

Perception and Change in Update Logic

Jan van Eijck

Abstract Three key ways of updating one’s knowledge are (i) perception of states of affairs, e.g., seeing with one’s own eyes that something is the case, (ii) reception of messages, e.g., being told that something is the case, and (iii) drawing new conclusions from known facts. If one represents knowledge by means of Kripke models, the implicit assumption is that drawing conclusions is immediate. This assumption of logical omniscience is a useful abstraction. It leaves the distinction between (i) and (ii) to be accounted for. In current versions of Update Logic (Dynamic Epistemic Logic, Logic of Communication and Change) perception and message reception are not distinguished. This paper proposes an extension of Update Logic that makes this distinction explicit. The logic deals with three kinds of updates: announcements, changes of the world, and observations about the world in the presence of witnesses. The resulting logic is shown to be complete by means of a reduction to epistemic propositional dynamic logic by a well known method.

1 The Riddle of the Caps

‘I see nobody on the road,’ said Alice. ‘I only wish I had such eyes,’ the King remarked in a fretful tone. ‘To be able to see Nobody! And at that distance too!’

Lewis Carroll, *Alice in Wonderland*.

We start with a variation on the so-called ‘wise men puzzle’ [22]. Imagine four people standing in line, with three of them looking to the left, and one looking to the right. These fellows, let us call them 1, 2, 3, 4, each wear a cap. The leftmost guy, agent number 1, can see no-one. Agent 2 can see agent 1. Agent 3 can see agents

Jan van Eijck
CWI, Amsterdam, e-mail: jve@cwi.nl

1 and 2. Agent 4, finally, who has his head turned in the other direction, can see no-one. (I will assume in the story that these caps are worn by men, for ladies wear hats.)



Assume it is common knowledge what the perceptive capabilities of the agents are. So it is common knowledge that 1 can see no-one, that 2 can see 1 but no others, that 3 can see 1 and 2 but no others, and that 4, can see no-one. In particular it is commonly known that they all wear a cap, and that no-one can see the colour of his own cap.

Now let there be a public announcement ‘there are two white caps and two black caps’. After this, 3 knows the colour his cap, for he can reason: I see two white caps in front of me. There are only two white caps. So my own cap must be black. If 3 publicly announces ‘I know the colour of my cap’, then as a result 2 gets to know the colour of his cap as well. For 2 can reason as follows. If the guy behind me knows the colour of his cap, then this means that he sees two caps of the same colour in front of him. I see that 1 has a white cap, so my own cap must be white as well.

Now suppose the initial situation is like this: the participants all face in the same directions as before, but now 2 and 4 have swapped caps.



In this case, obviously, 3 does not know the colour of his cap. But suppose 3 now publicly announces ‘I do not know the colour of my cap’. The riddle of the caps is this: how does this announcement reveal to 2 the colour of *his* cap? And for people interested in epistemic model checking the riddle becomes: find a series of updates corresponding to what goes on in this scenario, and such that in the model at the end of the series of updates, the third agent knows the colour of his cap.

The ingredients that are used here are public announcements (‘There are two white caps and two black caps’, ‘I know/do not know the colour of my cap’) and common knowledge about the perceptive capabilities of the participants.

‘Common knowledge’ is a term coined by David Lewis in [21]. It is sometimes called ‘mutual knowledge’ ([27, 8, 9]), and it is essentially different from general knowledge. If a group of people get separate messages that there will be dinner party on Thursday, this generates general knowledge: we all know that there is a dinner, but I do not know that you know there is a dinner, you do not know that I know that there is a dinner, . . . If, instead, we all get the same email about the dinner, *and we see that we are all on the addressee list*, this generates common knowledge.

A similarity relation on a set of worlds W is an equivalence relation on W . Given individual similarity relations \sim_i for a $i \in I$, where I is a group of agents, I -common knowledge has the similarity relation given by:

$$\left(\bigcup_{i \in I} \sim_i \right)^*$$

The semantics of common knowledge is expressed in terms of a fixpoint operation, so it seems natural to assume that common knowledge is something that emerges in the limit of a series of steps of knowledge aggregation, that it is some idealized version of what happens when we acquire knowledge about knowledge in real life. Nothing could be further from the truth. Common knowledge is ubiquitous, and it comes about in a single step, in one of the following ways:

- by public announcement [25];
- by common witnessing of an event, when it is already common knowledge that all can perceive the event [8];
- by variations and combinations of the above ('indirect co-presence', 'cultural co-presence' [8]).

Since common knowledge emerges in no other ways than these (in particular, in a distributed environment, where messages or observations can be missed by an agent, common knowledge cannot emerge [18]), it would seem that creating common knowledge about perceptive abilities prepares the ground for creating common knowledge about perceived events.

The philosophical literature on perception — in particular the problems caused by perceptual illusion or hallucination — is vast. See [1, 10, 12, 13, 26] for a small sample. Rather than take all of this on board, I will draw some methodological limitations from the outset, in that I will assume that perceptual experiences are *events* rather than *states* (not all that is known in some state is perceived in that state), and that these events lead to true knowledge: no deception resulting in illusion takes place. The first of these is uncontroversial. It merely states that there is more to knowledge than perception. The second assumption is more questionable, but it has the considerable merit that it allows us to sidestep puzzles about hallucination. Since the limitation to accurate perception can easily be lifted in an extension of the logic that I will propose below, I feel there is no harm in these assumptions.

As one of my reviewers reminded me, accuracy of perception does not always lead to knowledge: one can still misjudge what one accurately perceives. Indeed, the link between perception and knowledge is rather involved, and perceptual evidence cannot always serve as a basis for knowledge, as was pointed out by Gettier in [17]. Also, clearly, it is possible in certain cases to know p without perceiving p : we might have learnt p by being told. If one of my colleague has a headache then she can directly perceive this, but I cannot. Still, I will know about it as soon as she tells me. In any case, I will assume that there is a sense of 'perceiving p ' in which perceiving p is sufficient for knowing p .

Analyzing the example puzzle with this in mind, can we see what is the nature of the common knowledge about perceptive abilities that is needed? Focussing on the example, it seems that the following should somehow become common knowledge between agents two and three:

- that the second person can see the first,

- that the third person can see the first and the second.

Common knowledge between sets of agents of these facts should be expressible in the language. Moreover, updates that make these facts common knowledge should also be expressible. It is therefore reasonable to start out with a logic of public announcement and common knowledge.

