

Endogenizing Epistemic Actions

Will Nalls

Carnegie Mellon University
Pittsburgh, USA
will.nalls@gmail.com

Adam Bjorndahl

Carnegie Mellon University
Pittsburgh, USA
abjorn@andrew.cmu.edu

Through a series of examples, we illustrate some important drawbacks that the *action logic* framework suffers from in its ability to represent the dynamics of information updates. We argue that these problems stem from the fact that the *action model*, a central construct designed to encode agents' uncertainty about actions, is itself effectively common knowledge amongst the agents. In response to these difficulties, we motivate and propose an alternative semantics that avoids them by (roughly speaking) endogenizing the action model. We discuss the relationship to action logic, and provide a sound and complete axiomatization.

1 Introduction

Action Logic (AL) is a framework for reasoning about how knowledge and belief changes on the basis of incoming information [2, 1, 4].¹ Information is conveyed in the form of “epistemic actions”, with a canonical example being *public announcements* [7]. Unlike *Public Announcement Logic* (PAL), however, AL does not presume that all epistemic actions are public in the sense of becoming common knowledge, nor that such actions can always be distinguished from one another by the agents. Uncertainty about epistemic actions is explicitly encoded in a structure called an *action model*; this allows for the representation of scenarios in which agents can be uncertain about which action has taken place.

By design, this formalism is well-equipped to capture uncertainty about actions themselves; however, we argue in this paper that the AL framework is ill-suited to the representation of *higher-order* uncertainty about actions. Roughly speaking, this is because the action model that captures uncertainty about actions is itself effectively common knowledge amongst the agents, making it awkward to encode, for example, one agent's uncertainty about another agent's uncertainty about actions.

We expose this difficulty through a series of motivating examples. We demonstrate that AL can capture higher-order uncertainty about actions only by expanding the action model in such a way as to essentially “pre-encode” the desired uncertainty; this makes choosing an appropriate action model for any given application problematically post hoc. Furthermore, we show that in such cases small variations in the background epistemic conditions require corresponding alterations to the action models in order to ensure that the “pre-encoded” uncertainty maintains the right form. These observations seriously undermine the practical applicability of AL as a tool for reasoning about information updates.

In response to these challenges, we reformulate the semantics by “endogenizing” the action model; somewhat more precisely, we allow each agent's uncertainty about actions to be world-dependent, and therefore itself subject to uncertainty. Revisiting our examples, we show that these revised semantics completely circumvent the earlier difficulties; our semantics capture *formally* the informal process underlying the aforementioned post hoc expansion of the action model.

¹Terminology varies; some authors instead use *Dynamic Epistemic Logic* to refer to this framework [6, §2.2.2]. We follow van Ditmarsch et al. [4] in using it instead as an umbrella term for a collection of thematically related logics of information change, including Action Logic.

The idea of representing higher-order uncertainty about epistemic actions by encoding extra information about the agents (and their perceptions of such actions) into the state space is a natural one; it occurs also in [3], in which Bolander et al. study a more general class of announcements that may not be entirely public in that (loosely speaking) some agents may not be “listening”. In essence, their semantics work by encoding into each possible world whether or not each agent is “paying attention” at that world. As the authors show, such “attention-based” announcements can also be described using action models. The present work is more general in that we begin with the full action model framework rather than PAL, and we encode into each world the full spectrum of each agent’s uncertainty regarding epistemic actions, not just whether or not they are attentive.

The rest of the paper is organized as follows. In Section 2, we motivate and define the basic AL framework, and present a series of examples intended to illustrate its limitations. In Section 3, we present our new semantics and show how it deals with the problems discussed in the previous section. Section 4 presents an axiomatization.

2 Action Logic

We begin by reviewing the foundational definitions and motivations of AL, largely following van Ditmarsch et al. [4, Chapter 6]. This logic is an extension of standard epistemic logic, so we begin there.

Fix a countable set of primitive propositions PROP and a finite set of agents G . Let \mathcal{L}_{EL} denote the language recursively defined as follows:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_j\varphi,$$

where $p \in \text{PROP}$ and $j \in G$. We read $K_j\varphi$ as “agent j knows that φ ”. Thus, \mathcal{L}_{EL} is a language for reasoning about the knowledge of the agents in G . The other Boolean connectives can be defined in the usual way; we write \hat{K}_j to abbreviate $\neg K_j \neg$, and read $\hat{K}_j\varphi$ as “agent j considers it possible that φ ”.

An **epistemic model** is a structure of the form $M = \langle W, \{\sim_j : j \in G\}, V \rangle$, where:

- W is a (nonempty) set of *states*,
- \sim_j is an equivalence relation on W ,
- $V : \text{PROP} \rightarrow 2^W$ is a *valuation function*.

Intuitively, V specifies for each primitive proposition those states where it is true, while the relations \sim_j capture indistinguishability from the perspective of agent j . These intuitions are formalized in the following semantic clauses:

$$\begin{aligned} (M, w) \models p & \quad \text{iff} \quad w \in V(p) \\ (M, w) \models \neg\varphi & \quad \text{iff} \quad (M, w) \not\models \varphi \\ (M, w) \models \varphi \wedge \psi & \quad \text{iff} \quad (M, w) \models \varphi \text{ and } (M, w) \models \psi \\ (M, w) \models K_j\varphi & \quad \text{iff} \quad (\forall w' \in [w]^j)((M, w') \models \varphi), \end{aligned}$$

where $[w]^j$ denotes the equivalence class of w under \sim_j . Thus, the Boolean connectives are interpreted as usual, and $K_j\varphi$ is true at w precisely when φ is true at all worlds that agent j cannot distinguish from w . Insisting that the indistinguishability relations be equivalence relations results in a logic of knowledge that is *factive* and *fully introspective*. For a more thorough development of differing logics of knowledge we direct the reader to [5].

