Awareness and forgetting of facts and agents

Hans van Ditmarsch
Fac. Philosophy, Dep. Logic
University of Sevilla
Sevilla, Spain
Email: hvd@us.es

Tim French
Computer Science and Software Engineering
University of Western Australia
Perth, Australia
Email: tim@csse.uwa.edu.au

***Abstract*—We propose various logical semantics for change of awareness. The setting is that of multiple agents that may become aware of facts or other agents, or forget about them. We model these dynamics by quantifying over propositional variables and agent variables, in a multi-agent epistemic language with awareness operators, employing a notion of bisimulation with a clause for 'same awareness'. The quantification is over all different ways in which an agent can become aware (or forget).**

***Keywords*-awareness; knowledge; multi-agent systems**

I. Introduction

Becoming aware of facts: When modelling uncertainty in a multi-agent system you only pay attention to the relevant facts. You may then become aware of more facts, and integrate those in the system: a refinement operation. It is not obvious how to do this in epistemic logic! For example, when an agent i is uncertain of the value of some fact (atom / propositional variable) p, a two-state structure suffices to represent that uncertainty where p is true in one state and p is false in another state, namely as follows:

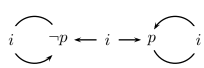

It may now occur that subsequent information about the agent's uncertainty comes to light. Apart from fact p, another fact q is also relevant—the modeller and curiously enough also the agent *become aware of* q. Now there are *many* ways in which this may happen, four of those are pictured below:

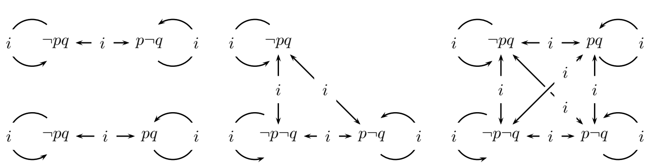

The four models have in common that ignorance about p is *still* the case after the transition. But they are all different in what has become known about q, sometimes the q of which i is aware is known, and sometimes not; and two, three, or four valuations for p and q may be considered possible.[1] We would like to have a logical operation with an existential and a dynamic modal flavour where any of these transitions are natural. But here we have a problem. Fact q does not drop out of thin air but should somehow already have a presence in the language (mathematical or logical) with which we describe the initial two state structure. This figure is informal in that no value for q, or other treatment of that fact, is given in the two states. For example, q may be true in both states. That choice would make the transition to the models wherein it is no longer known very unsatisfactory: the agent would initially know that q—we should apparently ignore that but how?—and after becoming aware of q it no longer knows q. But assuming that q is initially unknown, by giving it different values in the different states, is similarly unsatisfactory: in another of the resulting models agent i has become aware of q and *knows* q. We wish not to commit to either knowledge or ignorance of q, both should be allowed as ways to become aware of q. Other standard logical approaches fail too. Many-valued logic does not work. If we give q initially the value 'unknown', then that is compatible with q becoming known and with q becoming unknown. But we cannot get the increased complexity of three, or four (or more!) epistemically distinct states.

The solution is to regard the initial values of q as 'don't care': something the agent is unaware of. The transition to (e.g.) the model with three states can then be visualized as in Figure 1—what the agent is unaware of is between parentheses, the other transition is explained later. We propose

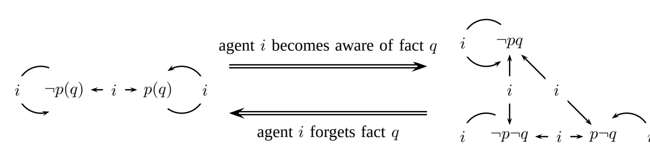

Figure 1. Agent i becomes aware of or forgets fact q

[1] In the single agent case, given this specific example where access is an equivalence, the maximum *factual* uncertainty between four valuations is also the maximum *epistemic* uncertainty that can result; but for for weaker structural conditions, or in the presence of more than one agent, any finite number of epistemically distinct states can result by an agent becoming aware of a single fact q in a two-state model as above.

a logic (i) in which the value of q is initially *irrelevant*, i.e., something the agent is *unaware* of; (ii) in which the depicted transition makes the agent aware of q, wherein any increased structure expressing uncertainty about q is allowed; and (iii) in which the agent is then aware of q.