Assume a set of agents N and a set of propositions P . Letting i range over N and p over P , we define our basic epistemic language as:

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

Assume the usual abbreviations. In particular, $\phi_1 \vee \phi_2$ is shorthand for $\neg(\neg\phi_1 \wedge \neg\phi_2)$, $\phi_1 \rightarrow \phi_2$ for $\neg(\phi_1 \wedge \neg\phi_2)$, $\langle \alpha \rangle \phi$ for $\neg[\alpha]\neg\phi$. Finally, $[i]\phi$ is abbreviated as $K_i\phi$, and $[(\cup_{i \in I} i)^*]\phi$ as $C\phi$.

This language is interpreted in the usual multimodal Kripke models (or labeled transition systems). Such a model M is a triple (W, R, V) where W is a nonempty set of worlds, $R : N \rightarrow \mathcal{P}(W^2)$ assigns a similarity relation \sim_i to each agent $i \in N$, and $V : W \rightarrow \mathcal{P}(P)$ assigns a valuation to each world $w \in W$.

If M is given, we refer to its worlds component as W_M , its relational component as R_M , and its valuation component as V_M . Let M be given, and let $w \in W_M$. Then the truth definition is given by:

$$\begin{aligned}M \models_w \top & \text{ always} \\ M \models_w p & ::= p \in V_M(w) \\ M \models_w \neg\phi & ::= \text{not } M \models_w \phi \\ M \models_w \phi_1 \wedge \phi_2 & ::= M \models_w \phi_1 \text{ and } M \models_w \phi_2 \\ M \models_w [\alpha]\phi & ::= \text{for all } w' \text{ with } w \xrightarrow{\alpha} w' \text{ } M \models_{w'} \phi \\ M \models_w [!\phi_1]\phi_2 & ::= \text{if } M \models_w \phi_1 \text{ then } M \upharpoonright \phi_1 \models_w \phi_2, \\ & \text{where } M \upharpoonright \phi_1 \text{ is given by } W_{M \upharpoonright \phi_1} = \{w \in W_M \mid M \models_w \phi_1\}, \\ & \text{where } V_{M \upharpoonright \phi_1} \text{ is the restriction of } V_M \text{ to } W_{M \upharpoonright \phi_1}, \\ & \text{and where } R_{M \upharpoonright \phi_1} \text{ is the restriction of } R_M \text{ to } W_{M \upharpoonright \phi_1}.\end{aligned}$$

In this definition, $\xrightarrow{\alpha}$ is given by:

$$\begin{aligned}w \xrightarrow{i} w' & ::= w R_M(i) w' \\ w \xrightarrow{?\phi} w' & ::= w = w' \text{ and } M \models_w \phi \\ w \xrightarrow{\alpha_1; \alpha_2} w' & ::= \exists w'' \text{ with } w \xrightarrow{\alpha_1} w'' \text{ and } w'' \xrightarrow{\alpha_2} w' \\ w \xrightarrow{\alpha_1 \cup \alpha_2} w' & ::= w \xrightarrow{\alpha_1} w' \text{ or } w \xrightarrow{\alpha_2} w' \\ w \xrightarrow{\alpha^*} w' & ::= \exists w_1, \dots, w_n \text{ with } w = w_1, w' = w_n, w_i \xrightarrow{\alpha} w_{i+1} \\ & \text{for } 1 \leq i < n.\end{aligned}$$

Use $\llbracket \phi \rrbracket_M$ for $\{w \in W_M \mid M \models_w \phi\}$.

The logic presented here is in fact the public announcement restriction of the format of ‘Logics of Communication and Change’ or LCC [6], which is in turn a version of BMS style update logic [2].

Using b_i for ‘the cap of the i -th guy (counting from left to right) is black’ and $D_i(p)$ for $K_i p \vee K_i \neg p$, we get that update with $!D_i(p)$ expresses the public announcement: ‘ i can distinguish p from $\neg p$ ’. Here are formal versions of the public announcements that make the relevant epistemic abilities common knowledge. Note that we are not making a distinction yet between epistemic and perceptive abilities. The ‘seeing’ in the paraphrases below is just a metaphor for knowing. ‘The first guy cannot see the colour of any hat’ is rendered by (1):

$$\neg D_1(b_1) \wedge \neg D_1(b_2) \wedge \neg D_1(b_3) \wedge \neg D_1(b_4). \quad (1)$$

‘The second guy can see the colour of cap 1, but not the colours of hats 2,3,4’ is formalized by (2):

$$D_2(b_1) \wedge \neg D_2(b_2) \wedge \neg D_2(b_3) \wedge \neg D_2(b_4). \quad (2)$$


Formula (3) expresses ‘The third guy can see the colours of hats 1,2, but not the colours of hats 3,4’:

$$D_3(b_1) \wedge D_3(b_2) \wedge \neg D_3(b_3) \wedge \neg D_3(b_4). \quad (3)$$

Finally, ‘The fourth guy cannot see the colour of any hat’ is rendered by (4):

$$\neg D_4(b_1) \wedge \neg D_4(b_2) \wedge \neg D_4(b_3) \wedge \neg D_4(b_4). \quad (4)$$

In the semantics above, we have given a semantics of knowledge, not perception. One possible way to go in order to incorporate perception would be to extend each \sim_i knowledge relation to an \approx_i perception relation, add an operation for individual perception, say $\hat{\sim}$, stipulate that $\hat{\sim}$ is interpreted by \approx_i , and impose suitable conditions on the relations, in particular $\sim_i \subseteq \approx_i$. An example would look like this (solid lines for knowledge, dotted lines for perception without knowledge; the convention is that the \approx relation is given as the transitive closure of the union of the solid and the dotted lines):



Here and henceforth, we use shading to single out the actual world, and we leave out reflexive links. So the model pictures a situation where bb is the case, where the subject knows that ww is ruled out, while this knowledge is not backed up by perception. As far as the subject can perceive, all of bb , bw , wb and ww are possible.

The difficulty with this approach is that it is an attempt to accommodate perception as part of the statics of what goes on. Rather than going this way, we will model observations as actions that change epistemic models. This is in line with the general approach of update logic; compare also the remark made in [4] that epistemic update logic is already a logic of observations.

Given that perception is always perception of phenomena in a fleeting world, it is very natural to view perceptions as actions. What is eternal and unchanging cannot be perceived. To be perceivable is always the hallmark of being part of the world of phenomena, the world described by ‘esse est percipi’. Perception is simply *change in the world that gets noticed*, where the noticing may either be simultaneous with the change or may take place after the change. Given this, the distinction between knowledge and perception reveals itself when the world itself changes. If the changes in the world go unnoticed, knowledge may be lost. If the changes in the world get noticed, we can say that change is perceived, and as a consequence knowledge gets updated. The next sections will make clear what is meant by this.