On top of this basic epistemic framework, Action Logic adds a layer of structure that aims to capture the epistemic dynamics of information update. An **action model** is a structure of the form

$$A = \langle \Sigma, \{\approx_j : j \in G\}, Pre \rangle$$

where:

- Σ is a (nonempty) set of *epistemic actions*,
- \approx_j is an equivalence relation on Σ ,
- $Pre : \Sigma \rightarrow \mathcal{L}_{EL}$ is a *precondition function*.

Intuitively, the relation \approx_j captures indistinguishability of actions from the perspective of agent j , while the function Pre captures the background conditions $Pre(\sigma)$ that must hold for a given action σ to be successfully performed. In short, an action model specifies a set of epistemic actions that can be executed together with their preconditions and the extent to which they can be individuated by the agents.

To formalize these intuitions, we must define the process by which an epistemic model M is updated based on the performance of an epistemic action from A . This is captured in the **updated model**

$$M^A = \langle W_\Sigma, \{\sim'_j : j \in G\}, V' \rangle$$

defined as follows:

- $W_\Sigma = \{(w, \sigma) : (M, w) \models Pre(\sigma)\}$,
- $(w_0, \sigma_0) \sim'_j (w_1, \sigma_1)$ iff $w_0 \sim_j w_1$ and $\sigma_0 \approx_j \sigma_1$,
- $(w, \sigma) \in V'(p)$ iff $w \in V(p)$.

Thus, the states of the updated model consist of those state-action pairs (w, σ) such that the precondition of the action σ is satisfied by the state w in M ; intuitively, (w, σ) represents the state of the world w after σ has been performed. The definition of V' ensures that such “updated states” satisfy the same primitive propositions as they did before (corresponding to the intuition that epistemic actions can change the *information* agents have access to, but cannot change basic facts about the world). Finally, the definition of \sim'_j specifies that updated state-action pairs are indistinguishable for agent j precisely when the constituent states and actions were indistinguishable for j in M and A , respectively. It is easy to see that M^A is an epistemic model.

The language for Action Logic, \mathcal{L}_{AL} , extends the basic epistemic language with an update operator:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_j\varphi \mid [A, \sigma]\varphi.$$

$[A, \sigma]\varphi$ is read “if action σ can be performed, then afterwards, φ is true”.² This language therefore lets us reason about agents’ knowledge and how it can change as a result of epistemic actions. Formulas of \mathcal{L}_{AL} can be interpreted in epistemic models as before, with the additional semantic clause for $[A, \sigma]$ given by:

$$(M, w) \models [A, \sigma]\varphi \quad \text{iff} \quad (M, w) \models Pre(\sigma) \text{ implies } (M^A, (w, \sigma)) \models \varphi.$$

To make these definitions clear, we present a simple example that will recur in various forms throughout the paper.




Figure 1: M_0 – Anne and Bob’s initial uncertainty

Example 1. Colleagues Anne and Bob are discussing whether a particular company policy passed in this morning’s board meeting: both are ignorant of whether or not the policy passed. Taking p to be the proposition “the policy passed”, this scenario may be represented with the simple epistemic model M_0 presented in Figure 1.³

In this diagram and others like it, the circles represents the states of the model (in this case, w_0 and w_1), the formulas listed beside a state are true at that state (in this case, p and $\neg p$ are true at w_0 and w_1 , respectively), and edges between states are labelled with those agents that cannot distinguish those states (in this case, the edge labelled with a and b indicates that neither Anne nor Bob can distinguish between these two possible worlds). Indeed, we have $M_0 \models (\neg K_a p \wedge \neg K_a \neg p) \wedge (\neg K_b p \wedge \neg K_b \neg p)$, so this model properly captures the ignorance of Anne and Bob regarding p .

Along comes Carl, fresh from the board meeting with news of whether the policy passed. He takes one of two actions: he either tells Bob that p is true (σ_p), or he tells Bob that p is false ($\sigma_{\neg p}$). Anne watches intently: she knows that Carl is telling Bob what happened, but she is too far away to hear what Carl actually reports. We presume that Carl is honest, so the precondition of σ_p is p , and the precondition of $\sigma_{\neg p}$ is $\neg p$. This is all captured in the action model A_0 depicted in Figure 2, where the circles represent actions, the formulas listed beside the actions are the corresponding preconditions, and the edges represent action indistinguishability.





Figure 2: A_0 – Carl’s communication to Bob

In particular, Anne cannot distinguish σ_p from $\sigma_{\neg p}$ (though she knows that *one* of them happened), while Bob can. We therefore arrive at the following updated model $M_0^{A_0}$ given in Figure 3. Thus we see that when we update the epistemic model with the action model, the result is a model where Bob knows whether or not p , while Anne does not: $M_0^{A_0} \models K_b p \vee K_b \neg p$, and $M_0^{A_0} \models \neg(K_a p \vee K_a \neg p)$, as desired. \square

Note that Carl’s announcement is not *public* since Anne cannot tell whether he announced p or $\neg p$

²The “if-then” construction here is interpreted as a standard material conditional, i.e., $[A, \sigma]\phi$ is vacuously true when σ cannot be performed.

³Unless stated otherwise, all relations are equivalence relations, so reflexive loops and edges implied by transitivity are assumed to be present, even when suppressed in the diagrams.

Figure 3: $M_0^{A_0}$ – After Carl’s announcement

while Bob can.⁴ This demonstrates that AL can capture epistemic dynamics that PAL cannot.⁵

Nonetheless, as we have claimed, AL suffers from limitations of its own; we turn now to a discussion of these limitations.

2.1 Limitations of action logic

The limitations we have in mind can be summarized quite simply: the AL framework treats the action model as common knowledge. Indeed, the updated model M^A imports much of the structure of the action model: in it, all agents come to know what actions the *other* agents can distinguish. The result is that AL has difficulty capturing scenarios involving *higher-order uncertainty*—e.g., uncertainty about what other agents know. Although this limitation can be overcome, in a sense, by expanding the action model, we will see that in general this is not an appealing solution. Moreover, the examples we consider demonstrate that the AL framework is not as modular as it might appear to be: adjustments to the epistemic model will often require corresponding adjustments to the action model to preserve the intended semantic interpretations. To illustrate these points, we return to the scenario presented in Example 1.