Forgetting facts: Apart from becoming *aware* of a fact, the agent may become *unaware* of a fact. This can be for voluntary and involuntary reasons: for the purpose of abstraction, to focus computational resources on 'more relevant' facts, by gradually things slipping from the mind, or because information is received that all beliefs about that fact are unreliable. Becoming unaware can also be modelled by a bisimulation quantification, and we have pictured this as well in Figure 1.

Becoming aware of other agents: And apart from becoming aware of a *fact*, an agent can also become aware of another *agent* in the system, and of the uncertainties of that agent about facts (and about other agents, including the observing agent). In other words, not all agents may be visible to a given agent at a given moment. We provide a similar operation for 'the agent becomes aware of agent i', including its dual where visible agents slip into oblivion: 'the agent becomes unware of (forgets) agent i'. This employs bisimulation quantification over an *agent* variable—an idea originally proposed in [1] as far as we know, but so far not employed in dynamic epistemics.

Figure 2 gives an example. The initial state of information is a slight adjustment of the previous example, as there is now another agent j as well, of which agent i is unaware and who is knowledgable about p. (But who is not introspective, as he is not aware of himself.) The figure depicts how agent i can become aware of agent j, and also the dual option of forgetting. The unaware atoms and agents are in parentheses. An arrow with two agent variables stands for two arrows.

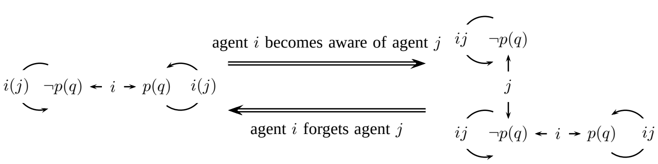

Figure 2. Agent i becomes aware of and forgets about other agent j. On the right, if agent j knows that p is false, he is uncertain if agent i knows that.

So far, the examples addressed change of awareness that is the same for all states and for all agents. We call this *public global awareness*. But awareness change can also be different for each agent (*individual global awareness*), and even in each state (*individual local awareness*, the common assumption in works like [2]). We present natural generalizations of our ideas to those settings.

Implicit and explicit knowledge: We see 'becoming aware' as a structural refinement, where our starting point is an initial minimal structure. 'Unaware' means 'don't care', and we are not interested in the structure of unawareness. This has as unintended consequence that the usual distinction between implicit and explicit knowledge is meaningless in our setting. Explicit knowledge is implicit knowledge of concepts of which you are *aware*. In the initial two-state structure in Figure 1 agent i implicitly knows q. After becoming aware of q, he has become explicitly ignorant about q! In other words, becoming aware does not mean that implicit knowledge will become explicit, an intuitively appealing requirement for awareness change. (Logics handling that properly, instead, have been omitted from this extended abstract.)

Comparison to standard research on awareness: Static awareness in our approach is according to the *semantically* flavoured proposals by [2], [3], [4]—the extraordinarily rich and exceptionally well-written [2] certainly inspired us along our chosen path. Our work is not related to *syntactically* flavoured proposals for awareness that model limited rationality of agents, wherein the set of formulas of which an agent is aware may not be deductively closed (as it can be simply *any* given set of formulas); this other approach is also pursued in [2] and in recent work like [5], [6]. *Dynamic* awareness is modelled by a bisimulation quantification on structures incorporating awareness, over the variable expressing the newly relevant fact. For bisimulation quantification see [7], [8], this concerns a further generalization of propositional quantification à la Fine [9]. We think that our ideas on dynamic awareness are novel. We have studied some of the recent literature on the topic [4], [10], [11], [12]. Only Hill's [11] explicitly addresses the *dynamics* of awareness (but there, the result of becoming aware of a new fact *must* be ignorance about that fact).

