])$.

An *action structure* (over $\mathcal{L}([\alpha])$) is a pair $\langle K, \text{PRE} \rangle$, where K is a finite Kripke frame over the set \mathcal{A} of agents, and PRE is a map $\text{PRE} : K \rightarrow \mathcal{L}$. We will usually write K for the action structure $\langle K, \text{PRE} \rangle$. An *action* (over $\mathcal{L}([\alpha])$) is a tuple $\alpha = \langle K, k, \text{PRE} \rangle$, where $\langle K, \text{PRE} \rangle$ is an action structure over $\mathcal{L}([\alpha])$, and $k \in K$. Each action α thus is a finite set with relations \mathcal{D}_D for $D \in \mathcal{A}$, together with a precondition function and a specified *actual* world.

The actions themselves constitute a Kripke frame Actions in the natural way, by setting

$$\langle K, k, \text{PRE} \rangle \mathcal{D}_D \langle L, l, \text{PRE}' \rangle \quad \text{iff} \quad K = L, \text{PRE} = \text{PRE}', \text{ and } k \mathcal{D}_D l \text{ in } K. \quad (2)$$

When $\alpha = \langle K, k, \text{PRE} \rangle$, we set $\text{PRE}(\alpha) = \text{PRE}(k)$. That is, $\text{PRE}(\alpha)$ is the precondition associated to the distinguished world of the action. For this reason, we often write $\text{PRE}(\alpha)$ instead of $\text{PRE}(k)$.

Examples 2.1 All of the sentences mentioned in Section 1.1 are sentences of $\mathcal{L}([\alpha])$, except for the ones that use $\Box_{\{A,B,C\}}^*$. This construct gives us a more expressive language, as we shall see. The structures $\alpha, \tau, \beta, \gamma, \gamma', \delta$, and δ' described informally in Section 1.2 are bona fide actions. As examples of the accessibility relation on the class of actions, we have the following facts: $\alpha \xrightarrow{D} \alpha$ and $\tau \xrightarrow{D} \tau$ for all $D \in \{A, B, C\}$; $\beta \xrightarrow{B} \beta$; $\beta \xrightarrow{C} \beta$; $\beta \xrightarrow{C} \tau$; $\gamma \xrightarrow{A\beta} \beta$; $\gamma, \gamma' \xrightarrow{C} \gamma, \gamma'$; $\gamma' \xrightarrow{A\beta} \tau$; $\delta \xrightarrow{A\beta} \delta$, $\delta' \xrightarrow{A\beta} \delta'$, and $\delta, \delta' \xrightarrow{C} \delta, \delta'$.

Many other types of examples are possible. We can represent misleading epistemic actions, e.g. lying, or more generally acting such that some people do not suspect that your action is possible. 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.

2.2 Semantics

As with the syntax, we define two things simultaneously: the semantic relation $\langle W, w \rangle \models \varphi$, and a partial operation $(\langle W, w \rangle, \alpha) \mapsto \langle W, w \rangle^\alpha$. Before this, we need another definition. Given a model W and an action structure K , we define the model W^K as follows:

1. The worlds of W^K are the pairs $(w, k) \in W \times K$ such that $\langle W, w \rangle \models \text{PRE}(k)$.
2. For such pairs,

$$(w, k) \xrightarrow{\Delta} (w', k') \quad \text{iff} \quad w \xrightarrow{\Delta} w' \text{ and } k \xrightarrow{\Delta} k'. \quad (3)$$

3. We interpret the atomic sentences by setting $v_{W^K}((w, k)) = v_W(w)$. That is, p is true at (w, k) in W^K iff p is true at w in W .

Given an action $\alpha = \langle K, k \rangle$ and a model-world pair $\langle W, w \rangle$, we say that $\langle W, w \rangle^\alpha$ is defined iff $\langle W, w \rangle \models \text{PRE}(k)$, and in that case we set $\langle W, w \rangle^\alpha = \langle W, w \rangle^{\langle K, k \rangle} = \langle W^K, (w, k) \rangle$. One can now check that the following holds for these definitions.

$$\langle W, w \rangle^\alpha \xrightarrow{\Delta} \langle W, x \rangle^\beta \quad \text{iff} \quad \langle W, w \rangle^\alpha \text{ and } \langle W, x \rangle^\beta \text{ are defined, } w \xrightarrow{\Delta} x \text{ in } W, \text{ and } \alpha \xrightarrow{\Delta} \beta.$$

The semantics is given by extending the usual clauses for modal logic by one for actions:

$$\langle W, w \rangle \models [\alpha]\varphi \quad \text{iff} \quad \langle W, w \rangle^\alpha \text{ is defined implies } \langle W, w \rangle^\alpha \models \varphi.$$

As is customary, we abbreviate $\neg[\alpha]\neg\varphi$ by $\langle \alpha \rangle \varphi$. Then we have

$$\langle W, w \rangle \models \langle \alpha \rangle \varphi \quad \text{iff} \quad \langle W, w \rangle^\alpha \text{ is defined and } \langle W, w \rangle^\alpha \models \varphi.$$

We also abbreviate the boolean connectives classically, and we let *true* denote some tautology such as $p \vee \neg p$.

The larger language $\mathcal{L}([\alpha], \Box^*)$ We also consider a larger language $\mathcal{L}([\alpha], \Box^*)$. This is defined by adding operators $\Box_{\mathcal{B}}^*$ for all subsets $\mathcal{B} \subseteq \mathcal{A}$. (When we do this, of course we get more actions as well.) The semantics works by taking $\Box_{\mathcal{B}}^* \varphi$ to abbreviate the infinitary conjunction

$$\bigwedge_{\langle A_1, \dots, A_n \rangle \in \mathcal{B}^*} \Box_{A_1} \cdots \Box_{A_n} \varphi.$$

Here \mathcal{B}^* is the set of all sequences from \mathcal{B} . This includes the empty sequence, so $\Box_{\mathcal{B}}^* \varphi$ logically implies φ .

Bisimulation Given two models, say K and L , over the same set of \mathcal{A} of agents, a *bisimulation* between K and L is a relation $R \subseteq K \times L$ such that if kRl and $A \in \mathcal{A}$, then:

1. For all atomic p , $\langle K, k \rangle \models p$ iff $\langle L, l \rangle \models p$.
2. For all $k \xrightarrow{A} k'$ there is some $l \xrightarrow{A} l'$ such that $k'Rl'$.
3. For all $l \xrightarrow{A} l'$ there is some $k \xrightarrow{A} k'$ such that $k'Rl'$.

Given two model-world pairs $\langle K, k \rangle$ and $\langle L, l \rangle$, we write $\langle K, k \rangle \equiv \langle L, l \rangle$ iff there is some bisimulation R such that kRl . It is a standard fact that if $\langle K, k \rangle \equiv \langle L, l \rangle$, then the two pairs agree on all sentences of standard modal logic. In our setting, we also can speak about actions being bisimilar: we change condition (1) above to refer to say that $\text{PRE}(k) = \text{PRE}(l)$. It is easy now to check two things simultaneously: (1) bisimilar pairs agree on all sentences of $\mathcal{L}([\alpha])$; and (2) if $\langle K, k \rangle \equiv \langle L, l \rangle$ and $\alpha \equiv \beta$, then $\langle K, k \rangle^\alpha \equiv \langle L, l \rangle^\beta$. Furthermore, these results extend to $\mathcal{L}([\alpha], \square^*)$.

Examples 2.2 We look back at Section 1.1 for some examples. We use \cong to denote the relation of isomorphism on model-world pairs. It is not hard to check the following: $\langle U, u_6 \rangle^\alpha \cong \langle V, v_6 \rangle$, $\langle U, u_6 \rangle^\beta \cong \langle X, x_6 \rangle$, $\langle U, u_6 \rangle^\gamma \cong \langle Y, y_6 \rangle$, and $\langle U, u_6 \rangle^\delta \cong \langle Z, z_6 \rangle$. For example, the isomorphism which shows that $\langle U, u_6 \rangle^\delta \cong \langle Z, z_6 \rangle$ is $(u_i, o) \mapsto z_i$ for $i \neq 2, 4$, $(u_2, q) \mapsto z_2$, $(u_4, p) \mapsto z_4$, and $(u_i, r) \mapsto z'_i$ for all i .

Let α' be the action of announcing to all agents that both A and B do know whether they are dirty. Then $\langle V, v_6 \rangle^{\alpha'} \cong \langle X, x_6 \rangle$. Moreover, $\langle Z, z_6 \rangle^{\alpha'} \equiv \langle X, x_6 \rangle$. Note that in this case we only have bisimilarity. However, we know that our languages will not discriminate between bisimilar pairs, so we can regard them as the same. This models our intuition that the epistemic states at the end of Scenarios 1.5 and 4.5 should be the same.

Finally, all of the semantic facts about the various models in Section 1.1 now turn into precise statements. For example, $\langle U, u_6 \rangle \models [\alpha] \diamond_A \square_B \text{D}_B$. Also, $\langle U, u_6 \rangle \models [\alpha][\alpha'] \square_{A,B,C}^* \square_C \text{D}_C$. This formalizes our intuition that if we start with $\langle U, u_6 \rangle$, first announce that each of A and B do not know their state, then second announce that they each do know it, then at that point it will be common knowledge to all three that C knows she is dirty.

2.3 More on Actions

In this section, we have a few remarks on actions. The point here is to clarify the relation between the scenarios of Section 1 and the intuitions concerning them, and the corresponding actions of Section 1.2.

First and foremost, here are the the conceptual points involved in our formalization. The idea is that epistemic actions present a lot of uncertainty. Indeed, what might be thought of as a single action (or event) is naturally interpreted by agents in different ways. The various agents might be unclear on what exactly happened, and again they might well have different interpretations on what is happening. Our formalization reflects this by making epistemic actions into Kripke models. So our use of possible-worlds modeling of actions is on a par with other uses of these models, and it inherits all of the features and bugs of those approaches.

Next, we want to spell out in words what our proposal amounts to. The basic problem is to decide how to represent what happens to a Kripke model W after an announcement α . (Of

course, we are modeling α by an action in our formal sense.) Our solution begins by considering copies of W , one for each action token k of α in which $\text{PRE}(\alpha)$ holds. We can think of tagging the worlds of W with the worlds of α , and then we must give an account of the accessibility relation between them. The intuition is that the agents' relations to alternative worlds should be independent from their relations to other possibilities for α . So the accessibility relations of K and W should be combined *independently*. This is expressed formally in (3).

The auxiliary language $\hat{\mathcal{L}}$ has as atomic sentences all sentences φ of $\mathcal{L}([\alpha], \square^*)$. It has all boolean connectives, standard modal operators \square_A for $A \in \mathcal{A}$, and also group knowledge operators $\square_{\mathcal{B}}^*$ for $\mathcal{B} \subseteq \mathcal{A}$.

We interpret $\hat{\mathcal{L}}$ on actions using the standard clauses for the connectives and modal operators, and by interpreting the atomic sentences as follows $\langle K, k \rangle \models p$ iff $\text{PRE}(k) = p$.