Example 2. Suppose again that Carl comes along with news of p ; however, instead of speaking so that only Bob can hear him, Carl speaks plainly for all to hear, but he delivers the message in French. As it happens, Bob speaks French and Anne does not.

As before, the apparent actions that Carl might take are telling Bob that p and telling Bob that $\neg p$, which presumably have preconditions p and $\neg p$, respectively. Bob can distinguish these actions, as he speaks French, while Anne cannot. This reasoning produces the same action model A_0 depicted in Figure 2, which therefore produces the same updated model $M_0^{A_0}$ shown in Figure 3. As expected, then, just as in Example 1, Bob ends up knowing whether or not p , and Anne does not. An unexpected result, however, is that Anne *knows this former fact*, and Bob knows the latter! That is, we have:

$$M_0^{A_0} \models K_a(K_b p \vee K_b \neg p)$$

and

$$M_0^{A_0} \models K_b(\neg(K_a p \vee K_a \neg p)).$$

The former says that Anne knows that Bob knows whether p , while the latter says that Bob knows that Anne is uncertain about p . But there was no assumption that Bob knows that Anne cannot understand Carl’s message, nor that Anne knows that Bob can. That is, we did not explicitly stipulate whether or not either knew about the other’s (in)ability to speak French. To capture this, the model must be refined. \square

⁴This particular sort of non-public announcement is sometimes referred to as “semi-private” in the literature.

⁵In fact, AL subsumes PAL: to capture a public announcement of φ , consider the action model A_φ consisting of a single node σ with the precondition φ . Then, given an epistemic model M and a state w therein, one can show that the resulting models $(M^{A_\varphi}, (w, \sigma))$ and $(M|_\varphi, w)$ are bisimilar: updating with A_φ effectively deletes the states in M where φ is false.

Loosely speaking, Bob’s ability to speak French and Anne’s inability to speak French are represented in the structure of the action model A_0 in Figure 2, and this is why they effectively become common knowledge in the updated model. But, of course, we may want to capture a scenario where one or both are *uncertain* about whether or not the other speaks French; indeed, uncertainties of this sort play an important role in everyday reasoning.

For the sake of simplicity, let us begin by aiming only to remove the consequence in the updated model that Anne knows that Bob knows whether p .

Example 3. A relevant proposition here is that Bob speaks French; call this q . If we wish to account for Anne’s uncertainty about q , we ought to expand the initial epistemic model M_0 to include the possible values this proposition might have. The result is the model M_1 depicted in Figure 4.




Figure 4: M_1 – An expanded model of Anne and Bob’s initial uncertainty

Anne does not know whether or not p is true, and also does not know whether or not q is true; thus, there are a -edges between all 4 nodes. We assume that Bob does know whether or not he speaks French, but, as before, does not know whether or not p ; thus there are horizontal b -edges, but no vertical b -edges.

It is easy to check that updating M_1 with A_0 produces an epistemic model $M_1^{A_0}$ in which Bob knows (at every state) the true value of p . So it seems we must modify the action model as well; we add a third action, σ , corresponding intuitively to the “unsuccessful” announcement in which Carl speaks his piece but no one (including Bob) understands him. The precondition for σ should therefore be $\neg q$: that Bob does not speak French. Furthermore, the preconditions for the actions σ_p and $\sigma_{\neg p}$ ought to be strengthened to include q , since these actions now represent Carl telling Bob p or $\neg p$ and *Bob understanding what was said*. Bob will be able to distinguish any of these three actions, since he knows whether or not he speaks French, and, given that he speaks French, he knows which announcement Carl is making. Anne, on the hand, will not be able to distinguish any of the three actions—it all sounds the same to her. All this is captured by the action model A_1 given in Figure 5.





Figure 5: A_1 – An expanded action model

The updated model $M_1^{A_1}$ is shown in Figure 6.

Figure 6: $M_1^{A_1}$ – After Carl’s announcement

As expected, we now have that Bob knows whether p only in those nodes where σ_p or $\sigma_{\neg p}$ was performed, and Anne does not know whether p at all. Furthermore, we have:

$$(w_0, \sigma_p) \models \hat{K}_a(K_b p \vee K_b \neg p) \wedge \hat{K}_a(\neg K_b p \wedge \neg K_b \neg p)$$

This reads: Anne considers it possible that Bob knows whether or not p , and also considers it possible that he does not. Thus, by expanding the initial epistemic model to include Anne’s uncertainty about Bob’s ability to speak French, as well as adding a third node to the action model corresponding (roughly speaking) to an “unsuccessful” announcement from Carl, we are able to capture the second-order uncertainty that we set out to capture. \square

One might reasonably feel some discomfort regarding the introduction of σ into the action model. On at least one intuition for what constitutes an “action”, Carl’s announcement (in French) of p ought to count as the *same* action regardless of who hears it or what languages they might understand. In other words, one might object to distinguishing σ_p from σ on the grounds that it builds into the ontology of actions properties that really have nothing to do with actions, but rather with agents.

This philosophical objection could perhaps be swept aside if the underlying formalism actually did the job we wanted it to: it is a hard case to make to let vague ontological concerns trump mathematical efficacy. What we now aim to demonstrate, however, is that this technique of expanding the action model is *not* an effective tool for capturing higher-order uncertainty.

Example 4. We alter the scenario of Examples 2 and 3 in only one respect: we assume now that Anne *does* speak French. This requires no change to the initial epistemic model M_1 (since whether or not Anne speaks French is not represented explicitly in this model), but it does, intuitively, require us to re-work the action model A_1 . In particular, the Anne-edge connecting σ_p and $\sigma_{\neg p}$ no longer seems appropriate, since now Anne *can* understand what Carl announces.