2 Publicly Observable Change in the World


To flesh out the connection between perception and change, what is needed is an account of *effective abilities* and *perceptive abilities*:

- Which changes can be observed by which agents?
- Which agents are aware of the perceptive abilities of which other agents?
- Which perceptive abilities are *common knowledge*?

Systematic frameworks for answering these questions bring in the aspect of *perspective* in a natural way.

There is an important distinction between individual awareness and common knowledge. Consider the hats example again. Imagine that the four guys with the hats are standing in line again, or rather, that someone has put them in line with their eyes closed. Now the first guy opens his eyes, and realizes: ‘I am in a position where I see nobody’. Next, the second guy opens his eyes, and realizes: ‘I am in a position where I can see person 1 in front of me’. Then, the third guy opens his eyes, and realizes: ‘I am in a position where I can see person 1 and person 2 in front of me’. Finally, the fourth guy opens his eyes, and realizes: ‘I am in a position where I see nobody.’ Unless they all *publicly announce* their perception capabilities the ‘reasoning about knowledge and ignorance’ scenario of the example never gets off the ground.

Suppose the world changes from



to



That is: the second and third guy swap caps.

Before we indicate how to model all this in our logic, we take one intermediate step, by discussing *publicly observable change*. For this, we add bindings (single substitutions). Substitutions (finite sets of bindings) as a mechanism for modelling actual change were first explored in [14], and subsequently adopted in [6], and later in [20].

If P is the set of basic propositions, and L_P is a language over P , then an L_P -binding is a pair (p, ϕ) , where $p \in P$ and $\phi \in L_P$. We will write binding (p, ϕ) as $p \mapsto \phi$.

If $M = (W, V, R)$ is an epistemic model and $p \mapsto \phi$ is an L_P binding, then $V^{p \mapsto \phi}$ is the valuation that is like V for all q different from p , and has $V^{p \mapsto \phi}(p) = \llbracket \phi \rrbracket^M$. In other words, $V^{p \mapsto \phi}$ assigns to p the set of worlds w in which ϕ is true. For $M = (W, V, R)$, call $M^{p \mapsto \phi}$ the model given by $(W, V^{p \mapsto \phi}, R)$. The model $M^{p \mapsto \phi}$ is the result of updating M with a publicly observable change of p to ϕ .

Now we can handle publicly perceived change as follows. Suppose it is common knowledge that an agent i can epistemically discriminate between p and $\neg p$. This means that formula (5) holds:

$$CD_i p. \tag{5}$$

Suppose the the facts of the world change: the truth value of p gets swapped, and this change is observable by all. This means the following substitution is applied:

$$p \mapsto \neg p.$$

Then if it was commonly known before the substitution that $D_i p$, then, since the change was publicly observable, afterwards this same fact will still be common knowledge. In other words, principle (6) holds:

$$CD_i p \rightarrow [p \mapsto \neg p] CD_i p. \quad (6)$$

And if it is commonly known before the substitution that $K_i p$, then it will be commonly known after the substitution that $K_i \neg p$:

$$CK_i p \rightarrow [p \mapsto \neg p] CK_i \neg p. \quad (7)$$

This illustrates the creation of common knowledge without public announcement. More precisely, this is an example of creation of common knowledge by means of *co-presence*.

The example illustrates one of the mechanisms behind the creation of common knowledge by co-presence: publicly perceived change in a situation where there is common knowledge about epistemic ability.


The simplest way to ensure that adding a change operator to an epistemic logic does not have the effect that knowledge degenerates into mere belief is to make all changes *public* changes. The public change $p \mapsto \perp$ changes the truth value of p to false everywhere in the model, so that $\neg p$ becomes common knowledge.

3 Perceptive Ability and Perspective Shifts

Notice that the assumption that all change is publicly observable does not quite fit our example: if the cap of the third guy changes colour, nobody can see this. What we need to model this kind of change is a way of keeping track of the agents that are able to observe the change.

These perceptive abilities themselves may change too, for a perspective shift is just another kind of change in the world.

Suppose that the discriminative abilities of the agents in the caps example are *common knowledge*, and suppose the situation then changes: the third guy turns around, and is now looking to the right. What this means is that the world changes from



to



Assume that nobody can see this. Then the effect is dramatic. Unless we assume that the other guys knew all along that it was possible that agent 3 might turn around, this already turns the model into a non-S5 model. For if the first guy gets a different colour cap, then the second guy will still believe that the third guy can observe this, and so agent 2 draws the wrong conclusions.

It follows that to prevent ‘knowledge degeneration’, in a system of knowledge and perception, loss of perceptive ability due to perspective shifts *must* either be made common knowledge, or the perspective shifts must be handled in such a way that everyone takes it at as conceivable that they occur.

Our modelling right now is not strong enough to capture this. What is needed is a new way to interpret change operations, a way that takes the perceptive abilities of the agents into account.

In a logic of perception we want to be able to distinguish between agents that can observe certain changes and agents that cannot. If we still want to ensure that updates with such partly perceived changes preserve S5 (do not change knowledge into mere belief), then we need a subtle definition for updates that change facts, by including information about who can observe the change.


To account for the facts that not all change is immediately perceived and that changes can only be perceived by agents who are in the position to perceive them, we will separate the changes from the observations about what has changed.

What we need in fact is a separation between (unobserved) change in the world and the acts of perception witnessing the change. Before we can go into details, we must briefly discuss a philosophical problem: how can one possibly maintain true knowledge of a fleeting world in which things are bound to happen of which one is not aware? The solution to this cannot be to model the infinity of ways in which the world might change, and update with that. Our representations are finite, anyway, as we are focussing on what a finite number of agents know about a finite number of elementary facts. How does one model that fact p changes without the agents being aware of the result of the change? We will model this as an action (the change that really happens) that is indistinguishable from the trivial action (where nothing happens). Suppose p changes to false, but no agent is aware of this change. Then the result will be a model where for every ‘old’ situation there now is a pair of ‘new’ situations, such that no agent can see the difference between the members of the pairs, and with the members of the pairs looking exactly the same, except for the possible difference that in one member of the pair p is false.

Suppose I locked the front door. Immediately after I locked it I *know* it is locked. But as soon as I turn my back, a change in the world might occur, and I *know* this, for I know there are others with keys to the door. Someone else with a key might unlock the door, and as soon as this *actually* happens, I *actually* lose my initial knowledge. This is an example where change in the world of which I am not even aware changes my state of knowledge. The way we cope with these situations in real life is by taking into account that certain changes might occur. That’s why I am not surprised if I find the door unlocked some time after I locked it.