Examples 2.3 The idea here is that the auxiliary language formalizes talk about what the different agents think is happening in our announcements. We refer back to the actions of Section 1.2. For example, $\alpha \models \square_{\{A,B,C\}}^* \psi$. Intuitively, in α , it is common knowledge that ψ was announced. Another example: that

$$\delta \models \square_{\{A,B,C\}}^* \diamond_C (\psi \vee \psi_A \vee \psi_B).$$

That is, in δ , it is common knowledge that C thinks it is possible that some non-trivial announcement happened. Recall that this was one of our basic intuitions about δ , one which is not in general statable in our main language $\mathcal{L}([\alpha], \square^*)$.

Definition Let $\langle K, k \rangle$ be a model-world pair, and let φ be a sentence of $\hat{\mathcal{L}}$. Then χ *characterizes* $\langle K, k \rangle$ iff for all $\langle L, l \rangle$, $\langle L, l \rangle \models \chi$ iff $\langle L, l \rangle \equiv \langle K, k \rangle$.

Proposition 2.4 *Let $\langle K, k \rangle$ be a model-world pair with K finite. Then there is a sentence χ of $\hat{\mathcal{L}}$ which characterizes $\langle K, k \rangle$.*

Proof By replacing $\langle K, k \rangle$ by its quotient under the largest auto-bisimulation, we may assume that if $l \neq m$, then $\langle K, l \rangle \not\equiv \langle K, m \rangle$. It is well-known that the relation of elementary equivalence in modal logic is a bisimulation on models in which each world has finitely many arrows coming in and out. It follows from this and the overall finiteness of K that we can find sentences φ_l for $l \in K$ with the property that for all l and m , $\langle K, m \rangle \models \varphi_l$ iff $m = l$. Let ψ be the following sentence

$$\psi \equiv \bigwedge_{l \in K, A \in \mathcal{A}} \left(\varphi_l \rightarrow \square_A \bigvee_{l \xrightarrow{A} l'} \varphi_{l'} \wedge \bigwedge_{l \xrightarrow{A} l'} \diamond_A \varphi_{l'} \right)$$

Going back to our original $\langle K, k \rangle$, let χ be $\varphi_k \wedge \square_{\mathcal{A}}^* \psi$. It is easy to check that each $\langle K, l \rangle$ satisfies ψ ; hence each satisfies $\square_{\mathcal{A}}^* \psi$. Therefore $\langle K, k \rangle \models \chi$. We claim that χ characterizes $\langle K, k \rangle$. To see this, suppose that $\langle J, j \rangle \models \chi$. Consider the relation $R \subseteq K \times J$ given by

$$k' R j' \quad \text{iff} \quad \langle J, j' \rangle \models \varphi_{k'} \wedge \square_{\mathcal{A}}^* \psi.$$

It is sufficient to see that R is a bisimulation. We'll verify half of this: suppose that $k' R j'$ and $j' \xrightarrow{A} j''$. By using ψ , we see that there is some k'' such that $k' \xrightarrow{A} k''$ and $j'' \models \varphi_{k''}$. And also, since $\models \square_{\mathcal{A}}^* \psi \rightarrow \square_A \square_{\mathcal{A}}^* \psi$, we see that $\langle J, j'' \rangle \models \square_{\mathcal{A}}^* \psi$. This completes the proof. \dashv

The connection of this result and our discussion of actions is that it is often difficult to go from an informal description of an epistemic action to a formal one along our lines. (For example, our formulation of δ was the last of several versions.) Presumably, one way to get a formal action in our sense is to think carefully about which properties the action should have, express them in the auxiliary language, and then write a characterizing sentence such as ψ in the proof of Proposition 2.4. Then one can construct the finite model by standard methods. Although this would be a tedious process, it seems worthwhile to know that it is available.

Our formalization of actions reflects some choices which one might wish to modify. One of these choices is to take the range of the function PRE to be some language. Another option would be to have the range to be the *power set* of that language. This would make actions into Kripke models over the whole set of sentences. (And so what we have done is like considering modal logic with the restriction that at any world satisfies exactly one atomic sentence.) Taking this other option thus brings actions and models closer. This idea is pursued in Baltag [1], a continuation of this work which develops a “calculus of epistemic actions.” This replaces the “semantic” actions of this paper with *action expressions*. These expressions have nicer properties than the auxiliary language of this paper, but it would take us too far afield to discuss this further.

On a different matter, it makes sense to restrict attention from the full collection of actions as we have defined it to the smaller collection of *S5 actions*, where each accessibility \triangleleft is an equivalence relation. This corresponds to the standard move of restricting attention to models with this property, and the reasons for doing this are similar. Intuitively, an S5 action is one in which every agent is introspective (with respect to their own suspicions about actions). Moreover, the introspection is accurate, and this fact is common knowledge.

A final modification which is quite natural is to allow actions which *change the world*. One would do this by adding to our notion of action a *sentential update* u . This would be a function defined on AtSen and written in terms of update equations such as $u(p) := p \wedge q$; $u(q) = \text{false}$, etc. We are confident that our logical systems can be modified to reflect this change, and we discuss this at certain points below. We decided not to make this change mostly in order to keep the basic notions as simple as possible.

With respect to both of the changes mentioned in the last two paragraphs, it is not hard to modify our logical work to get completeness results for the new systems. We discuss all of this in Section 5.3.

3 A Logic for $\mathcal{L}([\alpha])$

In Figure 2 below we present a logic for $\mathcal{L}([\alpha], \Box^*)$ which we shall study later. In this section, we shall restrict the logic to the simpler language $\mathcal{L}([\alpha])$. We do so partly to break up the study of a system with many axioms and rules, and partly to emphasize the significance of adding the infinitary operators \Box_B^* to $\mathcal{L}([\alpha])$. To carry out the restriction, we forget the axioms and rules of inference in Figure 2 which are marked by a *. In particular $\alpha \circ \beta$ will be defined later (Section 4).

The rules of the system are all quite standard from modal logic. The Action Axioms are the interesting new ones. In the Atomic Permanence axiom, p is an atomic sentence. The axiom then says that announcements do not change the brute fact of whether or not p holds. This

Basic Axioms	
All sentential validities	
($[\alpha]$ -normality)	$\vdash [\alpha](\varphi \rightarrow \psi) \rightarrow ([\alpha]\varphi \rightarrow [\alpha]\psi)$
(\Box_A -normality)	$\vdash \Box_A(\varphi \rightarrow \psi) \rightarrow (\Box_A\varphi \rightarrow \Box_A\psi)$
* (\Box_C^* -normality)	$\vdash \Box_C^*(\varphi \rightarrow \psi) \rightarrow (\Box_C^*\varphi \rightarrow \Box_C^*\psi)$
Action Axioms	
(Atomic Permanence)	$\vdash [\alpha]p \leftrightarrow (\text{PRE}(\alpha) \rightarrow p)$
(Partial Functionality)	$\vdash [\alpha]\neg\chi \leftrightarrow (\text{PRE}(\alpha) \rightarrow \neg[\alpha]\chi)$
(Action-Knowledge)	$\vdash [\alpha]\Box_A\varphi \leftrightarrow (\text{PRE}(\alpha) \rightarrow \bigwedge\{\Box_A[\beta]\varphi : \alpha \xrightarrow{\Delta} \beta\})$
* Mix Axiom	$\vdash \Box_C^*\varphi \rightarrow \varphi \wedge \bigwedge\{\Box_A\Box_C^*\varphi : A \in \mathcal{C}\}$
* Composition Axiom	$\vdash [\alpha][\beta]\varphi \leftrightarrow [\alpha \circ \beta]\varphi$
Modal Rules	
(Modus Ponens)	From $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$, infer $\vdash \psi$
($[\alpha]$ -necessitation)	From $\vdash \psi$, infer $\vdash [\alpha]\psi$
(\Box_A -necessitation)	From $\vdash \varphi$, infer $\vdash \Box_A\varphi$
* (\Box_C^* -necessitation)	From $\vdash \varphi$, infer $\vdash \Box_C^*\varphi$
* Action Rule	
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 $\alpha \rightarrow_C^* \beta$ (including α itself), and such that	
1. $\vdash \chi_\beta \rightarrow [\beta]\psi$.	
2. If $A \in \mathcal{C}$ and $\beta \xrightarrow{\Delta} \gamma$, then $\vdash (\chi_\beta \wedge \text{PRE}(\beta)) \rightarrow \Box_A\chi_\gamma$.	
From these assumptions, infer $\vdash \chi_\alpha \rightarrow [\alpha]\Box_C^*\psi$.	

Figure 2: The logical system for $\mathcal{L}([\alpha], \Box^*)$. For $\mathcal{L}([\alpha])$, we drop the * axioms and rules.

axiom reflects the fact that our actions do not change any kind of local state. (We discuss an extension of our system in Section 5.3 where this axiom is not sound.) The Partial Functionality Axiom corresponds to the fact that the operation $\langle W, w \rangle \mapsto \langle W, w \rangle^\alpha$ is a partial function. The key axiom of the system is the Action-Knowledge Axiom, giving a criterion for knowledge after an announcement. We will check soundness of this axiom leaving checking soundness of other unstarred axioms and rules to the reader.

Proposition 3.1 *The Action-Knowledge Axiom*

$$[\alpha]\Box_A\varphi \leftrightarrow (\text{PRE}(\alpha) \rightarrow \bigwedge\{\Box_A[\beta]\varphi : \alpha \triangleleft \beta\})$$

is sound.

Proof We remind the reader that the relevant definitions and notation are found in Section 2.2. Let α be the action $\langle K, k \rangle$. Fix a pair $\langle W, w \rangle$. If $\langle W, w \rangle \models \neg\text{PRE}(\alpha)$, then both sides of our biconditional hold. We therefore assume that $\langle W, w \rangle \models \text{PRE}(\alpha)$ in the rest of this proof. Assume that $\langle W, w \rangle^\alpha \models \Box_A\varphi$. Take some β such that $\alpha \triangleleft \beta$. This β is of the form $\langle K, k' \rangle$ for some k' such that $k \triangleleft k'$. Let $w \triangleleft w'$. We have two cases: $\langle W, w' \rangle \models \text{PRE}(k')$, and $\langle W, w' \rangle \models \neg\text{PRE}(k')$. In the latter case, $\langle W, w' \rangle \models [\beta]\varphi$ trivially. We'll show this in the former case, so assume $\langle W, w' \rangle \models \text{PRE}(k')$. Then $\langle w', k' \rangle$ is a world of W^K , and indeed $\langle w, k \rangle \triangleleft \langle w', k' \rangle$. Now our assumption that $\langle W, w \rangle^\alpha \models \Box_A\varphi$ implies that $\langle W^K, (w', k') \rangle \models \varphi$. This means that $\langle W, w' \rangle^\beta \models \varphi$. Hence $\langle W, w' \rangle \models [\beta]\varphi$. Since β and w' were arbitrary, $\langle W, w \rangle \models \bigwedge_\beta \Box_A[\beta]\varphi$.

The other direction is similar. \dashv

The rest of this section is devoted to the completeness result for $\mathcal{L}([\alpha])$. The reader not interested in this may omit the rest of this section, but at some points later we will refer back to the *term rewriting system* \mathcal{R} which we shall describe shortly. Our completeness proof is based on a translation of $\mathcal{L}([\alpha])$ to ordinary modal logic \mathcal{L} . And this translation is based on a term rewriting system to be called \mathcal{R} .