Now we are faced with a somewhat awkward question—should there be an Anne-edge between σ_p and σ ? Intuitively, there should be, since σ is supposed to encode the fact that Bob does not understand Carl’s announcement, and Anne is not supposed to be able to tell whether he does or not. Similar reasoning leads us to leave the Anne-edge between $\sigma_{\neg p}$ and σ in place, so the resulting relation fails to be transitive.

Perhaps we could relax the requirements placed on the relations \approx_i in action models to accommodate this type of problem, but in fact there is a deeper issue here that suggests an alternative resolution: it is easy to see that *any* reflexive relation \approx'_a for Anne on the set $\{\sigma_p, \sigma_{\neg p}, \sigma\}$ produces an action

model A'_1 such that the updated model $M_1^{A'_1}$ satisfies $(w_2, \sigma) \sim_a (w_3, \sigma)$. But this misrepresents the situation: in world w_2 , Carl’s announcement must have been that the policy passed, p ; as such, after his announcement, Anne should no longer be uncertain about p .

The problem here lies with σ : it was introduced originally to represent the possibility of an “unsuccessful” announcement by Carl. But in the present context, Carl’s announcement is always at least partially successful, in that it always informs Anne of the truth value of p . The natural fix to this problem is another adjustment to the action model: we “split” the action σ into two actions, $\sigma_{\neg p \neg q}$ and $\sigma_{p \neg q}$ (and relabel the other actions for clarity). The new action model A_2 is given in Figure 7.




Figure 7: A_2 – The adjusted action model for when Anne speaks French.

These new actions $\sigma_{p \neg q}$ and $\sigma_{\neg p \neg q}$ might be thought of as corresponding to situations where Carl announces p and Bob does not understand, and where Carl announces $\neg p$ and Bob does not understand, respectively. Anne can distinguish announcements based on their content, but not based on whether Bob understands them. Bob can distinguish announcements based on whether he understands them and, provided he understands them, based on their content as well. The updated model $M_1^{A_2}$ is given in Figure 8. As expected, Anne has learned whether or not p in the updated model (since she heard Carl), but she




Figure 8: $M_1^{A_2}$ – The updated model when Anne speaks French.

continues to be ignorant as to whether or not Bob speaks French (and, in turn, whether or not Bob has learned p). □

Note the duplication of effort in the construction of A_2 . The initial epistemic model M_1 already encodes the possibilities regarding Bob’s ability to speak French and Anne’s uncertainty about this. Yet our action model recapitulates this structure with actions that incorporate not just what Carl says, but also whether Bob understands it or not. Moreover, once our background assumptions are fixed (such

as whether Anne understands French or not), edges (and nodes!) in A_2 are determined, essentially, by examining M_1 and reading off what the uncertainties ought to be. Thus, while AL gives the impression of a clean, modular division between epistemic states and actions, in practice the two seems to be quite tangled, with unavoidable redundancies in their representations. To drive this point home, we sketch one further example.

Example 5. Consider an expanded epistemic model M_2 in which we take not just Bob's but also *Anne's* knowledge of French as endogenous: that is, suppose we also wish to represent Bob as being uncertain of whether or not Anne speaks French. A simple model of such a scenario might consist in eight worlds representing the possible combinations of truth values for primitive propositions p , q , and r , where p and q are interpreted as before and r stands for the proposition "Anne speaks French".

The action model A_2 of Example 4 is again inadequate. To see why, consider whether there ought to be an Anne-edge connecting σ_p and $\sigma_{\neg p}$. Intuitively, whether Anne can distinguish Carl announcing p from Carl announcing $\neg p$ depends on the world (i.e., it depends on whether or not Anne can speak French); it is not a fixed and unchanging truth that can be hard-coded into the model. As such, in order to capture this with a fixed action model we require, yet again, a proliferation of actions: e.g., actions of the form $\sigma_{pq\neg r}$, corresponding to something like Carl announcing p and Bob but not Anne understanding it. \square

These examples make it clear that the AL formalism is not well-suited to the practical task of building models to represent scenarios in which second-order knowledge is relevant: in addition to specifying the initial epistemic model, one must construct alongside it an elaborate space of actions fine-tuned to the specifics of the epistemic setting. These actions, rather than corresponding in a natural way with concrete events in the world (like Carl making an announcement), are individuated by details about agents' perceptions of them that seem less like part of the actions themselves and more like part of the background epistemics of the situation. We turn now to our proposed solution to this problem, which simplifies the action space considerably and imports the representation of higher-order uncertainty into the epistemic model, which is designed to handle it.

3 Adjusted Semantics

We propose a new semantics for modeling information update that subsumes AL and is able to capture higher-order uncertainty without the proliferation of actions illustrated in the previous section. In essence, we endogenize the action model, making the distinguishability of the actions world-dependent. It is therefore natural in our revised framework to drop the notion of a separate action model altogether.

In addition to a countable collection of primitive propositions PROP and a finite set of agents G , fix a set Σ of *epistemic actions* together with a *precondition* function $Pre : \Sigma \rightarrow \mathcal{L}_{EL}$ specifying the precondition for each action as before. A **dynamic model (over (Σ, Pre))** is a tuple

$$M = \langle W, \{\sim_j : j \in G\}, \{f_j : j \in G\}, V \rangle$$

where $\langle W, \{\sim_j : j \in G\}, V \rangle$ is an epistemic model, and for each agent j , $f_j : W \rightarrow 2^{\Sigma \times \Sigma}$ is a function from worlds to relations on actions. Intuitively, $(\sigma_1, \sigma_2) \in f_j(w)$ means that at world w , if action σ_1 is performed, then agent j cannot rule out σ_2 being the action performed. The f_j functions are the crucial novel component of our framework which allow action indistinguishability to vary from world to world.