Unlike explicit logical dynamics, the *comparison* of different levels of awareness in semantic structures is quite standard in the literature. Going from one to another level is then merely not expressed in the logical language. It is remarkable that the relation between the different static ways of unawareness in Heifetz et al.'s [4] can just as well be described in terms of bisimulation quantification. For example, take $S_{\{p\}}$ and $S_{\{p,q\}}$ in their Figure 2 [4, p.86]: $S_{\{p,q\}}$ sums up all different ways in which an agent only aware of p, as in $S_{\{p\}}$, can become aware of q (and our running example, above, can indeed be found among those). The notion of bisimulation seems a more succinct, technical tool to express the same. It inspired our research to see this correspondence with the acclaimed [4].

Awareness and information: The logics for change of awareness combine well with logics for change of knowledge [13]. A typical example is the announcement of a fact p of which the agents were unaware. The announcer *addresses an issue*, the truth about p, simultaneously with *revealing the truth* about that issue. A surprising result is that arbitrarily complex informational change (private announcements, suspicions, any action model execution [14]) can be seen as the *public announcement of a true fact of which the agents were unaware* [15].

Overview of content: In Section II we introduce epistemic awareness models wherein knowledge and awareness are encoded. An appropriately expanded notion of bisimulation is also introduced there (apart from atoms, back, and forth, there is a fourth clause involving awareness). In Section III we present the Logic of Public Global Awareness (LPGA) and the language $\mathcal{L}_0$: the awareness of facts and other agents at a given moment is the same for all agents. In Section IV we present the Logic of Individual Global Awareness (LIGA) and the language $\mathcal{L}$: the level of awareness can vary between agents, but is the same in all states. In Section V we summarily present the Logic of Individual Local Awareness (LILA) (also based on $\mathcal{L}$), wherein for each agent and each state the level of awareness may vary.

II. Structures

Given are a countably infinite set of propositional variables (facts) P and a countably infinite set of agents N. As we are also modelling 'becoming aware of an agent', any finite number of agents would be insufficient: whatever the finite number of agents in one's company, someone else can always turn up at any stage! The sets P and N are disjoint. The union $P \cup N$ is called the set of *concepts*. Propositional variables are named p, q, r, possibly indexed or quoted, and agent variables are named i, j, k, possibly indexed or quoted. For any set X, write $X + x$ for $X \cup \{x\}$ and write $X - x$ for $X \setminus \{x\}$. Write $\overline{Y}$ for $X \setminus Y$ and similarly $\overline{x}$ for $X - x$.

Epistemic awareness model: An *epistemic awareness model* $M = (S, R, \mathcal{A}, V)$ for N and P consists of a *domain* S of (factual) *states* (or 'worlds'), an *accessibility function* $R : N \to \mathcal{P}(S \times S)$, an *awareness function* $\mathcal{A} : N \to S \to \mathcal{P}(P \cup N)$ and a *valuation function* $V : P \to \mathcal{P}(S)$. For $R(i)$ we write R_i and for $\mathcal{A}(i)$ we write $\mathcal{A}_i$; accessibility function R can be seen as a set of *accessibility relations* R_i, and V as a set of *valuations* $V(p)$. A pointed epistemic awareness model (M, s) is an *epistemic awareness state*.

Given an agent i and a state s, a fact in $\mathcal{A}_i(s)$ (i.e., an element of $\mathcal{A}_i(s) \cap P$) is called *relevant* (for that agent, given that state), and a fact in $P \setminus \mathcal{A}_i(s)$ is called *irrelevant*. Similarly, an agent in $\mathcal{A}_i(s)$ is called *visible*, and an agent in $N \setminus \mathcal{A}_i(s)$ is called *invisible*.