In this paper I will assume that agents adapt to a changing world by being aware of the areas where knowledge may get lost. Specifically, what agents do is “not rule

out what actually happens” (always a wise strategy in life). In the action model $A_{p \mapsto \phi}$ for unobserved change $p \mapsto \phi$, all agents consider the change possible. The action model looks as follows (note that reflexive arrows are not drawn):




The action model has two states, 0 and 1, of which 0 is actual (indicated by the grey shading). The actual action has no precondition (or: its precondition equals \top), and changes the value of p to ϕ . However, the agents do not observe this immediately, for they confuse this action with \top (the action where nothing happens).

This was the action model for unobserved change. Next, we turn to the power of perception. We will represent a perception or observation of ϕ by agent i , witnessed by group of agents G , as (i, ϕ, G) , where it is assumed that $i \in G$ (agents are aware of their own perceptive abilities). Agents in G different from i do not themselves make the distinction between ϕ and $\neg\phi$, but they are aware that i makes that distinction.

We will use \mathbb{P}_i^p to express that $p \in \mathbb{P}(i)$, i.e., that agent i can perceive p . Let $\text{voc}(\phi)$ be the vocabulary (set of proposition letters) of ϕ . Then a perception of ϕ by i presupposes that the formula $\bigwedge_{p \in \text{voc}(\phi)} \mathbb{P}_i^p$ is true. Abbreviate this as $\mathbb{P}_i^{\text{voc}(\phi)}$. We will use this formula as a precondition of the actual perception.

An action model for perception (i, ϕ, G) takes the following shape:




The actual observation can be either ϕ or $\neg\phi$ (both actions are actual in the action model, indicated by the grey shading), and since it is i who is making this observation, other agents cannot distinguish between these two actions, for otherwise they too would have learned whether ϕ is true. But the witnesses learn that i now can distinguish the ϕ worlds (including the real world) from the $\neg\phi$ worlds. The agents that have not witnessed the observation cannot distinguish this from the event where nothing happens. This is given by the \overline{G} links to the \top event 2.

To see the effect of this, picture a situation where there are three agents i, j, k , and a true fact p that is known to i, j but not to k :



Assume that k has the power to observe p , and that a perception by k of the value of p takes place, witnessed by i but not by j . Then this action has the following form:




The effect of the update of M_0 with this act of perception is the following:

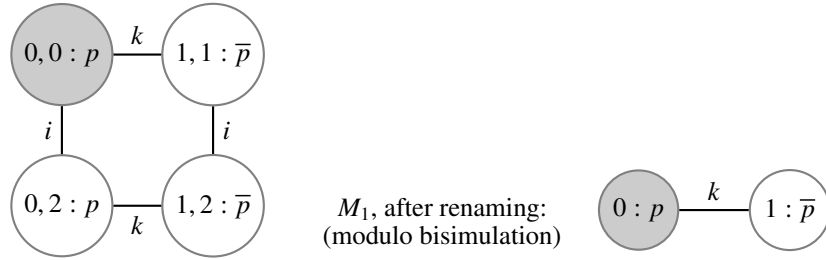


The result of the act of observation is that j , the agent who failed to witness the observation, no longer knows whether k knows about p or not.

Now let us play a different scenario. Assume that j has the power to observe the value of p , and that j actually observes the value of p , witnessed by k . The update model for $(j, p, \{j, k\})$, i.e., for the perception of the value of p by j , witnessed by k , but not i , has the form:



The result of the update of M_0 with this act of perception is the following model:




Perceptions of known facts change nothing, except when a witness of the perception does not yet know that the perceiver knows already. For an example of how a perception of a known fact can change a model, consider the initial model where both j and k are ignorant about p :



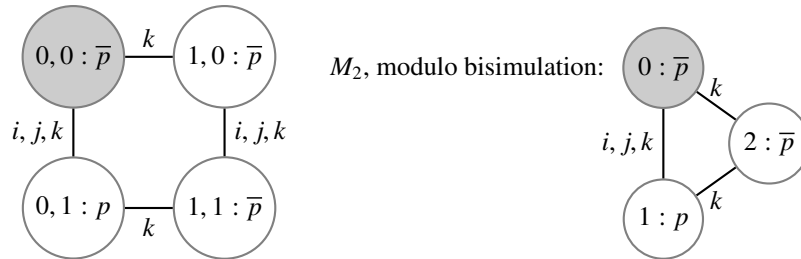
If we update this with the update model for $(j, p, \{j, k\})$, we get M_1 back. Note that the act of perception does not have as a result that j learns p , for j knew p already. The effect is in the witnessing of the perception by k but not i . In M_0 , k does not know p , but k does know that j can distinguish the p from the non- p situations. In other words, k has learned that j knows whether p , because k was a witness of the perception event.

Next, assume that the world changes, and p is made false everywhere. The action model for this is the following:




The result of making p false everywhere in model M_0 , publicly visible to all, would be a model with a single world where p is false, for everyone would know that p is false.

But the result of updating M_0 with the action model $A_{p \mapsto \perp}$ is more complicated. It is the following new epistemic model (call it M_2):




The knowledge of i, j about p has disappeared: i, j no longer know whether p is true or not. And indeed, the knowledge of k that i, j can distinguish p from $\neg p$ has disappeared as well.

Now suppose j makes a new observation. This time j observes the value of p , witnessed by k . This observation is called $(j, p, \{j, k\})$. Its action model looks like this:




The result of updating M_2 with this new observation looks rather complicated. Modulo bisimulation, it has 6 worlds $0, \dots, 5$, with p true in $\{4, 5\}$ and false in the other worlds:




Such models are a bit hard to construct by hand, but they can easily be found by using an epistemic model checker [15].

Observations that are witnessed by all are better behaved. The action model for $(j, p, \{i, j, k\})$, where $\{i, j, k\}$ equals the set of all agents, is given by:



Update of M_2 with this observation gives:



To construct an action model for the combination of a change and an act of perception directly after the change, we need a bit of Hoare style correctness reasoning [19]. A Hoare triple for postcondition $\mathbb{P}_i^p \wedge p$ after change $p \mapsto \phi$ looks as follows:

$$\{\mathbb{P}_i^{\text{voc}(\phi)} \wedge \phi\} p \mapsto \phi \{\mathbb{P}_i^p \wedge p\}.$$


This expresses that ϕ is the weakest precondition for making p true after the assignment $p \mapsto \phi$, and that the weakest precondition for being able to perceive p after the assignment is: being able to perceive every proposition letter in the vocabulary of ϕ . Note that it is not assumed that i perceives this change. But the Hoare triple can be used to construct the action model for change of p immediately perceived by i .