The rewriting rules of \mathcal{R} are:

$$\begin{aligned} [\alpha]p &\rightsquigarrow \text{PRE}(\alpha) \rightarrow p \\ [\alpha]\neg\psi &\rightsquigarrow \text{PRE}(\alpha) \rightarrow \neg[\alpha]\psi \\ [\alpha](\psi \wedge \chi) &\rightsquigarrow [\alpha]\psi \wedge [\alpha]\chi \\ [\alpha]\Box_A\psi &\rightsquigarrow \text{PRE}(\alpha) \rightarrow \bigwedge\{\Box_A[\beta]\psi : \alpha \triangleleft \beta\} \end{aligned}$$

As in all rewrite systems, we apply the rules of \mathcal{R} at arbitrary subsentences of a given sentence. (For example, consider what happens with something like $[\alpha][\beta]\varphi$. We might rewrite $[\beta]\varphi$ using some rule, say to ψ . Then we might rewrite $[\alpha]\psi$ to something like $[\gamma]\psi$, etc.)

Lemma 3.2 *There is a relation $<$ on the sentences of $\mathcal{L}([\alpha])$ such that*

1. $<$ is wellfounded.
2. For all rules $\varphi \rightsquigarrow \psi$ of \mathcal{R} , $\psi < \varphi$.
3. A sentence $\varphi \in \mathcal{L}([\alpha])$ is a normal form iff it is a modal sentence (that is, φ cannot be rewritten iff no actions occur in φ).

This takes some work, and because the details are less important than the facts themselves, we have placed the entire matter in an Appendix to this paper. (The Appendix also discusses an extension of the rewrite system \mathcal{R} to a system \mathcal{R}^* for the larger language $\mathcal{L}([\alpha], \Box^*)$, so if you read it at this point you will need to keep this in mind.)

In the next result, we let \mathcal{L} be ordinary modal logic over AtSen (where of course there are no actions).

Proposition 3.3 *There is a translation $t : \mathcal{L}([\alpha]) \rightarrow \mathcal{L}$ such that for all $\varphi \in \mathcal{L}([\alpha])$, φ is semantically equivalent to φ^t .*

Proof Every sentence φ of $\mathcal{L}([\alpha])$ may be rewritten to a normal form. By Lemma 3.2, the normal forms of φ is a sentence in \mathcal{L} . We therefore set φ^t to be any normal form of φ , say the one obtained by carrying out leftmost reductions. The semantic equivalence follows from the fact that the rewrite rules themselves are sound, and from the fact that semantic equivalence is preserved by substitutions. \dashv

Lemma 3.4 (Substitution) *Let φ be any sentence, and let $\vdash \chi \leftrightarrow \chi'$. Suppose that $\varphi[p/\chi]$ comes from φ by replacing p by χ at some point, and $\varphi[p/\chi']$ comes similarly. Then $\vdash \varphi[p/\chi] \leftrightarrow \varphi[p/\chi']$.*

Proof By induction on φ . The key point is that we have necessitation rules for each $[\alpha]$. \dashv

Theorem 3.5 *This logical system for $\mathcal{L}([\alpha])$ is strongly complete: $\Sigma \vdash \varphi$ iff $\Sigma \models \varphi$.*

Proof The soundness half being easy, we only need to show that if $\Sigma \models \varphi$, then $\Sigma \vdash \varphi$. First, $\Sigma^t \models \varphi^t$. Since our system extends the standard complete proof system of modal logic, $\Sigma^t \vdash \varphi^t$. Now for each χ of $\mathcal{L}([\alpha])$, $\vdash \chi \leftrightarrow \chi^t$. (This is an easy induction on $<$ using Lemma 3.4.) As a result, $\Sigma \vdash \chi^t$ for all $\chi \in \Sigma$. So $\Sigma \vdash \varphi^t$. As we know $\vdash \varphi^t \leftrightarrow \varphi$. So we have our desired conclusion: $\Sigma \vdash \varphi$. \dashv

Strong completeness results of this kind may also be found in Plaza [13] and in Gerbrandy and Groeneveld [7]. We discuss some of the history of the subject in Section 7.

4 A Logic for $\mathcal{L}([\alpha], \Box^*)$

At this point, we turn to the completeness result for $\mathcal{L}([\alpha], \Box^*)$. It is easy to check that there is no hope of getting a *strong* completeness result (where one has arbitrary sets of hypotheses). The best one can hope for is weak completeness: $\vdash \varphi$ if and only if $\models \varphi$. Also, in contrast to our translations results for $\mathcal{L}([\alpha])$, the larger language $\mathcal{L}([\alpha], \Box^*)$ cannot be translated into \mathcal{L} or even to $\mathcal{L}(\Box^*)$ (modal logic with extra modalities $\Box_{\mathcal{B}}^*$). We prove this in Theorem 6.2 below. So completeness results for $\mathcal{L}([\alpha], \Box^*)$ cannot simply be based on translation.

Our logical system is listed in Figure 2 above. We discussed the fragment of the system which does not have the $*$ axioms and rules in Section 3. The $\Box_{\mathcal{C}}^*$ -normality Axiom and $\Box_{\mathcal{C}}^*$ -necessitation Rule are standard, as is the Mix Axiom. We leave checking their soundness to the reader. The key features of the system are thus the Composition Axiom and the Action Rule. We begin with the Action Rule, restated below:

The Action Rule 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 $\alpha \rightarrow_{\mathcal{C}}^* \beta$ (including α itself), and such that

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

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

Remark We use $\rightarrow_{\mathcal{C}}^*$ as an abbreviation for the reflexive and transitive closure of the relation $\bigcup_{A \in \mathcal{C}} \xrightarrow{A}$. Recall that there are only finitely many β such that $\alpha \rightarrow_{\mathcal{C}}^* \beta$, since each is determined by a world of the same finite Kripke frame that determines α . So even though the Action Rule might look like it takes infinitely many premises, it really only takes finitely many.

Another point: if one so desires, the Action Rule could be replaced by a (more complicated) axiom scheme which we will not state here.

Lemma 4.1 $\langle W, w \rangle \models \langle \alpha \rangle \Diamond_{\mathcal{C}}^* \varphi$ iff there is a sequence of worlds from W

$$w = w_0 \xrightarrow{A_1} w_1 \xrightarrow{A_2} \cdots \xrightarrow{A_{k-1}} w_{k-1} \xrightarrow{A_k} w_k$$

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

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

such that $A_i \in \mathcal{C}$ and $\langle W, w_i \rangle \models \text{PRE}(\alpha_i)$ for all $0 \leq i \leq k$, and $\langle W, w_k \rangle \models \langle \alpha_k \rangle \varphi$.

Remark The case $k = 0$ just says that $\langle W, w \rangle \models \langle \alpha \rangle \Diamond_{\mathcal{C}}^* \varphi$ is implied by $\langle W, w \rangle \models \langle \alpha \rangle \varphi$.

Proof $\langle W, w \rangle \models \langle \alpha \rangle \Diamond_{\mathcal{C}}^* \varphi$ iff $\langle W, w \rangle \models \text{PRE}(\alpha)$ and $\langle W^\alpha, (w, \alpha) \rangle \models \Diamond_{\mathcal{C}}^* \varphi$; iff $\langle W, w \rangle \models \text{PRE}(\alpha)$ and there is a sequence in W^α ,

$$(w, \alpha) = v_0 \xrightarrow{A_1} v_1 \xrightarrow{A_2} \cdots \xrightarrow{A_{k-1}} v_{k-1} \xrightarrow{A_k} v_k$$

where $k \geq 0$ such that $A_i \in \mathcal{C}$ and $\langle W^\alpha, v_k \rangle \models \varphi$. Now suppose such sequences exist in W^α . Then we get a sequence of worlds w_i in W and actions α_i such that $v_i = (w_i, \alpha_i)$ and $\langle W, w_i \rangle \models \text{PRE}(\alpha_i)$. The condition that $\langle W^\alpha, v_k \rangle \models \varphi$ translates to $\langle W, w_k \rangle \models \langle \alpha_k \rangle \varphi$. Conversely, if we have a sequence in W with these properties, we get one in W^α by taking $v_i = (w_i, \alpha_i)$. \dashv

Proposition 4.2 *The Action Rule is sound.*

Proof Assume that $\langle W, w \rangle \models \chi_\alpha$ but also $\langle W, w \rangle \models \langle \alpha \rangle \Diamond_{\mathcal{C}}^* \neg\psi$. According to Lemma 4.1, there is a labeled sequence of worlds from W

$$w = w_0 \xrightarrow{A_1} w_1 \xrightarrow{A_2} \cdots \xrightarrow{A_{k-1}} w_{k-1} \xrightarrow{A_k} w_k$$

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

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

such that $\langle W, w_i \rangle \models \text{PRE}(\alpha_i)$ for all $0 < i \leq k$, and $\langle W, w_k \rangle \models \langle \alpha_k \rangle \neg\psi$. If $k = 0$, we have $\langle W, w \rangle \models \langle \alpha \rangle \neg\psi$. But since $\vdash \chi_\alpha \rightarrow [\alpha]\psi$, we have $\langle W, w \rangle \models [\alpha]\psi$. This is a contradiction.

Now we argue the case $k > 0$. We show by induction on $1 \leq i \leq k$ that $\langle W, w_i \rangle \models \chi_{\alpha_i} \wedge [\alpha_i]\psi$. In particular, $\langle W, w_k \rangle \models [\alpha_k]\psi$. This is a contradiction. \dashv

We close with a discussion of the Composition Rule, beginning with a general definition.

Definition Let $\alpha = \langle K, k \rangle$ and $\beta = \langle L, l \rangle$ be actions. Then the *action composition* $\alpha \circ \beta$ is the action defined as follows. Consider the product set $K \times L$. We turn this into a Kripke frame using the restriction of the product arrows. We get an action structure by setting

$$\text{PRE}(\langle k', l' \rangle) = \text{PRE}(k') \wedge [\langle K, k' \rangle] \text{PRE}(l').$$

Finally, we set $\alpha \circ \beta = \langle K \times L, (k, l) \rangle$.

Proposition 4.3 *Concerning the composition operation:*

1. $(W^\alpha)^\beta \cong W^{\alpha \circ \beta}$ via the restriction of $((w, k'), l') \mapsto (w, (k', l'))$ to $(W^\alpha)^\beta$.
2. The Composition Axiom is sound: $[\alpha][\beta]\varphi \leftrightarrow [\alpha \circ \beta]\varphi$.
3. $\alpha \circ (\beta \circ \gamma) \cong (\alpha \circ \beta) \circ \gamma$.
4. $\alpha \circ \tau \cong \alpha \cong \tau \circ \alpha$, where the null action τ is from Section 1.2.