The awareness function $\mathcal{A}$ satisfies *public global awareness* iff the value of $\mathcal{A}$ is the same for all agents and for all states. Slightly abusing the (mathematical) language, we then write $\mathcal{A}(S)$ for the set of (globally) relevant facts and visible agents. The awareness function $\mathcal{A}$ satisfies *individual global awareness* iff the awareness is the same in all states, but maybe different between agents. We then write $\mathcal{A}_i(S)$ for the set of (globally) relevant facts and visible agents *for that agent*. If the awareness may be different for all agents and in all states (the usual assumption in the literature) we call it *individual local awareness*.

Epistemic uncertainty over awareness is sometimes ruled out: an agent is supposed to know whether she is aware of a fact or of another agent. This corresponds to the property on epistemic awareness models of *no uncertain awareness* defined as "If $(s,t), (s,u) \in R_i$, then $\mathcal{A}_i(t) = \mathcal{A}_i(u)$." If R_i is an equivalence relation (to interpret knowledge of an agent), then, if 'no uncertain awareness' is satisfied, the partition induced by R_i on the domain is a refinement of the partition induced by $\mathcal{A}_i$.

Given an epistemic awareness model (M, s) with awareness function $\mathcal{A}$, we write $(M, s)^{\mathcal{A}_i + p}$ for the model that is like (M, s) except for its awareness function $\mathcal{A}'$ for which we have $\mathcal{A}'_i(s) = \mathcal{A}_i(s) + p$. When 'no uncertain awareness' holds we assume that p is added for all states in i's equivalence class, when 'individual global awareness' holds we assume that p is added for all states in the model, and for 'public global' conditions we write $(M, s)^{\mathcal{A}+p}$. Similarly for $(M, s)^{\mathcal{A}_i + p}$, and for $(M, s)^{\mathcal{A}_i - p}$ (to be used in a definition for forgetting), etc.

Bisimulation: Let two models $M = (S, R, \mathcal{A}, V)$ and $M' = (S', R', \mathcal{A}', V')$ be given. A non-empty relation $\mathfrak{R} \subseteq S \times S'$ is a *bisimulation*, iff for all $s \in S$ and $s' \in S'$ with $(s, s') \in \mathfrak{R}$:

- atoms $s \in V(p)$ iff $s' \in V'(p)$ for all $p \in P$;
- aware for all $i \in N$, $\mathcal{A}_i(s) = \mathcal{A}'_i(s')$;
- forth for all $i \in N$ and $t \in S$, if $R_i(s,t)$ then there is a $t' \in S'$ such that $R'_i(s', t')$ and $(t, t') \in \mathfrak{R}$;
- back for all $i \in N$ and $t' \in S'$, if $R'_i(s', t')$ then there is a $t \in S$ such that $R_i(s,t)$ and $(t, t') \in \mathfrak{R}$.

We write $(M, s) \leftrightarroweq (M', s')$, iff there is a bisimulation between M and M' linking s and s', and we call (M, s) and (M', s') *bisimilar*. The novelty of our definition is the clause **aware**, that requires that bisimilar states have the same level of awareness.

A bisimulation *except for fact* p satisfies **atoms** for $P - p$, and **aware** to the extent that $\mathcal{A}_i(s) - p = \mathcal{A}_i(s') - p$. The value of p may vary, including uncertainty of the agents about p and about each other's uncertainty, and the agents may have different awareness of fact p. We write $(M, s) \leftrightarroweq_{\overline{p}} (M', s')$, and $\mathfrak{R}|\overline{p}$ (the index / restriction makes the subset explicit for which the bisimulation should hold). A bisimulation *except for agent* i satisfies **back** and **forth** for $N - i$, and **aware** to the extent that $\mathcal{A}_i(s) - i = \mathcal{A}_i(s') - i$. The accessibility relation and awareness for agent i may vary, including uncertainty of other agents about i's uncertainty and awareness. We write $(M, s) \leftrightarroweq_{\overline{i}} (M', s')$, and $\mathfrak{R}|\overline{i}$. The notion of restricted bisimilarity and its notation generalize in the obvious way to more facts and agents, also in combination.