We will use dynamic models to interpret the language \mathcal{L}_{DL} recursively defined by

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_j\varphi \mid [\sigma]\varphi,$$

where $[\sigma]\varphi$ is read, as before, “if action σ can be performed, then afterwards, φ is true”.⁶ Of course, in order to interpret the update modalities, we must define the notion of an updated model. Let

$$M^+ = \langle W^+, \{\sim_j^+ : j \in G\}, \{f_j^+ : j \in G\}, V^+ \rangle,$$
⁷

where:

$$\begin{aligned} W^+ &= \{(w, \sigma) : w \models \text{Pre}(\sigma)\} \\ (w, \sigma) \sim_j^+ (w', \sigma') &\Leftrightarrow w \sim_j w' \text{ and } (\sigma, \sigma') \in f_j(w) \\ f_j^+(w, \sigma) &= f_j(w) \\ (w, \sigma) \in V^+(p) &\Leftrightarrow w \in V(p). \end{aligned}$$

It is easy to see that M^+ is itself a dynamic model over (Σ, Pre) provided the relations \sim_j^+ are equivalence relations. In general, this need not be the case; however, the following conditions guarantee that it is.

(C1) If $w \sim_j w'$, then $f_j(w) = f_j(w')$.

(C2) Each $f_j(w)$ is an equivalence relation on Σ .

Proposition 6. (C1) and (C2) together imply that M^+ is a dynamic model.

Proof. As noted, it suffices to show that \sim_j^+ is an equivalence relation. Reflexivity follows immediately from reflexivity of \sim_j and $f_j(w)$. For symmetry, suppose that $(w, \sigma) \sim_j^+ (w', \sigma')$. Then $w \sim_j w'$ so also $w' \sim_j w$; moreover, by (C1) we have $(\sigma, \sigma') \in f_j(w) = f_j(w')$, so (C2) implies $(\sigma', \sigma) \in f_j(w')$, whence $(w', \sigma') \sim_j^+ (w, \sigma)$. For transitivity, suppose that $(w, \sigma) \sim_j^+ (w', \sigma')$ and $(w', \sigma') \sim_j^+ (w'', \sigma'')$. Clearly $w \sim_j w''$. Moreover, we have $(\sigma, \sigma') \in f_j(w)$ and $(\sigma', \sigma'') \in f_j(w')$; by (C1) $f_j(w) = f_j(w')$, and by (C2) this relation is transitive, so we deduce that $(\sigma, \sigma'') \in f_j(w)$, which shows that $(w, \sigma) \sim_j^+ (w'', \sigma'')$, as desired.⁸ □

Henceforth, we assume (C1) and (C2). Now define

$$(M, w) \models [\sigma]\varphi \quad \text{iff} \quad (M, w) \models \text{Pre}(\sigma) \text{ implies } (M^+, (w, \sigma)) \models \varphi.$$

This completes our specification of the new semantics for epistemic actions. It is not hard to see that update by an action model is essentially a special case of update in this framework: given an action model A , one simply defines each f_j to be the constant function such that $f_j(w) = \approx_j$.⁹ This directly realizes the intuition that the action model is common knowledge in AL.

⁶We no longer need to “tag” the update modality with an action model since our framework does not employ action models; this turns out to have important implications in our proof of completeness.

⁷This notation does not specify *which* action σ the update is being performed with respect to, since (as in AL) our update procedure effectively performs all available updates simultaneously.

⁸In fact, condition (C1) is stronger than it needs to be to establish Proposition 6: since \sim_j is assumed to be symmetric; weakening (C1) to:

(C1*) If $w \sim_j w'$, then $f_j(w) \subseteq f_j(w')$

has no effect. Of course, this argument would not work for weaker epistemic logics where agents are not assumed to be negatively introspective and \sim_j is not assumed to be symmetric. Exploring the interplay between properties of the relations \sim_j and properties of the functions f_j is intriguing, but takes us too far afield in the present work.

⁹Successive updates by different action models A_1, \dots, A_n can be recast as repeated updates by the same action model A , where A is an appropriate disjoint union of the A_k 's.

3.1 Revisiting the examples

We now revisit the problematic examples of Section 2.1 and show that the new framework we have developed actually addresses the deficiencies we demonstrated. Let $\Sigma = \{\sigma_p, \sigma_{\neg p}\}$ and set $Pre(\sigma_p) = p$ and $Pre(\sigma_{\neg p}) = \neg p$, corresponding to the two intuitive actions of Carl announcing p or announcing $\neg p$, respectively. Recall that in Examples 2 and 3, Carl delivers his message in French, we assume that Anne does not know French, and we let q represent the proposition that Bob knows French. We therefore define the dynamic model \tilde{M}_1 that we will use to reason about this scenario by extending the epistemic model M_1 depicted in Figure 4. In particular, we set

- $f_a(w_0) = f_a(w_1) = f_a(w_2) = f_a(w_3) = \Sigma \times \Sigma$,
- $f_b(w_0) = f_b(w_1) = id_\Sigma$, and
- $f_b(w_2) = f_b(w_3) = \Sigma \times \Sigma$,

where id_X denotes the identity relation on X . Thus, this dynamic model encodes the fact that Anne can never distinguish the two actions, whereas Bob can distinguish them just in case he speaks French. It is easy to see that (the epistemic part of) \tilde{M}_1^+ looks exactly like the model M_1^{A1} depicted in Figure 6, except with the nodes (w_2, σ) and (w_3, σ) relabeled (w_2, σ_p) and $(w_3, \sigma_{\neg p})$, respectively. In other words, our update produces the “right” epistemic results, and it does so using a simple and natural set of actions and without requiring any fine-tuning of the model beyond the basic association between worlds where Bob speaks French and worlds where he can distinguish Carl’s two possible announcements.

Next consider the scenario of Example 4, which is just like the previous one except it is assumed that Anne *does* speak French. To capture this, we need only change one line of the previous specifications for \tilde{M}_1 :

- $f_a(w_0) = f_a(w_1) = f_a(w_2) = f_a(w_3) = id_\Sigma$.