Proof Let $\alpha = \langle K, k \rangle$ and $\beta = \langle L, l \rangle$. For (1), note that the worlds of $(W^\alpha)^\beta$ are of the form $((w, k'), l')$, where $(w, k') \in W^\alpha$ and $\langle W^\alpha, (w, k') \rangle \models \text{PRE}(l')$. For such $((w, k'), l')$, $\langle W, w \rangle \models \text{PRE}(k')$ and $\langle W, w \rangle \models [\langle K, k' \rangle] \text{PRE}(l')$. That is, $(w, (k', l')) \in W^{\alpha \circ \beta}$. The converse is similar, and the rest of the isomorphism properties are easy.

Part (2) follows from (1). We use the obvious isomorphism $((k, l), m) \mapsto (k, (l, m))$ in part (3). We use the Composition and $[\alpha]$ -necessitation axioms to show that this isomorphism preserves the PRE function up to logical equivalence. Part (4) is easy, using the fact that $\models [\tau]\varphi \leftrightarrow \varphi$. \dashv

Extending the rewriting system \mathcal{R} to $\mathcal{L}([\alpha], \square^*)$. We consider $\mathcal{L}([\alpha], \square^*)$. The rewriting system \mathcal{R} extends naturally to this larger language, taking new symbols for the operators $\square_{\mathcal{B}}^*$. We also add a rule corresponding to the Composition Axiom: $[\alpha][\beta]\varphi \rightsquigarrow [\alpha \circ \beta]\varphi$. We call this rewriting system \mathcal{R}^* .

Lemma 4.4 *There is a relation $<$ on the sentences and actions of $\mathcal{L}([\alpha], \square^*)$ such that*

1. $<$ is wellfounded.
2. For all rules $\varphi \rightsquigarrow \psi$ of \mathcal{R}^* , $\psi < \varphi$.
3. If ψ is a proper subsentence of φ , then $\psi < \varphi$.
4. A sentence $\varphi \in \mathcal{L}([\alpha], \square^*)$ is a normal form iff it is built from atomic sentences using \neg , \wedge , \square_A , and $\square_{\mathcal{B}}^*$, or if it is of the form $[\alpha]\square_{\mathcal{B}}^*\psi$, where α is an action in normal form, and ψ too is in normal form.
5. An action α is a normal form if whenever $\alpha \xrightarrow{*} \beta$, $\text{PRE}(\beta)$ is a normal form sentence.

6. If $\alpha \rightarrow \beta$, then $[\alpha]\Box_{\mathcal{C}}^*\psi > [\beta]\psi$.

7. $nf(\varphi) \leq \varphi$.

Once again, the details are in the Appendix.

In Section 3, we saw a translation t from $\mathcal{L}([\alpha])$ to \mathcal{L} can be extended to a translation from $\mathcal{L}([\alpha], \Box^*)$ to the infinitary language \mathcal{L}_∞ , where we have countable conjunctions and disjunctions. This extension is defined using Part (4) of Lemma 4.4. The additional clauses in the definition of t are

$$\begin{aligned} (\Box_{\mathcal{B}}^*\varphi)^t &= \bigwedge_{\langle A_1, \dots, A_n \rangle \in \mathcal{B}^*} (\Box_{A_1} \cdots \Box_{A_n} \varphi)^t \\ ([\alpha]\Box_{\mathcal{B}}^*\psi)^t &= \bigwedge_{\langle A_1, \dots, A_n \rangle \in \mathcal{B}^*} ([\alpha]\Box_{A_1} \cdots \Box_{A_n} \psi)^t \end{aligned}$$

In this way, we see that $\mathcal{L}([\alpha], \Box^*)$ may be regarded as a fragment of infinitary modal logic.

Remark It is possible to drop the Composition Axiom in favor of a more involved version of the Action Rule. The point is the Composition Axiom simplifies the normal forms of the $\mathcal{L}([\alpha], \Box^*)$: Without the Composition Axiom, the normal forms of sentences of $\mathcal{L}([\alpha], \Box^*)$ would be of the form $[\alpha_1][\alpha_2] \cdots [\alpha_r]\psi$, where each α_i is a normal form action and ψ is a normal form sentence. The Composition Axiom insures that the normal forms are of the form $[\alpha]\psi$. So if we were to drop the Composition Axiom, we would need a formulation of the Action Rule 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.

This completes the discussion of the axioms and rules of our logical system for $\mathcal{L}([\alpha], \Box^*)$.

5 Completeness for $\mathcal{L}([\alpha], \Box^*)$

In this section, we prove the completeness of the logical system for $\mathcal{L}([\alpha], \Box^*)$. Section 5.1 has some technical results which culminate in the Substitution Lemma 5.3. This is used in some of our work on normal forms in the Appendix, and that work figures in the completeness theorem of Section 5.2.

5.1 Some Syntactic Results

Lemma 5.1 *For all $A \in \mathcal{C}$ and all β such that $\alpha \rightarrow_A \beta$,*

1. $\vdash [\alpha]\Box_{\mathcal{C}}^*\psi \rightarrow [\alpha]\psi$.
2. $\vdash [\alpha]\Box_{\mathcal{C}}^*\psi \wedge \text{PRE}(\alpha) \rightarrow \Box_A[\beta]\Box_{\mathcal{C}}^*\psi$.

Proof Part (1) follows easily from the Mix Axiom and modal reasoning. For part (2), we start with a consequence of the Mix Axiom: $\vdash \Box_{\mathcal{C}}^*\psi \rightarrow \Box_A\Box_{\mathcal{C}}^*\psi$. Then by modal reasoning, $\vdash [\alpha]\Box_{\mathcal{C}}^*\psi \rightarrow [\alpha]\Box_A\Box_{\mathcal{C}}^*\psi$. By the Action-Knowledge Axiom, , we have $\vdash [\alpha]\Box_{\mathcal{C}}^*\psi \wedge \text{PRE}(\alpha) \rightarrow \Box_A[\beta]\Box_{\mathcal{C}}^*\psi$. ◻

Definition Let α and α' be actions. We write $\vdash \alpha \leftrightarrow \alpha'$ if α and α' are based on the same Kripke frame W and the same world w , and if for all $v \in W$, $\vdash \text{PRE}(v) \leftrightarrow \text{PRE}'(v)$, where PRE is the announcement function for α , and PRE' for α' .

We note the following bisimulation-like properties:

1. If $\vdash \alpha \leftrightarrow \alpha'$, then also $\vdash \text{PRE}(\alpha) \leftrightarrow \text{PRE}(\alpha')$.
2. Whenever β' is such that $\alpha' \rightarrow_{\mathcal{C}}^* \beta'$, then there is some β such that $\vdash \beta \leftrightarrow \beta'$ and $\alpha \rightarrow_{\mathcal{C}}^* \beta$.

These follow easily from the way we defined PRE on actions in terms of functions on frames.

Lemma 5.2 *If $\vdash \alpha \leftrightarrow \alpha'$, then for all ψ , $\vdash [\alpha]\psi \leftrightarrow [\alpha']\psi$.*

Proof By induction on ψ . For ψ 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. Assuming the result for ψ gives the result for $[\chi]\psi$, using the Composition Axiom and the induction hypothesis. This leaves the step for sentences of the form $\Box_{\mathcal{B}}^* \psi$, assuming the result for ψ . We use the Action Rule to show that $\vdash [\alpha]\Box_{\mathcal{C}}^* \psi \rightarrow [\alpha']\Box_{\mathcal{C}}^* \psi$. For each β' , we let $\chi_{\beta'}$ be $[\beta]\Box_{\mathcal{C}}^* \psi$, where β is such that $\vdash \beta \leftrightarrow \beta'$. We need to show that for all relevant β' and γ' ,

- a. $\vdash [\beta]\Box_{\mathcal{C}}^* \psi \rightarrow [\beta']\psi$; and
- b. If $\beta' \triangleq \gamma'$, then $\vdash [\beta]\Box_{\mathcal{C}}^* \psi \wedge \text{PRE}(\beta') \rightarrow \Box_A[\gamma]\Box_{\mathcal{C}}^* \psi$.

For (a), we know from Lemma 5.1, part (1) that $\vdash [\beta]\Box_{\mathcal{C}}^* \psi \rightarrow [\beta]\psi$. By induction hypothesis on ψ , $\vdash [\beta]\psi \leftrightarrow [\beta']\psi$. And this implies (a). For (b), Lemma 5.1, part (2) tells us that under the assumptions,

$$\vdash [\beta]\Box_{\mathcal{C}}^* \psi \wedge \text{PRE}(\beta) \rightarrow \Box_A[\gamma]\Box_{\mathcal{C}}^* \psi.$$

As we know, $\vdash \text{PRE}(\beta) \leftrightarrow \text{PRE}(\beta')$. This implies (b).

This completes the induction on ψ . ⊣

Lemma 5.3 (Substitution) *Let t be a sentence or action of $\mathcal{L}([\alpha], \Box^*)$, and let $\vdash \chi \leftrightarrow \chi'$. Suppose that $t[p/\chi]$ comes from t by replacing p by χ at some point, and $t[p/\chi']$ comes similarly. Then $\vdash t[p/\chi] \leftrightarrow t[p/\chi']$.*

Proof By induction on t , using Lemma 5.2. ⊣

Lemma 5.4 *For every sentence $\varphi \in \mathcal{L}([\alpha], \Box^*)$ there is some normal form $\text{nf}(\varphi) \leq \varphi$ such that $\vdash \varphi \leftrightarrow \text{nf}(\varphi)$.*

Proof Given φ , there is a finite sequence $\varphi_0 \rightsquigarrow \dots \rightsquigarrow \varphi_n = \varphi'$ such that $\varphi_0 = \varphi$, and φ_n is in normal form. This is a consequence of the fact that $<$ is wellfounded and the rules of the system are reducing. By Lemma 5.3, we see that for all i , $\vdash \varphi_i \leftrightarrow \varphi'_i$. ⊣

5.2 Completeness

The proof of completeness and decidability is based on the filtration argument for completeness of PDL due to Kozen and Parikh [10]. 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 4.