Similarly, we have:


$$\{\mathbb{P}_i^{\text{voc}(\phi)} \wedge \neg\phi\} p \mapsto \phi \{\mathbb{P}_i^p \wedge \neg p\}.$$

The weakest precondition for making p false after assignment $p \mapsto \phi$ is $\neg\phi$, while the weakest precondition that making p observable by i after the assignment is $\mathbb{P}_i^{\text{voc}(\phi)}$.


Combining the models for unnoticed change and perception with witnesses, we can construct an action model for perceived change. The action model for perceived change $(i, p, p \mapsto \phi, G)$ (perception by i of p after a change in p has taken place in the model, with G as witnesses of the act of perception) takes the following shape:



As an example of the effect of updating with this, here is the updating result of model M_0 with $A_{(i,p,p \mapsto \perp,(i,j))}$:



Notice that the effect is not different from that of $A_{(j,p,p \mapsto \perp,(i,j))}$, for the action models for these two actions boil down to the same thing:



This is the kind of perceived change that goes on in the swapping caps example that we started with. If the second guy in the row gets his white cap replaced by

a black cap, then this change is noticed by the third guy. This is an example of an application of the perceived change action model $A_{(3,b_2,b_2 \mapsto \top, \{3\})}$. The change that takes place concerns variable b_2 representing the blackness of the cap of the second guy, and agent 3 is the only one witnessing the change.

It should be clear how the present framework can be extended with perspective changes. A perspective change is nothing but a change in the function \mathbb{P} . Such changes can themselves be modelled by means of bindings $\mathbb{P}(i) \mapsto Q$ with $Q \subseteq P$. The model-changing effect of such an operation is that the \mathbb{P} function of the model gets modified. Again, there are various versions of this that preserve the S5 properties of models, with the two extremes being (i) perspective changes witnessed by all, and (ii) unobserved perspective changes. We settle for the first in our logic of the next section, but the treatment of the second is similar.

A perspective change witnessed by all can be represented by an action model as follows:

$$A_{\mathbb{P}(i) \mapsto Q} \quad \boxed{\top, \mathbb{P}(i) \mapsto Q}$$

To handle the effect of this, we assume that every Kripke model has an associated \mathbb{P} function, with type $\mathbb{P} : N \rightarrow \mathcal{P}(P)$. The model change that $A_{\mathbb{P}(i) \mapsto Q}$ brings about is a map from models M to models $M^{\mathbb{P}(i) \mapsto Q}$, where $M^{\mathbb{P}(i) \mapsto Q}$ is the result of modifying the \mathbb{P} function of M to \mathbb{P}' , where \mathbb{P}' is given by: $\mathbb{P}'(i) = Q$, and $\mathbb{P}'(j) = \mathbb{P}(j)$ for j different from i .

Common perception of change can be represented as change immediately followed by observation witnessed by others. Common perception by i and j of a change in p can get modelled as the following sequence of updates:

$$(p \mapsto \phi), (i, p, \{i, j\}), (j, p, \{i, j\}).$$

The logic we are about to present accurately models how common perception of change can create common knowledge.

A related issue is to explain how *common knowledge about limitations to perception* can create common knowledge. If we can *all* see that a person is blind, then her resulting lack of perception of change can become common knowledge (among all who see that she is blind). In the present framework a change of affairs where i cannot perceive p anymore is modelled as $\mathbb{P}_i^p \mapsto \perp$. In a framework where some perceptive abilities are unperceived, an update with $!\neg\mathbb{P}_i^p$, i.e. a public announcement that i cannot perceive p , would be needed to make this change visible to all.

Our modelling assumes that perceptive abilities are always common knowledge. If one wants to drop that assumption, to allow for (partially) unobserved perspective shifts, one needs a separate \mathbb{P} function for every world in a model. We leave the details of this to the reader.

4 A Logic of Change and Perception

Summing up, what do we need to develop a logic of perception and perceived change?

Here is a language that allows for public announcements, for unperceived changes, and for observations on those changes made by agents. This is just an example of a possible set-up: we fix a particular set of action operators, suitable for the issues we want to deal with, and we develop its logic.

$$\begin{aligned} \phi &::= \top \mid p \mid \mathbb{P}_i^p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\alpha]\phi \mid [!\phi_1]\phi_2 \mid [\beta]\phi \mid [\gamma]\phi \mid [\delta]\phi \\ \alpha &::= i \mid ?\phi \mid \alpha_1 \cup \alpha_2 \mid \alpha_1 ; \alpha_2 \mid \alpha^* \\ \beta &::= (\mathbb{P}(i) \mapsto Q) \text{ where } Q \subseteq P \\ \gamma &::= (p \mapsto \phi) \\ \delta &::= (i, \phi, G) \text{ where } i \in G \subseteq N \end{aligned}$$

The new elements are basic propositions of the form \mathbb{P}_i^p , expressing that i has the power to observe or perceive (changes in) p , and the modal operators β for perspective shifts witnessed by all, γ for unobserved change and δ for an observation by an agent witnessed by a group. Again, abbreviate as usual.


To interpret the language, we need to extend epistemic models with a function \mathbb{P} , as sketched above. If M is an epistemic model, \mathbb{P}_M is its \mathbb{P} function. The truth conditions for the new modal constructs were all given above. Note that we do not require that \mathbb{P}_i^p implies $K_i p \vee K_i \neg p$, for we wish to allow for the possibility that an unperceived change of p has occurred that i has not yet observed.

The logic of perception and change has the same expressive power as Epistemic PDL. This can be seen by performing a reduction in two stages:

- Translate into LCC, by replacing the β , γ and δ operators by the corresponding LCC action models. Using $^\circ$ for the translation operation, key clauses are: $([\beta]\phi)^\circ$ equals $[A_\beta]\phi^\circ$, $([\gamma]\phi)^\circ$ equals $[A_\gamma, 0]\phi^\circ$, and $([\delta]\phi)^\circ$ equals $[A_\delta, 0]\phi^\circ$, where A_γ is the action model for unobserved change γ , and 0 is the actual action of that action model, and where A_δ is the action model for observation δ , and 0 is the actual action of that action model. (Translate $[!\phi_1]\phi_2$ as $[A_{!\phi_1}]\phi_2$, where $A_{!\phi_1}$ is the action model for public announcement.)
- Reduce the resulting fragment of LCC to PDL, using the approach of [6]. See below.