This corresponds directly to the assumption that Anne knows French (i.e., this is valid in the model), so she can always distinguish the two actions in question. Call this dynamic model \tilde{M}'_1 . Now as before, it is straightforward to check that (the epistemic part of) \tilde{M}'_1^+ looks exactly like the model M_1^{A2} given in Figure 8, provided we replace every instance of σ_{pq} and $\sigma_{p\neg q}$ with σ_p , and every instance of $\sigma_{\neg p q}$ and $\sigma_{\neg p \neg q}$ with $\sigma_{\neg p}$. So again, without the confusion of defining new, abstract actions, our framework reproduces the intended epistemic consequences of Carl’s announcement.

Finally, it is not hard to figure out how to define a dynamic model \tilde{M}_2 extending the epistemic model M_2 of Example 5:

- $f_a(w) = \begin{cases} id_\Sigma & \text{if } w \models r \\ \Sigma \times \Sigma & \text{if } w \models \neg r, \end{cases}$
- $f_b(w) = \begin{cases} id_\Sigma & \text{if } w \models q \\ \Sigma \times \Sigma & \text{if } w \models \neg q. \end{cases}$

4 Axiomatization

Our approach here follows much of the literature in seeking an axiomatization by way of *reduction schemes* that effectively transform statements in \mathcal{L}_{DL} into equivalent statements in some other language, and then axiomatizing that other language. Unlike PAL, however, it is easy to show that we cannot hope to reduce \mathcal{L}_{DL} to the basic epistemic language \mathcal{L}_{EL} —it is easy to produce pairs of dynamic models whose epistemic parts are bisimilar but which satisfy different formulas of \mathcal{L}_{DL} at bisimilar worlds.

Let \mathcal{L}_{DL}^+ and \mathcal{L}_{EL}^+ denote the languages \mathcal{L}_{DL} and \mathcal{L}_{EL} , respectively, augmented with additional primitive formulas $\xi_{j,\sigma,\sigma'}$ for each $j \in G$ and $\sigma, \sigma' \in \Sigma$. Interpret these new formulas in dynamic models as follows:

$$(M, w) \models \xi_{j,\sigma,\sigma'} \quad \text{iff} \quad (\sigma, \sigma') \in f_j(w).$$

Intuitively, $\xi_{j,\sigma,\sigma'}$ says that if action σ is performed, agent j cannot rule out that σ' was the action performed. Thus, \mathcal{L}_{EL}^+ can talk about both knowledge of the agents and action indistinguishability.

Somewhat surprisingly, \mathcal{L}_{DL} is reducible to \mathcal{L}_{EL}^+ : every formula of \mathcal{L}_{DL} is equivalent to a formula in \mathcal{L}_{EL}^+ , and this equivalence can be captured by reduction schemes that allow us to provide a sound and complete axiomatization of \mathcal{L}_{DL}^+ with respect to the class of all dynamic models that satisfy (C1) and (C2). It can be shown, however, that \mathcal{L}_{DL}^+ is strictly more expressive than \mathcal{L}_{DL} , so this result leaves something to be desired. This is the subject of ongoing research.

4.1 Soundness and completeness of \mathcal{L}_{EL}^+

Fix a finite set of actions Σ together with a precondition function $Pre : \Sigma \rightarrow \mathcal{L}_{EL}$. We begin by axiomatizing \mathcal{L}_{EL}^+ and then turn to \mathcal{L}_{DL}^+ . Consider the following axioms:

1. All instantiations of propositional tautologies
2. $K_a(\varphi \rightarrow \psi) \rightarrow (K_a\varphi \rightarrow K_a\psi)$
3. $K_a\varphi \rightarrow \varphi$
4. $K_a\varphi \rightarrow K_aK_a\varphi$
5. $\neg K_a\varphi \rightarrow K_a\neg K_a\varphi$
6. ξ axioms:
 - (a) $\xi_{j,\sigma,\sigma}$
 - (b) $\xi_{j,\sigma,\sigma'} \rightarrow \xi_{j,\sigma',\sigma}$
 - (c) $\xi_{j,\sigma,\sigma'} \rightarrow (\xi_{j,\sigma',\sigma''} \rightarrow \xi_{j,\sigma,\sigma''})$
7. Interaction axiom:¹⁰
 - (a) $\xi_{j,\sigma,\sigma'} \rightarrow K_j\xi_{j,\sigma,\sigma'}$
8. From φ and $\varphi \rightarrow \psi$, infer ψ .
9. From φ , infer $K_j\varphi$.

Soundness is easy, while completeness can be established by a fairly standard canonical model construction. Let $M^c = \langle W^c, \{\sim_j^c : j \in G\}, \{f_j^c : j \in G\}, V^c \rangle$ be defined as follows:

- $M^c = \{\Gamma \subseteq \mathcal{L}_{EL}^+ : \Gamma \text{ is maximally consistent}\}$
- $\Gamma \sim_j^c \Delta$ iff $\{K_j\varphi : K_j\varphi \in \Gamma\} = \{K_j\varphi : K_j\varphi \in \Delta\}$
- $(\sigma, \sigma') \in f_j^c(\Gamma)$ iff $\xi_{j,\sigma,\sigma'} \in \Gamma$
- $\Gamma \in V^c(p)$ iff $p \in \Gamma$.

¹⁰In fact, the scheme $\neg\xi_{j,\sigma,\sigma'} \rightarrow K_j\neg\xi_{j,\sigma,\sigma'}$ is also valid, but it can be proved from 7(a) together with the S5 axioms for K_j .

We must show that M^c is a dynamic model over (Σ, Pre) . That \sim_j^c is an equivalence relation follows from standard proofs of the completeness of S5. We show that f_j^c satisfies conditions (C1) and (C2).