Definition Let $s(\varphi)$ be the set of subsentences of φ , including φ itself. This includes all sentences occurring in actions which occur in φ and their subsentences. For future use, we note that

$$s([\alpha]\square_{\mathcal{C}}^*\varphi) = \{[\alpha]\square_{\mathcal{C}}^*\varphi, \square_{\mathcal{C}}^*\varphi\} \cup s(\varphi) \cup \bigcup\{s(\text{PRE}(\beta)) : \alpha \rightarrow_{\mathcal{C}}^* \beta\} \quad (4)$$

We define a function $f : \mathcal{L}([\alpha], \square^*) \rightarrow \mathcal{P}(\mathcal{L}([\alpha], \square^*))$ by recursion on the wellfounded relation $<$ as follows: For *normal forms*, f works as follows:

$$\begin{aligned} f(p) &= \{p\} \\ f(\neg\varphi) &= f(\varphi) \cup \{\neg\varphi\} \\ f(\varphi \wedge \psi) &= f(\varphi) \cup f(\psi) \cup \{\varphi \wedge \psi\} \\ f(\square_A\varphi) &= f(\varphi) \cup \{\square_A\varphi\} \\ f(\square_{\mathcal{B}}^*\varphi) &= f(\varphi) \cup \{\square_{\mathcal{B}}^*\varphi\} \cup \{\square_A\square_{\mathcal{B}}^*\varphi : A \in \mathcal{B}\} \\ f([\alpha]\square_{\mathcal{C}}^*\varphi) &= \{\square_A[\beta]\square_{\mathcal{C}}^*\varphi : \alpha \rightarrow_{\mathcal{C}}^* \beta \ \& \ A \in \mathcal{C}\} \\ &\cup \{[\beta]\square_{\mathcal{C}}^*\varphi : \alpha \rightarrow_{\mathcal{C}}^* \beta \ \& \ A \in \mathcal{C}\} \\ &\cup \bigcup\{f(\chi) : (\exists\beta) \alpha \rightarrow_{\mathcal{C}}^* \beta \ \& \ \chi \in s(\text{PRE}(\beta))\} \\ &\cup f(\square_{\mathcal{C}}^*\varphi) \\ &\cup \bigcup\{f([\beta]\varphi) : \alpha \rightarrow_{\mathcal{C}}^* \beta\} \end{aligned}$$

For φ not in normal form, let $f(\varphi) = f(nf(\varphi))$. (Note that we need to define f on sentences which are not normal forms, because $f([\beta]\psi)$ figures in $f([\alpha]\square_{\mathcal{C}}^*\varphi)$. Also, the definition makes sense because the calls to f on the right-hand sides are all $<$ the arguments on the left-hand sides, and since $nf(\varphi) \leq \varphi$ for all φ ; see Lemma 4.4.)

Lemma 5.5 *For all φ :*

1. $f(\varphi)$ is a finite set of normal form sentences.
2. $nf(\varphi) \in f(\varphi)$.
3. If $\psi \in f(\varphi)$, then $f(\psi) \subseteq f(\varphi)$.
4. If $\psi \in f(\varphi)$, then $s(\psi) \subseteq f(\varphi)$.
5. If $[\gamma]\square_{\mathcal{C}}^*\chi \in f(\varphi)$, $\gamma \rightarrow_{\mathcal{C}}^* \delta$, and $A \in \mathcal{C}$, then $f(\varphi)$ also contains $\square_A[\delta]\square_{\mathcal{C}}^*\chi$, $[\delta]\square_{\mathcal{C}}^*\chi$, $\text{PRE}(\delta)$, and $nf([\delta]\chi)$.

Proof All of the parts are by induction on φ in the well-order $<$. For part (1), note that if

$[\alpha]\Box_{\mathcal{C}}^*\psi$ is a normal form, then each sentence $\Box_A[\beta]\Box_{\mathcal{C}}^*\psi$ and all subsentences of this sentence are normal forms. For part (2), note that when φ is a normal form, $\varphi \in f(\varphi)$.

In part (3), 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

$$f(\psi) = f(\Box_{\mathcal{B}}^*\varphi) \cup \{\Box_A\Box_{\mathcal{B}}^*\varphi\} \subseteq f(\Box_{\mathcal{B}}^*\varphi).$$

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 are easily done by the induction hypothesis. This covers all of the cases except for $\psi = [\beta]\Box_{\mathcal{C}}^*\varphi$ and $\psi = \Box_A[\beta]\Box_{\mathcal{C}}^*\varphi$. For the first of these, we use the transitivity of $\rightarrow_{\mathcal{C}}^*$ to check that $f([\beta]\Box_{\mathcal{C}}^*\varphi) \subseteq f([\alpha]\Box_{\mathcal{C}}^*\varphi)$. And now the second case follows:

$$f(\Box_A[\beta]\Box_{\mathcal{C}}^*\varphi) = f([\beta]\Box_{\mathcal{C}}^*\varphi) \cup \{\Box_A[\beta]\Box_{\mathcal{C}}^*\varphi\} \subseteq f([\alpha]\Box_{\mathcal{C}}^*\varphi).$$

Part (4) is similar to part (3), using equation (4) at the beginning of this subsection.

For part (5), assume that $[\gamma]\Box_{\mathcal{C}}^*\chi \in f(\varphi)$. By part (1), $[\gamma]\Box_{\mathcal{C}}^*\chi$ is a normal form. We show that $\Box_A[\delta]\Box_{\mathcal{C}}^*\chi$, $[\delta]\Box_{\mathcal{C}}^*\chi$, $\text{PRE}(\delta)$, and $nf([\delta]\chi)$ all belong to $f([\gamma]\Box_{\mathcal{C}}^*\chi)$, and then use part (3). The first two of these sentences are immediate by the definition of f ; the third one follows from part (4); and the last comes from part (2) since $nf([\delta]\chi) \in f([\delta]\chi) \subseteq f([\gamma]\Box_{\mathcal{C}}^*\chi$. \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 $\Delta = \{\psi_1, \dots, \psi_n\}$. Given a maximal consistent set U of $\mathcal{L}([\alpha], \Box^*)$, let

$$\llbracket U \rrbracket = \pm\psi_1 \wedge \dots \wedge \pm\psi_n,$$

where the signs are taken in accordance with membership in U . That is, if $\psi_i \in U$, then ψ is a conjunct of $\llbracket U \rrbracket$; but if $\psi_i \notin U$, then $\neg\psi_i$ is a conjunct.

Two (standard) observations are in order. Notice that if $\llbracket U \rrbracket \neq \llbracket V \rrbracket$, then $\llbracket U \rrbracket \wedge \llbracket V \rrbracket$ is inconsistent. Also, for all $\psi \in \Delta$,

$$\vdash \psi \leftrightarrow \bigvee \{ \llbracket W \rrbracket : W \text{ is maximal consistent and } \psi \in W \}. \quad (5)$$

and

$$\vdash \neg\psi \leftrightarrow \bigvee \{ \llbracket W \rrbracket : W \text{ is maximal consistent and } \neg\psi \in W \}. \quad (6)$$

(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 $\mathcal{L}([\alpha], \Box^*)$, and the equivalence relation is $U \equiv V$ iff $\llbracket U \rrbracket = \llbracket V \rrbracket$ (iff $U \cap \Delta = V \cap \Delta$). We set $\langle \mathcal{F}, [U] \rangle \models p$ iff $p \in U \cap \Delta$. Furthermore,

$$[U] \triangleleft [V] \text{ in } \mathcal{F} \quad \text{iff} \quad \text{whenever } \Box_A\psi \in U \cap \Delta, \text{ then also } \psi \in V. \quad (7)$$

This condition is independent of the choice of representatives: we use part (4) of Lemma 5.5 to see that if $\Box_A \chi \in \Delta$, then also $\chi \in \Delta$.

A good path from $[V_0]$ for $\langle \alpha \rangle \diamond_{\mathcal{C}}^* \psi$ is a path in \mathcal{F}

$$[V_0] \rightarrow_{A_1} [V_1] \rightarrow_{A_2} \cdots \rightarrow_{A_{k-1}} [V_{k-1}] \rightarrow_{A_k} [V_k]$$

such that $k \geq 0$, each $A_i \in \mathcal{C}$, and such that there exist actions

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

such that $\text{PRE}(\alpha_i) \in V_i$ 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 4.1 in \mathcal{F} . Of course, the special case of that result would require that $\langle \mathcal{F}, [V_i] \rangle \models \text{PRE}(\alpha_i)$ rather than $\text{PRE}(\alpha_i) \in V_i$, 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 5.8).

Lemma 5.6 *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 5.1, part (1), we have $[\alpha] \psi \in V_0$. 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 5.5, 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 $[\alpha] \Box_{\mathcal{C}}^* \psi \in V_0$. By Lemma 5.1, part (2), V_0 contains $[\alpha] \Box_{\mathcal{C}}^* \psi \wedge \text{PRE}(\alpha) \rightarrow \Box_A [\alpha_1] \Box_{\mathcal{C}}^* \psi$. So V_0 contains $\Box_A [\alpha_1] \Box_{\mathcal{C}}^* \psi$. Again, this sentence belongs to Δ by Lemma 5.5, 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. \dashv

Lemma 5.7 *If $\llbracket V_0 \rrbracket \wedge \langle \alpha \rangle \diamond_{\mathcal{C}}^* \psi$ 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

$$\chi_\beta = \bigvee \{ \llbracket W \rrbracket : W \in S_\beta \}.$$

Note that $\neg \chi_\beta$ is logically equivalent to $\bigvee \{ \llbracket W' \rrbracket : [W'] \in \mathcal{F} \text{ and } W' \notin S_\beta \}$. Since we assumed $V_0 \in S_\alpha$, we have $\vdash \llbracket V_0 \rrbracket \rightarrow \chi_\alpha$.

We first claim that $\chi_\beta \wedge \langle \beta \rangle \psi$ 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, $\vdash \langle \beta \rangle \psi \rightarrow \text{PRE}(\beta)$. 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 will need the following standard claim, an argument for which can be found in Kozen and Parikh [10]. We will also use this claim in the proof of Lemma 5.8.

Claim If $\llbracket U \rrbracket \wedge \Box_A \llbracket V \rrbracket$ is consistent, then $[U] \rightarrow_A [V]$.

Proof of Claim Assume $\Box_A \psi \in U \cap \Delta$. If $\psi \notin V$, then $\neg\psi \in V$, so since $\psi \in \Delta$, $\vdash \llbracket V \rrbracket \rightarrow \neg\psi$. Thus, $\vdash \Diamond_A \llbracket V \rrbracket \rightarrow \Diamond_A \neg\psi$, and so $\vdash \llbracket U \rrbracket \wedge \Diamond_A \llbracket V \rrbracket \rightarrow \Box_A \psi \wedge \Diamond_A \neg\psi$, whence $\llbracket U \rrbracket \wedge \Diamond_A \llbracket V \rrbracket$ is inconsistent. This contradiction establishes the claim.

We next show that for all $A \in \mathcal{C}$ and all β such that $\beta \rightarrow_A \gamma$, $\chi_\beta \wedge \text{PRE}(\beta) \wedge \Diamond_A \neg\chi_\gamma$ is inconsistent. Otherwise, there would be $[W] \in S_\beta$ with χ_β , $\text{PRE}(\beta)$, and $\Diamond_A \neg\chi_\gamma$ in it. Then $\bigvee \{\Diamond_A \llbracket W' \rrbracket : W' \notin S_\gamma\}$, being equivalent to $\Diamond_A \neg\chi_\beta$, would belong to W . It follows that $\Diamond_A \llbracket W' \rrbracket \in W$ for some $W' \notin S_\gamma$. By the claim, $[W] \rightarrow_A [W']$. Since $[W'] \notin S_\gamma$, there is a good path from $[W']$ for $\langle \gamma \rangle \Diamond_C^* \psi$. But since $\beta \rightarrow_A \gamma$ and W contains $\text{PRE}(\beta)$, we also have a good path from $[W]$ for $\langle \beta \rangle \Diamond_C^* \psi$. This again contradicts $[W] \in S_\beta$. As a result, for all relevant A , β , and γ , $\vdash \chi_\beta \wedge \text{PRE}(\beta) \rightarrow \Box_A \chi_\gamma$.