For the reduction clauses for public announcements we refer to the literature. A key clause in the reduction clauses for perspective shifts is: $[A_{\mathbb{P}_i^p \mapsto Q}]\mathbb{P}_i^p := \top$ if $p \in Q$, and \perp otherwise.

We will give the reduction clauses for unobserved change in detail. Consider the action model for this as a finite automaton:




Define functions $T_{00}^{p \mapsto \phi}, T_{01}^{p \mapsto \phi}, T_{11}^{p \mapsto \phi}, T_{10}^{p \mapsto \phi}$ on the set of regular epistemic expressions, where $T_{ij}^{p \mapsto \phi}(\alpha)$ describes the joint effect of a step-by-step parallel transition through the action model and through α , starting in state i and ending in state j of the action model. Abbreviating, $p \mapsto \phi$ as γ , the definitions run as follows:

$$\begin{aligned}
 T_{00}^{\gamma}(i) &:= i, T_{01}^{\gamma}(i) := i, & T_{11}^{\gamma}(i) &:= i, T_{10}^{\gamma}(i) := i \\
 T_{00}^{\gamma}(\psi) &:= ?([p \mapsto \phi]\psi) \\
 T_{01}^{\gamma}(\psi) &:= ?\perp \\
 T_{11}^{\gamma}(\psi) &:= ?\psi \\
 T_{10}^{\gamma}(\psi) &:= ?\perp.
 \end{aligned}$$

Next, for composite epistemic relations:

$$\begin{aligned}
 T_{00}^{\gamma}(\alpha_1; \alpha_2) &:= (T_{00}^{\gamma}(\alpha_1); T_{00}^{\gamma}(\alpha_2)) \cup (T_{01}^{\gamma}(\alpha_1); T_{10}^{\gamma}(\alpha_2)) \\
 T_{01}^{\gamma}(\alpha_1; \alpha_2) &:= (T_{00}^{\gamma}(\alpha_1); T_{01}^{\gamma}(\alpha_2)) \cup (T_{01}^{\gamma}(\alpha_1); T_{11}^{\gamma}(\alpha_2)) \\
 T_{11}^{\gamma}(\alpha_1; \alpha_2) &:= (T_{11}^{\gamma}(\alpha_1); T_{11}^{\gamma}(\alpha_2)) \cup (T_{10}^{\gamma}(\alpha_1); T_{01}^{\gamma}(\alpha_2)) \\
 T_{10}^{\gamma}(\alpha_1; \alpha_2) &:= (T_{10}^{\gamma}(\alpha_1); T_{00}^{\gamma}(\alpha_2)) \cup (T_{11}^{\gamma}(\alpha_1); T_{10}^{\gamma}(\alpha_2)) \\
 T_{00}^{\gamma}(\alpha_1 \cup \alpha_2) &:= T_{00}^{\gamma}(\alpha_1) \cup T_{00}^{\gamma}(\alpha_2) \\
 T_{01}^{\gamma}(\alpha_1 \cup \alpha_2) &:= T_{01}^{\gamma}(\alpha_1) \cup T_{01}^{\gamma}(\alpha_2) \\
 T_{11}^{\gamma}(\alpha_1 \cup \alpha_2) &:= T_{11}^{\gamma}(\alpha_1) \cup T_{11}^{\gamma}(\alpha_2) \\
 T_{10}^{\gamma}(\alpha_1 \cup \alpha_2) &:= T_{10}^{\gamma}(\alpha_1) \cup T_{10}^{\gamma}(\alpha_2)
 \end{aligned}$$

For the final case, the reduction for expressions of the form α^* , we must bear in mind that the regular languages generated by the generic two-state automaton with the same structure as the change action model, and with various choices for start state and final state.



The language for the automaton with start state 0 and final state 0 can be characterized as:

$$A^*(BC^*DA^*)^*$$

For the same automaton, with 0 as start state and 1 as final state, we get:

$$A^*BC^*(DA^*BC^*)^*$$

For the automaton with 1 as start and as stop state:

$$C^*(DA^*BC^*)^*$$

And finally, for the automaton with 1 as start state and 0 as stop state:

$$C^*DA^*(BC^*DA^*)^*$$

Replacing A, B, C, D by $T_{00}^\gamma, T_{01}^\gamma, T_{11}^\gamma$ and T_{10}^γ , respectively, we get the following recipe for transforming an epistemic expression of the form α^* :

$$T_{00}^\gamma(\alpha^*) := (T_{00}^\gamma(\alpha))^*; (T_{01}^\gamma(\alpha); (T_{11}^\gamma(\alpha))^*; T_{10}^\gamma(\alpha); (T_{00}^\gamma(\alpha))^*)^*$$

$$T_{01}^\gamma(\alpha^*) := (T_{00}^\gamma(\alpha))^*; T_{01}^\gamma(\alpha); (T_{11}^\gamma(\alpha))^*; (T_{10}^\gamma(\alpha); (T_{00}^\gamma(\alpha))^*; T_{01}^\gamma(\alpha); (T_{11}^\gamma(\alpha))^*)^*$$


$$T_{11}^\gamma(\alpha^*) := (T_{11}^\gamma(\alpha))^*; (T_{10}^\gamma(\alpha); (T_{00}^\gamma(\alpha))^*; T_{01}^\gamma(\alpha); (T_{11}^\gamma(\alpha))^*)^*$$

$$T_{10}^\gamma(\alpha^*) := (T_{11}^\gamma(\alpha))^*; T_{10}^\gamma(\alpha); (T_{00}^\gamma(\alpha))^*; (T_{01}^\gamma(\alpha); (T_{11}^\gamma(\alpha))^*; T_{10}^\gamma(\alpha); (T_{00}^\gamma(\alpha))^*)^*$$

Reduction axioms for $[\gamma][\alpha]\psi$ now run as follows:

$$\begin{aligned}
[\gamma]\psi &\leftrightarrow [A_\gamma, 0]\psi \\
[A_\gamma, 0]\top &\leftrightarrow \top \\
[A_\gamma, 1]\top &\leftrightarrow \top \\
[A_\gamma, 0]p &\leftrightarrow \phi \\
[A_\gamma, 0]q &\leftrightarrow q \quad (q \text{ different from } p) \\
[A_\gamma, 1]q &\leftrightarrow q \\
[A_\gamma, 0]\neg\psi &\leftrightarrow \neg[A_\gamma, 0]\psi \\
[A_\gamma, 1]\neg\psi &\leftrightarrow \neg[A_\gamma, 0]\psi \\
[A_\gamma, 0](\psi_1 \wedge \psi_2) &\leftrightarrow [A_\gamma, 0]\psi_1 \wedge [A_\gamma, 0]\psi_2 \\
[A_\gamma, 1](\psi_1 \wedge \psi_2) &\leftrightarrow [A_\gamma, 1]\psi_1 \wedge [A_\gamma, 1]\psi_2 \\
[A_\gamma, 0][\alpha]\psi &\leftrightarrow [T_{00}^\gamma(\alpha)][A_\gamma, 0]\psi \wedge [T_{01}^\gamma(\alpha)][A_\gamma, 1]\psi \\
[A_\gamma, 1][\alpha]\psi &\leftrightarrow [T_{11}^\gamma(\alpha)][A_\gamma, 1]\psi \wedge [T_{10}^\gamma(\alpha)][A_\gamma, 0]\psi
\end{aligned}$$