(C1) Suppose that $\Gamma \sim_j^c \Delta$; we show that $f_j^c(\Gamma) = f_j^c(\Delta)$. Since \sim_j^c is an equivalence relation, suppose without loss of generality that $(\sigma, \sigma') \in f_j^c(\Gamma)$. Then $\xi_{j,\sigma,\sigma'} \in \Gamma$. By axiom 7a, $K_j \xi_{j,\sigma,\sigma'} \in \Gamma$; by the way that \sim_j^c is defined, we know then that $K_j \xi_{j,\sigma,\sigma'} \in \Delta$. Finally, by axiom 3, we have that $\xi_{j,\sigma,\sigma'} \in \Delta$, which gives the desired result that $(\sigma, \sigma') \in f_j^c(\Delta)$.

(C2) follows easily from the ξ axioms.

The equivalence

$$(M^c, \Gamma) \models \varphi \text{ iff } \varphi \in \Gamma$$

(i.e., the Truth Lemma) is proved in the standard way by structural induction.

4.2 Soundness and completeness of \mathcal{L}_{DL}^+

Let pre_σ abbreviate the epistemic formula $Pre(\sigma)$, and consider the following additional axioms:

10. Action axioms:

- (a) $[\sigma](\varphi \rightarrow \psi) \rightarrow ([\sigma]\varphi \rightarrow [\sigma]\psi)$
- (b) $[\sigma]p \leftrightarrow (pre_\sigma \rightarrow p)$
- (c) $[\sigma]\neg\varphi \leftrightarrow (pre_\sigma \rightarrow \neg[\sigma]\varphi)$
- (d) $[\sigma](\varphi \wedge \psi) \leftrightarrow ([\sigma]\varphi \wedge [\sigma]\psi)$
- (e) $[\sigma]K_a\varphi \leftrightarrow (pre_\sigma \rightarrow \bigwedge_{\sigma' \in \Sigma} (\xi_{a,\sigma,\sigma'} \rightarrow K_a[\sigma']\varphi))$

11. From φ , infer $[\sigma]\varphi$.

The soundness of most of the above axioms is immediate; we show the soundness of (10e).

(\Rightarrow) Suppose that $(M, w) \models [\sigma]K_a\varphi$ and assume that $(M, w) \models pre_\sigma$ (otherwise the equivalence is trivial). Let $\sigma' \in \Sigma$ be such that $w \models \xi_{a,\sigma,\sigma'}$ and consider any world w' with $w \sim_a w'$. We wish to show that $w' \models [\sigma']\varphi$. We know that $(M^+, (w, \sigma)) \models K_a\varphi$. Since $w \models \xi_{a,\sigma,\sigma'}$, we also know that $(\sigma, \sigma') \in f_a(w)$, so $(w, \sigma) \sim_a^+(w', \sigma')$; it follows that $(w', \sigma') \models \varphi$, which means that $w' \models [\sigma']\varphi$.

(\Leftarrow) Suppose that $w \models pre_\sigma \rightarrow \bigwedge_{\sigma' \in \Sigma} (\xi_{a,\sigma,\sigma'} \rightarrow K_a[\sigma']\varphi)$, and that $w \models pre_\sigma$. We wish to show that $(w, \sigma) \models K_a\varphi$. Consider any $(w', \sigma') \sim_a^+(w, \sigma)$; we will show that $(w', \sigma') \models \varphi$. Since $(\sigma, \sigma') \in f_a(w)$, we know that $w \models \xi_{a,\sigma,\sigma'}$, so by assumption, $w \models K_a[\sigma']\varphi$. Then since $w \sim_a w'$, we have that $w' \models [\sigma']\varphi$, which implies that $(w', \sigma') \models \varphi$, as desired.

For completeness, consider the following translation:

$$\begin{aligned}
t(p) &= p \\
t(\xi_{j,\sigma,\sigma'}) &= \xi_{j,\sigma,\sigma'} \\
t(\neg\varphi) &= \neg t(\varphi) \\
t(\varphi \wedge \psi) &= t(\varphi) \wedge t(\psi) \\
t(K_a\varphi) &= K_a t(\varphi) \\
t([\sigma]p) &= pre_\sigma \rightarrow p \\
t([\sigma]\neg\varphi) &= pre_\sigma \rightarrow \neg t([\sigma]\varphi) \\
t([\sigma](\varphi \wedge \psi)) &= t([\sigma]\varphi) \wedge t([\sigma]\psi) \\
t([\sigma]K_a\varphi) &= pre_\sigma \rightarrow \bigwedge_{\sigma' \in \Sigma} (\xi_{a,\sigma,\sigma'} \rightarrow K_a t([\sigma']\varphi)) \\
t([\sigma][\sigma']\varphi) &= t([\sigma]t([\sigma']\varphi))
\end{aligned}$$

Proposition 7. For all formulas $\varphi \in \mathcal{L}_{DL}^+$, $t(\varphi)$ is provably equivalent to φ and $t(\varphi) \in \mathcal{L}_{EL}^+$.

Proof. We proceed by induction on the *action nesting depth* of φ , defined in the obvious way:

$$\begin{aligned}
d(p) &= 0 \\
d(\xi_{j,\sigma,\sigma'}) &= 0 \\
d(\neg\varphi) &= d(\varphi) \\
d(\varphi \wedge \psi) &= \max(d(\varphi), d(\psi)) \\
d(K_j\varphi) &= d(\varphi) \\
d([\sigma]\varphi) &= d(\varphi) + 1.
\end{aligned}$$

The case $d(\varphi) = 0$ is immediate. So suppose the result holds for all formulas with nesting depth less than n , and let $\varphi \in \mathcal{L}_{DL}^+$ be such that $d(\varphi) \leq n$.