Awareness bisimulation: The notion of bisimulation will be exactly what we need to capture awareness change under conditions of public global and individual global awareness. It is also effective under conditions of individual local awareness, but in that case too restrictive—now we need a coarser notion of structural similarity to cover all different ways in which an agent can become aware of facts or agents.

Consider the following example: in the actual state s agent i is aware of agent j and of fact p, and state t is i-accessible from the actual state. In state t, agent j is aware of p and q. That agent j is also aware of q should leave agent i indifferent, as she was not aware of q in the actual state. Therefore, in case agent i were to become aware of q in state s, she should consider it possible that j is unaware of q in that i-accessible state t. Under conditions of public or individual global awareness this is not a variation we care to consider: if j is aware of q in t, then he is already aware of q in the actual state s. Clearly, we do not want to change the value of atoms of which agents are aware in the actual state.

This sort of similarity is captured in the following notion, named *awareness bisimulation*. If two models are awareness bisimilar, they cannot be distinguished (as we will see) by the fragment of the language of which the agents are aware: therefore, they are described by the same explicit knowledge. In the following, we use the notational abbreviation $\mathcal{A}(s)$ for $\lambda i.\mathcal{A}_i(s)$, and the abbrevation $N(s)$ for the set of all agents of which some agent is aware in state s, i.e., $\{j \in N \mid \text{there is an } k \in N \text{ such that } j \in \mathcal{A}_k(s)\}$.

Let epistemic awareness states $(M, u) = ((S, R, \mathcal{A}, V), u)$ and $(M', u') = ((S', R', \mathcal{A}', V'), u')$ be given. A non-empty relation $\mathfrak{R}^A \subseteq S \times S'$ (where A in $\mathfrak{R}^A$ stands for 'aware') is an *awareness bisimulation* between (M, u) and (M', u'), notation $(M, u) \leftrightarrows^A (M', u')$, iff $(u, u') \in \mathfrak{R}^A$ and $\mathfrak{R}^A = \bigcap_{j \in N(u)} \mathfrak{R}^A_j[\mathcal{A}(u)]$. We continue by defining $\mathfrak{R}^A_j[\mathcal{A}'']$ for any $\mathcal{A}'' : N \to \mathcal{P}(P \cup N)$ (write $\mathcal{A}''_i$ for $\mathcal{A}''(i)$). Let such a $\mathcal{A}''$ be given, $s \in S$, and $s' \in S'$, then $(s, s') \in \mathfrak{R}^A_j[\mathcal{A}'']$ iff:

- atoms $s \in V(p)$ iff $s' \in V'(p)$ for all $p \in \mathcal{A}''_j$;
- aware for all $i \in \mathcal{A}''_j$, $\mathcal{A}_i(s) \cap \mathcal{A}''_j = \mathcal{A}'_i(s') \cap \mathcal{A}''_j$;
- forth for all $i \in \mathcal{A}''_j$ and $t \in S$, if $R_i(s, t)$ then there is a $t' \in S'$ such that $R_i(s', t')$ and $(t, t') \in \mathfrak{R}^A_j[\mathcal{A}'' \cap \mathcal{A}'(t)]$;
- back for all $i \in \mathcal{A}''_j$ and $t' \in S'$, if $R_i(s', t')$ then there is a $t \in S$ such that $R_i(s, t)$ and $(t, t') \in \mathfrak{R}^A_j[\mathcal{A}'' \cap \mathcal{A}'(t)]$.

In the **back** and **forth** clauses, the relation $\mathfrak{R}^A_j[\mathcal{A}'' \cap \mathcal{A}'(t)]$ is inductively assumed to be already defined. Note that $\mathcal{A}'' \cap \mathcal{A}'(t)$ is a function from the set of agents to a possibly smaller subset of facts and agents than in $\mathcal{A}''$ (we overload the $\cap$ notation), and that these functions are downwardly closed; and if $\mathcal{A}''_j = \emptyset$, then $\mathfrak{R}^A_j[\mathcal{A}''] = \emptyset$. In finite multi-$S5$ structures a stable point will be finitely reached in any chain where $\mathcal{A}'' \cap \mathcal{A}'(t) = \mathcal{A}''$. The relation $\mathfrak{R}^A$ is indeed an equivalence (proof omitted).