Next, we can perform the same trick for action models A_δ . Here, the action model automaton is a bit more complicated:



But the procedure is the same. Only this time we need epistemic program transformers $T_{00}^\delta, T_{01}^\delta, T_{02}^\delta, T_{10}^\delta, T_{11}^\delta, T_{12}^\delta, T_{20}^\delta, T_{21}^\delta, T_{22}^\delta$. The basic reduction axiom for $[\delta]\psi$ runs

$$[\delta]\psi \leftrightarrow [A_\delta, 0]\psi \wedge [A_\delta, 1]\psi$$

reflecting the fact that A_δ has two actual states 0 and 1. Assuming δ equals (i, ϕ, G) , $T_{02}^\delta(a)$ will be defined as $?(P^{\text{voc}(\phi_i)} \wedge \phi); a$ if $a \notin G$, and $?\perp$ otherwise. Rather than spell out further details we refer to the general description given in [6].

Reduction axioms + PDL axioms + S5 knowledge axioms for $[i]$ now give a complete axiomatisation.

Theorem 1. *The logic of change and perception is complete.*

Theorem 2. *The logic of change and perception has the same expressive power as epistemic PDL.*

We briefly return to the motivating example: the puzzle of the wise men. Using the vocabulary b_1, b_2, b_3, b_4 , we can model the perceptive abilities as the function \mathbb{P} given by:

$$\begin{aligned}\mathbb{P}(1) &= \emptyset \\ \mathbb{P}(2) &= \{b_1\} \\ \mathbb{P}(3) &= \{b_1, b_2\} \\ \mathbb{P}(4) &= \emptyset\end{aligned}$$

This says that the first and fourth agent can observe the colour of no-one, the second agent can observe the colour of the first agent, and the third agent can observe the colour of the second agent.

The perceptive abilities are assumed to be common knowledge; indeed this assumption is a crucial element in the solution of the puzzle. This means that all observations that take place have the set of *all* agents as witnesses. Consider again the first version of the puzzle.



Assume an initial situation with common knowledge of universal ignorance: a model with $4^2 = 16$ worlds, for all possible combinations of cap colours.

To get at the correct epistemic representation for the epistemic situation where the epistemic exchange ('I do not know my colour', 'I do know my colour now') can take place, we need two public observations plus one public announcement. The formal version of 'there are two white caps and two black caps' is:

$$\begin{aligned}\phi &= (b_1 \wedge b_2 \wedge \neg b_3 \wedge \neg b_4) \vee (\neg b_1 \wedge \neg b_2 \wedge b_3 \wedge b_4) \vee (b_1 \wedge \neg b_2 \wedge \neg b_3 \wedge b_4) \\ &\quad \vee (\neg b_1 \wedge b_2 \wedge b_3 \wedge \neg b_4) \vee (\neg b_1 \wedge b_2 \wedge \neg b_3 \wedge b_4) \vee (b_1 \wedge \neg b_2 \wedge b_3 \wedge \neg b_4).\end{aligned}$$

The four updates we need are:

$$(2, \neg b_1, \{1, 2, 3\}), (3, \neg b_1, \{1, 2, 3\}), (3, \neg b_2, \{1, 2, 3\}), !\phi.$$

After the first update, it is common knowledge that agent 2 has observed the colour of the first cap. After the second and the third update, it is common knowledge that agent 3 has observed the colours of the first and the second cap, and the fourth update is the public announcement "Two of the caps are black and the other two are white".

Note that the first update can also be represented as $(2, b_1, \{1, 2, 3\})$, and that the second update can also be represented as $(3, b_1, \{1, 2, 3\})$. Also note that the pair of updates $(3, \neg b_1, \{1, 2, 3\}), (3, \neg b_2, \{1, 2, 3\})$ is stronger than the single update $(3, \neg b_1 \wedge \neg b_2, \{1, 2, 3\})$.

In our logic the following statement will hold in the initial situation of common knowledge of universal ignorance:

$$[(2, b_1, \{1, 2, 3\})][(3, \neg b_1, \{1, 2, 3\})][(3, \neg b_2, \{1, 2, 3\})][!\phi]K_3 b_3.$$

Next assume that the caps of the second and the third agent get swapped, but the agents are unaware of this: they are not paying attention, or they have closed their eyes. This is modelled by the following updates:

$$b_2 \mapsto \top, b_3 \mapsto \perp.$$

After updating with this, agent 3 has lost his knowledge:

$$[b_2 \mapsto \top][b_3 \mapsto \perp]\neg K_3 b_3.$$

The actual world has changed into:



The epistemic effect of the unobserved changes is that we are back at square one: no agent knows anything anymore about who is wearing which hat.

So one gets from the initial model of the first example to the initial model for the second example by means of unobserved change. The knowledge from the earlier witnessed observations has got lost.

New observations can remedy this. In the new universal ignorance model the following statements will hold:

$$[(2, b_1, \{1, 2, 3\})][(3, \neg b_1, \{1, 2, 3\})][(3, b_2, \{1, 2, 3\})][!\phi]\neg K_3 b_3 \wedge \neg K_3 \neg b_3.$$

That is, after a witnessed observation by agent two that the cap of agent one is white and witnessed observations by agent three that the hat of agent one is white and that of agent two is black, and a public announcement that there are still two black hats and two white hats, we get into a situation where the third agent does not know his hat colour.

But when *this* gets publicly announced, the announcement of ignorance by the third agent reveals to the second agent what is his cap colour:

$$[(2, b_1, \{1, 2, 3\})][(3, \neg b_1, \{1, 2, 3\})][(3, b_2, \{1, 2, 3\})][!\phi][!\phi][!\phi]\neg K_3 b_3 \wedge \neg K_3 \neg b_3]K_2 b_2.$$

Properties such as these can be (and have been) verified by means of epistemic model checking.