We now proceed via a subinduction on the *weight* of φ , defined as follows:

$$\begin{aligned}
w(p) &= 1 \\
w(\xi_{j,\sigma,\sigma'}) &= 1 \\
w(\neg\varphi) &= w(\varphi) + 1 \\
w(\varphi \wedge \psi) &= \max(w(\varphi), w(\psi)) + 1 \\
w(K_j\varphi) &= w(\varphi) + 1 \\
w([\sigma]\varphi) &= w(\varphi) + 1
\end{aligned}$$

The base case where $w(\varphi) = 1$ is again immediate. So suppose inductively the result holds for formulas of weight less than $w(\varphi)$. The proof now breaks into cases depending on the structure of φ , since this determines which recursive clause of the definition of t is relevant. The inductive steps corresponding to the Boolean connectives and the K_a modalities are straightforward, so we move to the case where $\varphi = [\sigma]\psi$; this in turn naturally breaks into several subcases depending on the structure of ψ :

- If $\psi = p$, then $t(\varphi) = t([\sigma]p) = pre_\sigma \rightarrow p$ and we are done by axiom (10b).

- If $\psi = \neg\chi$, then $t(\varphi) = t([\sigma]\neg\chi) = pre_\sigma \rightarrow \neg t([\sigma]\chi)$. Clearly $w([\sigma]\chi) < w(\varphi)$, so by the inductive hypothesis we know that $t([\sigma]\chi) \in \mathcal{L}_{EL}^+$ and is provably equivalent to $[\sigma]\chi$. It follows immediately that $t(\varphi) \in \mathcal{L}_{EL}^+$ and, by axiom (10c), that $t(\varphi)$ is provably equivalent to φ .
- If $\psi = \chi_1 \wedge \chi_2$, then $t(\varphi) = t([\sigma]\chi_1) \wedge t([\sigma]\chi_2)$. Clearly $w([\sigma]\chi_1) < w(\varphi)$ and $w([\sigma]\chi_2) < w(\varphi)$, so by the inductive hypothesis we know that $t([\sigma]\chi_1) \in \mathcal{L}_{EL}^+$ and $t([\sigma]\chi_2) \in \mathcal{L}_{EL}^+$, and they are provably equivalent to $[\sigma]\chi_1$ and $[\sigma]\chi_2$, respectively. It follows immediately that $t(\varphi) \in \mathcal{L}_{EL}^+$ and, by axiom (10d), that $t(\varphi)$ is provably equivalent to φ .
- If $\psi = K_j\chi$, then $t(\varphi) = t([\sigma]K_j\chi) = pre_\sigma \rightarrow \bigwedge_{\sigma' \in \Sigma} (\xi_{j,\sigma,\sigma'} \rightarrow K_j t([\sigma']\chi))$. Clearly, for each $\sigma' \in \Sigma$, $w([\sigma']\chi) < w(\varphi)$, so by the inductive hypothesis we know that $t([\sigma']\chi) \in \mathcal{L}_{EL}^+$ and is provably equivalent to $[\sigma']\chi$. It follows immediately that $t(\varphi) \in \mathcal{L}_{EL}^+$ and, by axiom (10e), that $t(\varphi)$ is provably equivalent to φ , as desired.
- Finally, if $\psi = [\sigma']\chi$, then $t(\varphi) = t([\sigma][\sigma']\chi) = t([\sigma]t([\sigma']\chi))$. Let $\tilde{\psi} = t([\sigma']\chi)$; clearly $w([\sigma']\chi) < w(\varphi)$, so by the inductive hypothesis we know that $\tilde{\psi} = t([\sigma']\chi) \in \mathcal{L}_{EL}^+$ and $\tilde{\psi}$ is provably equivalent to $[\sigma']\chi$. So we have $t(\varphi) = t([\sigma]t([\sigma']\chi)) = t([\sigma]\tilde{\psi})$. Now the weight of $\tilde{\psi}$ may be very large, so we can't apply our inner inductive hypothesis again here. However, since $\tilde{\psi} \in \mathcal{L}_{EL}^+$, it is easy to see that $d([\sigma]\tilde{\psi}) = 1 < d(\varphi)$, so we can appeal to our *outer* inductive hypothesis to conclude that $t(\varphi) = t([\sigma]\tilde{\psi}) \in \mathcal{L}_{EL}^+$ and is provably equivalent to $[\sigma]\tilde{\psi}$. Moreover, since $\tilde{\psi}$ is provably equivalent to $[\sigma']\chi$, using axiom (10a) it is easy to show that $[\sigma]\tilde{\psi}$ is provably equivalent to $[\sigma][\sigma']\chi$, whence $t(\varphi)$ is provably equivalent to φ , as desired. \square

Completeness is an immediate corollary.

References

- [1] A. Baltag & L. S. Moss (2004): *Logics for Epistemic Programs*. *Synthese* 139(2), pp. 165–224, doi:10.1023/B:SYNT.0000024912.56773.5e.
- [2] A. Baltag, L. S. Moss & S. Solecki (1998): *The logic of public announcements, common knowledge, and private suspicions*. In I. Gilboa, editor: *Proceedings of the 11th conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, pp. 43–56.
- [3] T. Bolander, H. van Ditmarsch, A. Herzig, E. Lorini, P. Pardo & F. Schwarzentruber (2016): *Announcements to Attentive Agents*. *Journal of Logic, Language, and Information* 25(1), pp. 1–25, doi:10.1007/s10849-015-9234-3.
- [4] H. van Ditmarsch, W. van der Hoek & B. Kooi (2008): *Dynamic Epistemic Logic*. Springer.
- [5] R. Fagin, J. Y. Halpern, Y. Moses & M. Y. Vardi (1995): *Reasoning About Knowledge*. MIT Press, Cambridge, Mass. A slightly revised paperback version was published in 2003.
- [6] B. Kooi & E. Pacuit (2011): *Logics of Rational Interaction*. In P. Girard, O. Roy & M. Marion, editors: *Dynamic Formal Epistemology*, Springer, pp. 5–32, doi:10.1007/978-94-007-0074-1_2.
- [7] J. Plaza (2007): *Logics of Public Communications*. *Synthese* 158, pp. 165–179, doi:10.1007/s11229-007-9168-7.