By the Action Rule, $\vdash \chi_\alpha \rightarrow [\alpha] \Box_C^* \neg\psi$. Now $\vdash \llbracket V_0 \rrbracket \rightarrow \chi_\alpha$. So $\vdash \llbracket V_0 \rrbracket \rightarrow [\alpha] \Box_C^* \neg\psi$. This contradicts the assumption with which we began this proof. \dashv

Lemma 5.8 (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 $< \chi$. Note that by soundness, we may assume that χ is in normal form. 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 5.5, part (4). 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 $\llbracket U \rrbracket \wedge \Diamond_A \neg\psi$ is consistent. We use equation (6) and the fact that \Diamond_A distributes over disjunctions to see that $\llbracket U \rrbracket \wedge \Diamond_A \neg\psi$ is logically equivalent to $\bigvee (\llbracket U \rrbracket \wedge \Diamond_A \llbracket V \rrbracket)$, where the disjunction is taken over all V which contain $\neg\psi$. Since $\llbracket U \rrbracket \wedge \Diamond_A \neg\psi$ is consistent, one of the disjuncts $\llbracket U \rrbracket \wedge \Diamond_A \llbracket V \rrbracket$ must be consistent. The induction hypothesis again applies, and we use it to see that $\langle \mathcal{F}, [V] \rangle \models \neg\psi$. By the claim in the proof of Lemma 5.7, $[U] \xrightarrow{A} [V]$. We conclude that $\langle \mathcal{F}, [U] \rangle \models \Diamond_A \neg\psi$, and this is a contradiction.

For χ of the form $\Box_C^* \psi$, we use the standard argument for PDL (see Kozen and Parikh [10]). This is based on lemmas that parallel Lemmas 5.6 and 5.7. The work is somewhat easier than what we do below for sentences of the form $[\alpha] \Box_C^* \psi$, and so we omit these details.

We conclude with the case when χ is a normal form sentence of the form $[\alpha] \Box_C^* \psi \in \Delta$. Assume that $[\alpha] \Box_C^* \psi \in \Delta$. First, suppose that $[\alpha] \Box_C^* \psi \notin U$. Then by Lemma 5.7, there is a good path from $[U]$ for $\langle \alpha \rangle \Diamond_C^* \neg\psi$. We want to apply Lemma 4.1 in \mathcal{F} to assert that $\langle \mathcal{F}, [U] \rangle \models \langle \alpha \rangle \Diamond_C^* \neg\psi$. Let k be the length of the good path. For $i \leq k$, $\text{PRE}(\alpha_i) \in U_i$. Now each $\text{PRE}(\alpha_i)$ belongs to Δ by Lemma 5.5, part (5), and is $< [\alpha] \Box_C^* \psi$. So by induction hypothesis, $\langle \mathcal{F}, [U_i] \rangle \models \text{PRE}(\alpha_i)$. We also need to check that $\langle \mathcal{F}, [U_k] \rangle \models \langle \alpha_k \rangle \neg\psi$. For this, recall from Lemma 5.5, part (5) that Δ contains $nf([\alpha_k] \psi) \leq [\alpha_k] \psi$. By Lemma 4.4, $nf([\alpha_k] \psi) \leq [\alpha_k] \psi < [\alpha] \Box_C^* \psi$. Since the path is good, U_k contains $\langle \alpha_k \rangle \neg\psi$ and hence $\neg[\alpha_k] \psi$. It also must contain the normal form of this, by Lemma 5.4. So by induction hypothesis, $\langle \mathcal{F}, [U_k] \rangle \models nf(\neg[\alpha_k] \psi)$. By soundness, $\langle \mathcal{F}, [U_k] \rangle \models \langle \alpha_k \rangle \neg\psi$. Now it does follow from Lemma 4.1 that $\langle \mathcal{F}, [U] \rangle \models \langle \alpha \rangle \Diamond_C^* \neg\psi$.

Going the other way, suppose that $\langle \mathcal{F}, [U] \rangle \models \langle \alpha \rangle \diamond_c^* \neg \psi$. By Lemma 4.1, 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 $\langle \alpha \rangle \diamond_c^* \neg \psi$. By Lemma 5.6, U contains $\langle \alpha \rangle \diamond_c^* \neg \psi$. This completes the proof. \dashv

Theorem 5.9 (Completeness) *For all φ , $\vdash \varphi$ iff $\models \varphi$. Moreover, this relation is decidable.*

Proof By Lemma 5.4, $\vdash \varphi \leftrightarrow nf(\varphi)$. Let φ be consistent. By the Truth Lemma, $nf(\varphi)$ holds at some world in the filtration \mathcal{F} . So $nf(\varphi)$ 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 . \dashv

5.3 Two Extensions

We briefly mention two extensions of the Completeness Theorem 5.9. These extensions come from our discussion at the end of Section 2.3.

First, consider the case of S5 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 $\alpha \rightarrow_A \beta$, 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 \triangleleft is an equivalence relation. This is a standard exercise in modal logic (see, e.g., Fagin et al [4], Theorem 3.3.1).

Our second 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 $[\alpha]p \leftrightarrow (\text{PRE}(\alpha) \rightarrow p \wedge \neg q)$. 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.

6 Results on Expressive Power

In this section, we present two results which show that adding announcements to modal logic with \diamond^* adds expressive power as does adding private announcements to modal logic with \diamond^* and public announcements. To show these results it will be sufficient to take the set \mathcal{A} of agents to be $\{A, B\}$ and consider only languages contained in a language built-up from the atomic sentences p and q , using \diamond_A , \diamond_B , \diamond_A^* , \diamond_B^* , and \diamond_{AB}^* , and the actions $[\varphi]_A$, $[\varphi]_B$ of announcing φ to A or B privately, and $[\varphi]_{AB}$ the action of announcing φ to A and B publicly. Let \mathcal{L}_{all} stand for this language. We use here the customary notation ($[\varphi]_A$, $[\varphi]_B$, $[\varphi]_{AB}$) for announcements, but $[\varphi]_A$ is simply the action with the Kripke structure $K = \{k\}$ with \triangleleft from k to k and $\text{PRE}(k) = \varphi$. We think of $[\varphi]_B$ similarly. $[\varphi]_{AB}$ is the action with the Kripke structure $K = \{k\}$ with \triangleleft and \triangleright going from k to k and $\text{PRE}(k) = \varphi$.

We need to define a *rank* $|\varphi|$ on sentences from \mathcal{L}_{all} . Let $|p| = 0$ for p atomic, $|\neg\varphi| = |\varphi|$, $|\varphi \wedge \psi| = \max(|\varphi|, |\psi|)$, $|\neg\varphi| = |\varphi|$, $|\diamond_X \varphi| = 1 + |\varphi|$, for $X = A$ or $X = B$, $|\diamond_X^* \varphi| = 1 + |\varphi|$ for $X = A$, $X = B$, or $X = AB$, and $||[\varphi]_X \psi| = \max(|\varphi|, |\psi|)$ for $X = A$, $X = B$, or $X = AB$.

First 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 equivalent to any theory in \mathcal{L}_0 .

Lemma 6.1 *Let \mathcal{L}_0 be a language included in \mathcal{L}_{all} , and let ψ be a sentence in \mathcal{L}_{all} . 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. +

6.1 Announcements add Expressive Power to Modal Logic with \square^*

In the result below, there will be only one agent A , and so we omit the letter A from the notation. We let $\mathcal{L}([\], \diamond^*)$ be modal logic with announcements (to this A) and $\diamond^* = \diamond_A^*$. We also let $\mathcal{L}(\diamond^*)$ be the obvious sublanguage.

Theorem 6.2 *There is a sentence of $\mathcal{L}([\], \diamond^*)$ which cannot be expressed by any set of sentences of $\mathcal{L}(\diamond^*)$.*

Proof We show first that $[p]\diamond^+q = [p]\diamond\diamond^*q$ cannot be expressed by any single sentence of $\mathcal{L}(\square^*)$. (Incidentally, the same holds for $[p]\diamond^*q$.) Fix a natural number n . We define structures $\mathcal{A} = \mathcal{A}_n$ and $\mathcal{B} = \mathcal{B}_n$ as follows. First \mathcal{B} has $2n + 3$ points arranged cyclically as

$$0 \rightarrow 1 \rightarrow \dots \rightarrow n \rightarrow n+1 \rightarrow -n \rightarrow \dots \rightarrow -1 \rightarrow 0.$$

For the atomic sentences, we set p true at all points except $n+1$, and q true only at 0.

The structure \mathcal{A} is a copy of \mathcal{B} with n more points $\bar{1}, \dots, \bar{n}$ arranged as

$$0 \rightarrow \bar{1} \rightarrow \dots \rightarrow \bar{n} \rightarrow 0.$$

The shape of \mathcal{A} is a figure-8. In both structures, every point is reachable from every point by the transitive closure of the \rightarrow relation. At the points \bar{i} , p is true and q is false. Notice that $1 \models [p]\diamond^+q$ in \mathcal{A} , but $1 \not\models [p]\diamond^+q$ in \mathcal{B} .

The main technique in the proof is an adaptation of Fraisse-Ehrenfeucht games to the setting of modal logic. Here is a description of the relevant game $G_n(\langle U, u \rangle, \langle V, v \rangle)$. For $n = 0$, II immediately wins if the following holds: for all $p \in \text{AtSen}$, $\langle U, u \rangle \models p$ iff $\langle V, v \rangle \models p$. And if u and v differ on some atomic sentence, I immediately wins. Continuing, here is how we define $G_{n+1}(\langle U, u \rangle, \langle V, v \rangle)$. As in the case of G_0 , we first check if u and v 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. \diamond -move

I has a choice of playing from U or from V . If I chooses U , then I continues by choosing some u' such that $u \rightarrow u'$ in U . Then II replies with some $v' \in V$ such that $v \rightarrow v'$. Of course, if I had chosen in V , then II would have chosen in U . Either way, points u' and v' are determined, and the two players then play $G_n(\langle U, u' \rangle, \langle V, v' \rangle)$.

2. \diamond^* -move

I plays by selecting U (or V , but we ignore this symmetric case below), and then playing some u' (say) reachable from u in the reflexive-transitive closure \rightarrow^* of \rightarrow ; II responds with a point in the other model, V , which is similarly related to v .

We write $\langle U, u \rangle \sim_n \langle V, v \rangle$ if II has a winning strategy in the game $G_n(\langle U, u \rangle, \langle V, v \rangle)$. It is easy to check that by induction on m that if $\langle U, u \rangle \sim_n \langle V, v \rangle$ and $m < n$, then $\langle U, u \rangle \sim_m \langle V, v \rangle$.

Claim 1 If $\langle U, u \rangle \sim_n \langle V, v \rangle$, then for all φ with $|\varphi| \leq n$, $\langle U, u \rangle \models \varphi$ iff $\langle V, v \rangle \models \varphi$.

The proof will be done by induction on φ . Let φ be atomic. Suppose $\langle U, u \rangle \sim_n \langle V, v \rangle$. Then since II has a winning strategy, the atomic sentences satisfied by u and v must be the same. So we are done in this case.