5 Conclusion and Further Work

Related work in (dynamic) epistemic logic can be found in Gasqueta and Schwarzen-truber [16], Pacuit and Parikh [23], and Parikh, Moss and Steinsvold [24]. But all this remains very close to the general frameworks of DEL and LCC. A more difficult challenge would be to model limitations of perception. Any sophisticated account of perceptive ability will at some point have to address the issue of modelling minimum perceivable difference (MPD). This relation is hard to handle in an epistemic framework, for it is not transitive. To tackle this, one need to get around the ‘paradox of the heap.’ See [28] for an early attempt in this area. Bonnay and Egré show in [7] how imprecise knowledge can be represented with non-transitive non-euclidean K45 models. Exploring the connection between the present approach and theirs is future work.

The proposal of the present paper already makes a distinction between actual perception and capability of perception, and should be compared with the logic of sensors of [11], while the connection with the logic of perceptibilia [3] is also worth pursuing. It may be possible to represent uncertainty in perception in terms of probabilities; this would give an obvious link to [5].

A more radically different approach to uncertainty or vagueness in perception is taken in the paper by Dunin-Kępicz and Andrzej Szalas elsewhere in this volume, which is phrased in terms of rough sets.

Finally, awareness of basic propositions allows for the modelling of *communication channels* in DEL: if i is aware that j can perceive changes in p , then p is a communication channel between i and j . Establishing a communication channel can now be done by means of $K_i\mathbb{P}_j^p$. If this is transmitted from i to j , then i can henceforth communicate with j by changes in the value of p .

Acknowledgements

I am grateful to the Netherlands Institute for Advanced Studies (NIAS) for hosting the GASS project (2005-2006), during which the draft version of this paper was written. Thanks to Rineke Verbrugge and to two anonymous reviewers for very helpful feedback. It is my pleasure to acknowledge Floor Sietsma’s help with the correction of the final version.

References

1. AUSTIN, J. *Sense and Sensibilia*. Oxford University Press, 1962.
2. BALTAG, A., MOSS, L., AND SOLECKI, S. The logic of public announcements, common knowledge, and private suspicions. Tech. Rep. SEN-R9922, CWI, Amsterdam, 1999. With many updates.
3. BELL, J. Continuity and the logic of perception. *Transcendent Philosophy* 1, 2 (2000), 1–7.
4. BENTHEM, J. V. Open problems in logic and games. Tech. Rep. PP-2005-06, ILLC, Amsterdam, 2005.
5. BENTHEM, J. V., GERBRANDY, J., AND KOOI, B. Dynamic update with probabilities. Manuscript, ILLC, Amsterdam, 2006.

6. BENTHEM, J. v., VAN EIJCK, J., AND KOOI, B. Logics of communication and change. *Information and Computation* 204, 11 (2006), 1620–1662.
7. BONNAY, D., AND EGRÉ, P. Inexact knowledge with introspection. *Journal of Philosophical Logic* 38, 2 (2009), 179–228.
8. CLARK, H., AND MARSHALL, C. Reference diaries. Tech. rep., Stanford University, 1978.
9. CLARK, H., AND MARSHALL, C. Definite reference and mutual knowledge. In *Elements of discourse understanding*, A. Joshi, B. Webber, and I. Sag, Eds. Cambridge University Press, 1981, pp. 10–63.
10. DAVIDSON, D. Mental events. In *Essays on Actions and Events*. Oxford University Press, 1982.
11. DEL VAL, A., II, P. M.-R., AND SHOHAM, Y. Qualitative reasoning about perception and belief. In *IJCAI (1)* (1997), pp. 508–513.
12. DRETSKE, F. *Seeing and Knowing*. University of Chicago Press, 1969.
13. DRETSKE, F. *Knowledge and The Flow of Information*. Bradford/MIT Press, Cambridge, MA, 1981.
14. EIJCK, J. v. Guarded actions. Tech. Rep. SEN-E0425, CWI, Amsterdam, December 2004. Available from <http://repository.cwi.nl/>.
15. EIJCK, J. v. DEMO — a demo of epistemic modelling. In *Interactive Logic — Proceedings of the 7th Augustus de Morgan Workshop*, J. v. Benthem, D. Gabbay, and B. Löwe, Eds., no. 1 in Texts in Logic and Games. Amsterdam University Press, 2007, pp. 305–363.
16. GASQUET, O., AND SCHWARZENTRUBER, F. Knowledge in lineland (extended abstract). In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)* (2010), W. v. d. Hoek, Kaminka, Lespérance, Luck, and Sen, Eds., pp. 1537–1538.
17. GETTIER, E. Is justified true belief knowledge? *Analysis* 23 (1963), 121–123.
18. HALPERN, J., AND MOSES, Y. Knowledge and common knowledge in a distributed environment. In *Proceedings of the 3rd ACM Symposium on Principles of Distributed Computing (PODS)* (1984), pp. 50–61. A newer version appeared in the *Journal of the ACM*, vol. 37:3, 1990, pp. 549–587.
19. HOARE, C. An axiomatic basis for computer programming. *Communications of the ACM* 12, 10 (1969), 567–580, 583.
20. KOOI, B. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics* 16, 2 (2007).
21. LEWIS, D. *Convention: A Philosophical Study*. Harvard University Press, 1969.
22. MCCARTHY, J. Formalization of two puzzles involving knowledge. In *Formalizing Common Sense: Papers by John McCarthy*, V. Lifschitz, Ed. Ablex Publishing Corporation, Norwood, New Jersey, 1990. Original manuscript dated 1978–1981.
23. PACUIT, E., AND PARIKH, R. Reasoning about communication graphs. In *Interactive Logic — Proceedings of the 7th Augustus de Morgan Workshop*, J. v. Benthem, D. Gabbay, and B. Löwe, Eds., no. 1 in Texts in Logic and Games. Amsterdam University Press, Amsterdam, 2007, pp. 135–157.
24. PARIKH, R., MOSS, L., AND STEINSVOLD, C. Topology and epistemic logic. In *Handbook of Spatial Logics*, M. Aiello, I. Pratt-Hartmann, and J. van Benthem, Eds. Springer, 2007, ch. 6.
25. PLAZA, J. A. Logics of public communications. In *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems* (1989), M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, Eds., pp. 201–216.
26. ROCK, I. *The Logic of Perception*. MIT Press, 1985.
27. SCHIFFER, S. *Meaning*. Oxford University Press, 1972.
28. SHOHAM, Y., AND DEL VAL, A. A logic for perception and belief. Tech. Rep. CS-TR-91-1391, Stanford University, 1991.