The induction steps for the boolean connectives are trivial. For $\Box\varphi$, suppose that $|\Box\varphi| \leq n$, $\langle U, u \rangle \sim_n \langle V, v \rangle$, and $\langle U, u \rangle \models \Box\varphi$. Suppose towards a contradiction that $\langle V, v \rangle \models \Diamond\neg\varphi$. Let v' be such that $v \rightarrow v'$ in V and $\langle V, v' \rangle \models \neg\varphi$. Let I begin a play of $G_{n-1}(\langle U, u \rangle, \langle V, v \rangle)$ by choosing to play $v' \in V$. Then II 's winning strategy responds with some u' such that $\langle U, u' \rangle \sim_{n-1} \langle V, v' \rangle$. Since $|\varphi| \leq n-1$, our induction hypothesis implies that $\langle U, u' \rangle \models \Diamond\neg\varphi$. This is a contradiction.

The argument for $\Box^*\varphi$ is similar and we leave it to the reader. The claim is proved.

We return to the models \mathcal{A} and \mathcal{B} described in the beginning of this proof. For $0 \leq i \leq n$, we let $S_i \subseteq \mathcal{A} \times \mathcal{B}$ be the following set

$$\begin{aligned} S_i &= \{ (0, 0), \dots, (n, n), (n+1, n+1), (-n, -n), \dots, (-1, -1) \} \\ &\cup \{ (\overline{n}, -1), (\overline{n-1}, -2), \dots, (\overline{2}, -(n-1)), (\overline{1}, -n) \} \\ &\cup \{ (\overline{1}, 1), \dots, (\overline{n-i}, n-i) \} \end{aligned}$$

In the case of $i = n$, then the last disjunct is empty. Note that $S_0 \supset S_1 \supset \dots \supset S_n$. Also, for $0 \leq i \leq n$, every point of one structure is related by S_i to some point of the other.

Claim 2 If $0 \leq i \leq n$ and $(a, b) \in S_i$, then $\langle \mathcal{A}, a \rangle \sim_i \langle \mathcal{B}, b \rangle$.

The proof is by induction on i . If $i = 0$, this is due to the fact that pairs in S_0 agree on the atomic formulas. Assume the statement for i , and that $i+1 \leq n$. Let $(a, b) \in S_{i+1}$. We only need to show that II can respond to any play and have the resulting pair belong to S_i . Suppose first that I plays a \Diamond -move. Suppose also that $a = b$, so that (a, a) comes from the first subset of S_{i+1} . In this case, we only need to notice that $(a+1, a+1) \in S_i$ if $|a| \leq n$, $(-n, -n) \in S_i$ if $a = n+1$, and $(\overline{1}, 1) \in S_i$ if $a = 0$, since $i < n$. The case of (a, b) from the second subset is similar. Finally, if (\overline{a}, a) belongs to the third subset of S_{i+1} , then $a \leq n - (i+1) = n - i - 1$. So $a+1 \leq n-i$, and $(\overline{a+1}, a+1)$ belongs to the third subset of S_i . This tells II how to play.

We remarked above that each S_i is a total relation. Moreover, each world can be reached from any other one in \mathcal{A} and in \mathcal{B} . This implies that if I makes a \Diamond^* -move, II can respond. This completes the proof of the claim.

It follows that $\langle \mathcal{A}_n, 1 \rangle \sim_n \langle \mathcal{B}_n, 1 \rangle$. So by Claim 1, for each sentence $\varphi \in \mathcal{L}(\Box^*)$ and all $n \geq |\varphi|$, $\langle \mathcal{A}_n, 1 \rangle \models \varphi$ iff $\langle \mathcal{B}_n, 1 \rangle \models \varphi$. This shows that $[p]\Diamond^+q$ cannot be expressed by a single sentence in $\mathcal{L}(\Diamond^*)$. To prove the stronger result as stated in Theorem 6.2, we only need to quote Lemma 6.1. \dashv

6.2 Private Announcements Add Expressive Power

In this section, $\mathcal{L}([\]_{AB}, \Diamond^*)$ denotes the set of sentences built from p using $[\varphi]_{AB}$, \Diamond_A , \Diamond_B , \Diamond_A^* , \Diamond_B^* , and \Diamond_{AB}^* . $\mathcal{L}([\]_A, \Diamond_A^*)$ denotes the set built from p using $[\varphi]_A$, \Diamond_A^* , and \Diamond_B .

Theorem 6.3 *There is a sentence of $\mathcal{L}([\]_A, \Diamond_A^*)$ which cannot be expressed by any set of sentences in $\mathcal{L}([\]_{AB}, \Diamond^*)$.*

Proof We consider $\chi \equiv [p]_A \Diamond_A^* \Diamond_B \neg p$.

Let G_n be the following model. We begin with a cycle in \xrightarrow{A} :

$$a_1 \xrightarrow{A} a_\infty \xrightarrow{A} b \xrightarrow{A} a_n \xrightarrow{A} a_{n-1} \xrightarrow{A} \cdots \xrightarrow{A} a_2 \xrightarrow{A} a_1 \quad (8)$$

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 $x = a_\infty$). 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 after a *private* update of p to A , $\langle G_n, a_i \rangle \models \chi$ 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, $\langle G_n, a_\infty \rangle \models \neg\chi$. This is because the only way to go from a_∞ to b is to go through 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}([\]_{AB}, \diamond^*)$. 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}([\]_{AB}, \diamond^*)$ and $|\varphi| < j$. Then $\langle G_n, a_j \rangle \models \varphi$ iff $\langle G_n, a_\infty \rangle \models \varphi$.

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 (8) 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 $[\varphi]_{AB}\psi$, assume the result for φ and ψ , and let $|[\varphi]_{AB}\psi| < j$. Then also $|\varphi| < j$ and $|\psi| < j$. Let $H = \{x : x \models \varphi\}$ be the updated model, and recall that $\langle G_n, x \rangle \models [\varphi]_{AB}\psi$ iff $x \in H$ and $\langle H, x \rangle \models \psi$. We have two cases: First, $H = G_n$. Then $\langle G_n, x \rangle \models [\varphi]_{AB}\psi$ 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 $[\varphi]_{AB}\psi$. 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 $(A_{\geq k} \times A_{\geq k}) \cup \{(b, b)\}$ is a bisimulation on H . The verification here is easy.

So in H , a_j and a_∞ agree on all sentences in *any* language which is invariant for bisimulation. Now $\mathcal{L}([\]_{AB}, \diamond^*)$ has this property (as do all the languages which we study: they are translatable into infinitary modal logic). In particular, $\langle H, a_j \rangle \models \psi$ iff $\langle H, a_\infty \rangle \models \psi$. This concludes the claim.

We get Theorem 6.3 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 6.1. \dashv

We feel that our two 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?

7 Conclusions and Historical Remarks

The work of this paper builds on the long tradition of epistemic logic as well as technical results in other areas. 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** [4] by Fagin, Halpern, Moses, and Vardi. We depart from [4] by introducing the new operators for epistemic actions, 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 [8] examines 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.

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" [13]. As the title suggests, the epistemic actions studied are announcements to the whole group, as in our α and α' . Perhaps the main result of the paper is a completeness theorem for the logic of public announcements and knowledge. This result is closely related to a special case of our Theorem 3.5. The difference is that Plaza restricts attention to the case when all of the accessibility relations are equivalence relations. Incidentally, Plaza's proof involves a translation to multi-modal logic, just as ours does. In addition to this, [13] 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.

Other predecessors to this paper are the papers of Gerbrandy [5, 6] and Gerbrandy and Groeneveld [7]. These study epistemic actions similar to our β , where an announcement is made to set of agents in a private way with no suspicions. They presented a logical system which included the common knowledge operators. An important result is that all of the reasoning in the original Muddy Children scenario can be carried out in their system. This shows that in order to get a formal treatment of the problem, one need not posit models which maintain histories. They did not obtain the completeness/decidability result for their system, but it would be the version of Theorem 5.9 restricted to actions which are compositions of private announcements. So it follows from our work that all of the reasoning in the Muddy Children can be carried out in a decidable system.

We should mention that the systems studied in [5, 6, 7] differ from ours in that they are variants of dynamic logic rather than propositional logic. That is, announcements are particular types of *programs* as opposed to modalities. This is a natural move, and although we have not followed it in this paper, we have carried out a study of expressive power issues of various fragments of a dynamic logic with announcement operators. We have shown, for example, that

the dynamic logic formulations are more expressive than the purely propositional ones. Details on this will appear in a forthcoming paper.

Incidentally, the semantics in [5, 6, 7] use non-wellfounded sets. In other words, they work with models modulo bisimulation. The advantages of moving from these to arbitrary Kripke models are that the logic can be used by those who do not know about non-wellfounded sets, and also that completeness results are slightly stronger with a more general semantics. The relevant equivalence of the two semantics is the subject of the short note [11].

The following are the new contributions of this paper:

1. We formulated a logical system with modalities corresponding to intuitive group-level epistemic actions. These actions include natural formalizations of announcements such as γ and δ , which allow various types of suspicion by outsiders. Our apparatus also permits us to study epistemic actions which apparently have not yet been considered in this line of work, such as actions in which nothing actually happens but one agent suspects that a secret communication took place.
2. We formulated a logical system with these modalities and with common knowledge operators for all groups. Building on the completeness of PDL and using a bit of term rewriting theory, we axiomatized the validities in our system.
3. We obtained some results on expressive power: in the presence of common knowledge operators, it is not possible to translate away public announcements, and in our framework, private announcements add expressive power to public ones.

8 Appendix: the lexicographic path order

In this appendix, we give the details on the *lexicographic path ordering* (LPO), both in general and in connection with $\mathcal{L}([\alpha])$ and $\mathcal{L}([\alpha], \Box^*)$.

Fix some many-sorted signature Σ of terms. In order to define the LPO $<$ on the Σ -terms, we must first specify a well-order $<$ on the set of function symbols of Σ . The LPO determined by such choices is the smallest relation $<$ such that:

(LPO1) If $(t_1, \dots, t_n) < (s_1, \dots, s_n)$ in the lexicographic ordering on n -tuples, and if $t_j < f(s_1, \dots, s_n)$ for $1 \leq j \leq n$, then $f(t_1, \dots, t_n) < f(s_1, \dots, s_n)$.

(LPO2) If $t \leq s_i$ for some i , then $t < f(s_1, \dots, s_n)$.

(LPO3) If $g < f$ and $t_i < f(s_1, \dots, s_n)$ for all $i \leq m$, then $g(t_1, \dots, t_m) < f(s_1, \dots, s_n)$.

Here is how this is applied in this paper. We shall take two sorts: *sentences* and *actions*. Our signature contains the usual sentence-forming operators p (for $p \in \text{AtSen}$) \neg , \wedge , and \Box_A for all $A \in \mathcal{A}$. Here each p is 0-ary, \neg and \Box_A are unary, and \wedge is binary. We also have an operator *app* taking actions and sentences to sentences. We think of $app(\psi, \alpha)$ as merely a variation on $[\alpha]\psi$. (The order of arguments to *app* is significant.) We further have a binary operator \circ on actions. (This is a departure from the treatment of this paper, since we used \circ as a metalinguistic abbreviation instead of as a formal symbol. It will be convenient to make this change because this leads to a smoother treatment of the Composition Axiom.) Finally,

for each finite Kripke frame K over $\mathcal{L}([\alpha])$ and each $1 \leq i \leq |K|$, we have a symbol F_K^i taking $|K|$ sentences and returning an action.

Each sentence φ has a formal version $\overline{\varphi}$ in this signature, and each action α also has a formal version $\overline{\alpha}$. These are defined by the recursion which is obvious except for the clauses

$$\begin{aligned} \overline{[\alpha]\varphi} &= \text{app}(\overline{\varphi}, \overline{\alpha}) \\ \overline{\alpha} &= F_K^i(\overline{\text{PRE}(k_1)}, \dots, \overline{\text{PRE}(k_n)}) \end{aligned}$$

Here $\alpha = \langle K, k_i, \text{PRE} \rangle$ with $K = \{k_1, \dots, k_n\}$ in some specified order. However, outside of the proof of Proposition 8.2 we shall not explicitly mention the formal versions at all, since they are harder to read than the standard notation.

We must also first fix a wellfounded relation $<$ on the function symbols. We set app to be greater than all other function symbols. In all other cases, distinct function symbols are unordered.

Theorem 8.1 (Kamin and Levy [9]; Dershowitz [3]) *Let $<$ be an LPO on Σ -terms.*

1. $<$ is transitive.
2. $<$ has the subterm property: if t is a proper subterm of u , then $t < u$.
3. $<$ is monotonic (it has the replacement property): if $y < x_i$ for some i , then

$$f(x_1, \dots, y, \dots, x_n) < f(x_1, \dots, x_i, \dots, x_n).$$

4. $<$ is wellfounded.
5. Consider a term rewriting system every rule of which of the form $l \rightsquigarrow r$ with $r < l$. Then the system is terminating: there are no infinite sequences of rewritings.

Proof Here is a sketch for part (1): We check by induction on the construction of the least relation $<$ that if $s < t$, then for all u such that $t < u$, $s < u$. For this, we use induction on the term u . We omit the details. Further, (2) follows easily from (1) and (LPO2), and (3) from (LPO1), (1) and (2). Moreover, (5) follows easily from (4) and (3), since the latter implies that any replacement according to the rewrite system results in a smaller term in the order $<$.

Here is a proof of the wellfoundedness property (4), taken from Buchholz [2]. (We generalized it slightly from the one-sorted to the many-sorted setting and from the assumption that $<$ is a finite linear order on Σ to the assumption that $<$ is any wellfounded order.)

Let W be the set of terms t such that the order $<$ is wellfounded below t . W is then itself wellfounded under $<$. So for all n , W^n is wellfounded under the induced lexicographic order. We prove by induction on the given wellfounded relation on function symbols of Σ that for all n -ary f , $f[W^n] \subseteq W$. So assume that for $g < f$, say with arity m , $g[W^m] \subseteq W$. We check this for f by using induction on W^n . Fix $\vec{s} \in W^n$, and assume that whenever $\vec{u} < \vec{s}$ in W^n , that $f(\vec{u}) \in W$. We prove that $f(\vec{s}) \in W$ by checking that for all t such that $t < f(\vec{s})$, $t \in W$. And *this* is done by induction on the structure of t . If $t = f(\vec{u}) < f(\vec{s})$ via (LPO1), then $\vec{u} < \vec{s}$ lexicographically, and each $u_i < s_i$. This last point implies that $\vec{u} \in W^n$ by induction

hypothesis on t , so $t \in W$ by induction hypothesis on W^n . If $t \leq s_i$ so that $t < f(\vec{s})$ via (LPO2), then $t \in W$ by definition of W . And if $t = g(u_1, \dots, u_m) < f(\vec{s})$ via (LPO3), then $g < f$ and each $u_i < f(\vec{s})$. By induction hypothesis on t , each $u_i \in W$. So by induction hypothesis on f , $g(\vec{u}) \in W$.

Now that we know that each f takes tuples in W^n to elements of W , it follows by induction on terms that all terms belong to W . \dashv

For more on the LPO, its generalizations and extensions, see the surveys Dershowitz [3] and Plaisted [12].

Proposition 8.2 *Consider the LPO $<$ on $\mathcal{L}([\alpha], \square^*)$ defined above.*

1. If $\alpha \dot{\rightarrow} \beta$, then $\text{PRE}(\beta) < \alpha$.
2. If $\alpha \dot{\rightarrow} \beta$, then $[\beta]\psi < [\alpha]\square_{\mathcal{C}}^*\psi$.
3. $\text{PRE}(\alpha) \rightarrow p < [\alpha]p$.
4. $\text{PRE}(\alpha) \rightarrow \neg[\alpha]\psi < [\alpha]\neg\psi$.
5. $[\alpha]\psi \wedge [\alpha]\chi < [\alpha](\psi \wedge \chi)$.
6. $\text{PRE}(\alpha) \rightarrow \bigwedge\{\square_A[\beta]\psi : \alpha \dot{\rightarrow} \beta\} < [\alpha]\square_A\psi$.
7. $[\alpha \circ \beta]\varphi < [\alpha][\beta]\psi$.

In particular, for all rules $\varphi \rightsquigarrow \psi$ of the rewriting system \mathcal{R}^* , $\psi < \varphi$.

Proof Part (1) holds because we regard α as a term $\alpha = F_K^i(\overline{\gamma}_1, \dots, \overline{\gamma}_n)$, for some frame K and i . So whenever $\alpha \dot{\rightarrow} \beta$, each $\text{PRE}(\beta)$ is a proper subterm of α .

Here is the argument for part (2): We need to see that $\text{app}(\psi, \overline{\beta}) < \text{app}(\square_{\mathcal{C}}^*\psi, \overline{\alpha})$. Now lexicographically, $(\psi, \overline{\beta}) < (\square_{\mathcal{C}}^*\psi, \overline{\alpha})$. So we only need to know that $\overline{\beta} < \text{app}(\square_{\mathcal{C}}^*\psi, \overline{\alpha})$. Let $\alpha = F_K^i(\overline{\gamma}_1, \dots, \overline{\gamma}_n)$. Now according to equation (2) in Section 2.1, $\overline{\beta}$ is $F_K^j(\overline{\gamma}_1, \dots, \overline{\gamma}_n)$, for the same K and $\gamma_1, \dots, \gamma_n$ but perhaps for $j \neq i$. Then it is clear by (LPO2) that $\overline{\gamma}_i < \text{app}(\square_{\mathcal{C}}^*\psi, \overline{\alpha})$ for all i . So by (LPO3), $\overline{\beta} < \text{app}(\square_{\mathcal{C}}^*\psi, \overline{\alpha})$. \dashv

The remaining parts are similar. \dashv

A *normal form* in a rewriting system is a sentence which cannot be rewritten in the system. Of course, we are interested in the systems \mathcal{R} and \mathcal{R}^* from Sections 3 and 4, respectively. It follows from the wellfoundedness of $<$ that for every φ there is a normal form $nf(\varphi) \leq \varphi$ obtained by rewriting φ in some arbitrary fashion until a normal form is reached.

Lemma 8.3 *A sentence $\varphi \in \mathcal{L}([\alpha])$ is a normal form of \mathcal{R}^* iff φ is a modal sentence (i.e., iff φ contains no actions). Moreover, the rule $[\alpha][\beta]\varphi \rightsquigarrow [\alpha \circ \beta]\varphi$ is not needed to reduce φ to normal form. So for $\mathcal{L}([\alpha])$, \mathcal{R} has the same normal forms as \mathcal{R}^* .*

A sentence $\varphi \in \mathcal{L}([\alpha], \square^)$ is a normal form of \mathcal{R}^* iff φ is built from atomic sentences using \neg , \wedge , \square_A , and $\square_{\mathcal{B}}^*$, or if φ is of the form $[\alpha]\square_{\mathcal{B}}^*\psi$, where α is a normal form action, and ψ is a normal form. An action α is a normal form if whenever $\alpha \dot{\rightarrow} \beta$, then $\text{PRE}(\beta)$ is a normal form.*

Proof It is immediate that every modal sentence is a normal form in $\mathcal{L}([\alpha])$, that every $[\alpha]\square_{\mathcal{C}}^*\varphi$

is a normal form in $\mathcal{L}([\alpha], \Box^*)$, and that if each $\text{PRE}(\beta)$, with $\alpha \xrightarrow{*} \beta$, is a normal form, then α is a normal form action. Going the other way, we check that if $\varphi \in \mathcal{L}([\alpha])$, $[\alpha]\varphi$ is *not* a normal form. So we see by an easy induction that the normal forms of $\mathcal{L}([\alpha])$ are exactly the modal sentences. We also argue by induction for $\mathcal{L}([\alpha], \Box^*)$, and we note that every $[\alpha][\beta]\varphi$ is *not* a normal form, using the rule $[\alpha][\beta]\varphi \rightsquigarrow [\alpha \circ \beta]\varphi$. \dashv

One fine point concerning \mathcal{R} and our work in Section 3 is that to reduce sentences of $\mathcal{L}([\alpha])$ to normal form we may restrict ourselves to rewriting sentences which are not subterms of actions. This simplification accounts for the differences between parallel results of Sections 3 and 4.

Acknowledgements

We thank Jelle Gerbrandy and Rohit Parikh for useful conversations on this work. An earlier version of this paper was presented at the 1998 Conference on Theoretical Aspects of Rationality and Knowledge.

References

- [1] Alexandru Baltag, A logic of epistemic actions, ms. CWI, Amsterdam, 1999.
- [2] Wilfried Buchholz, Proof-theoretic analysis of termination proofs, *Annals of Pure and Applied Logic* **75** (1995) 57–65.
- [3] Nachum Dershowitz, Orderings for term-rewriting systems, *Theoretical Computer Science* **17** (1982) 279–301.
- [4] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi, **Reasoning About Knowledge**, MIT Press, 1995.
- [5] Jelle Gerbrandy, Dynamic epistemic logic, in Lawrence S. Moss, et al (eds.) **Logic, Language, and Information, vol. 2**, CSLI Publications, Stanford University 1999.
- [6] Jelle Gerbrandy, *Bisimulations on Planet Kripke*, Ph.D. dissertation, University of Amsterdam, 1999.
- [7] Jelle Gerbrandy and Willem Groeneveld, Reasoning about information change, *Journal of Logic, Language, and Information* **6** (1997) 147–169.
- [8] 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, 195–237.
- [9] S. Kamin and J. J. Levy, Two generalizations of the recursive path orderings, Unpublished note, Department of Computer Science, University of Illinois, Urbana IL (1980).
- [10] D. Kozen and R. Parikh, An elementary proof of the completeness of PDL, *Theoretical Computer Science* (1981) 113–118.

- [11] 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 50th Birthday**, Vossiuspers, Amsterdam University Press, 1999.
- [12] David A. Plaisted, Termination orderings, in D. Gabbay (et al) eds., **Handbook of Logic in Artificial Intelligence and Logic Programming**, vol. I, 273–364.
- [13] Jan Plaza, Logics of public communications, *Proceedings, 4th International Symposium on Methodologies for Intelligent Systems*, 1989.
- [14] Frank Veltman, Defaults in update semantics, *Journal of Philosophical Logic* **25** (1996) 221–261.