The relation between bisimulation and awareness bisimulation is as follows. (i) Given *public global awareness*, awareness bisimulation reverts to bisimulation by way of $\mathfrak{R}|\mathcal{A}(S) = \mathfrak{R}^A$. (ii) Given *individual global awareness* we have that $\mathfrak{R}|\mathcal{A}_i(S) = \mathfrak{R}^A_i[\lambda i.\mathcal{A}_i(S)]$, from which directly follows that $\mathfrak{R}^A$ is the intersection of all $\mathfrak{R}|\mathcal{A}_i(S)$ such that some agent is aware of i somewhere: $\bigcap_{i \in N(S)} \mathfrak{R}|\mathcal{A}_i(S) = \mathfrak{R}^A$. (iii) *Otherwise*, $\mathfrak{R}$ is a refinement of $\mathfrak{R}^A$. (Proofs omitted.)

III. Public Global Awareness

Language: Given are a countably infinite set of propositional variables (facts) P, and a countably infinite set of agents N. The language $\mathcal{L}_0$ of public global awareness is defined as

$$\varphi ::= p \mid \varphi \wedge \varphi \mid \neg\varphi \mid K_i\varphi \mid \exists p\varphi \mid \exists i\varphi \mid A\varphi$$

where $p \in P$ and $i \in N$. By notational abbreviation are defined:

$$\begin{array}{lcl} \top & = & \exists p(p \vee \neg p) \\ \dot{K}_i\varphi & = & A\varphi \wedge K_i\varphi \\ \dot{\exists} p\varphi & = & \neg Ap \wedge \exists p(\varphi \wedge Ap) \\ \dot{\exists} i\varphi & = & \neg AK_i\top \wedge \exists i(\varphi \wedge AK_i\top) \\ \underset{\cdot}{\exists} p\varphi & = & Ap \wedge \exists p(\varphi \wedge \neg Ap) \\ \underset{\cdot}{\exists} i\varphi & = & AK_i\top \wedge \exists i(\varphi \wedge \neg AK_i\top) \end{array}$$

Construct $K_i\varphi$ stands for 'agent i implicitly knows φ'. We already pointed out in the introduction that this is a rather tentative phrasing in our setting. Construct $A\varphi$ stands for 'the agents are aware of φ', or rather more strictly, looking ahead to our semantics: 'the *visible* agents are *collectively* aware of φ'. The curiously non-standard definition of $\top$ is to make explicit knowledge of truth possible even if all facts are irrelevant. The meaning of the bisimulation quantifications $\exists p\varphi$ and $\exists i\varphi$ is less intuitive than that of their counterparts that were introduced by abbreviation:

$\dot{K}_i\varphi$	agent i (explicitly) knows φ
$\dot{\exists} p\varphi$	after the agents become aware of fact p, φ
$\dot{\exists} i\varphi$	after the agents become aware of agent i, φ
$\underset{\cdot}{\exists} p\varphi$	after the agents forget fact p, φ
$\underset{\cdot}{\exists} i\varphi$	after the agents forget agent i, φ

Let us explain one of these abbrevations. Explicit awareness $\dot{\exists} p\varphi$ is defined as $\neg Ap \wedge \exists p(\varphi \wedge Ap)$, which says that the agents are currently not aware of fact p, and there is a way to vary the valuation *and* the awareness of the p, such that afterwards φ is true and the agents are aware of fact p. We have to distinguish universal from existential readings of becoming aware: $\exists p$ says that '*there is a way to become aware of* p *after which ...*', but we need $\forall p$ for '*after any way to become aware of* p *...*'.

The semantics of the awareness operator A will be purely syntax-based, namely using the *free variables* of a formula.

These are defined as follows (note that $var(\varphi) \subseteq P \cup N$): $var(p) = \{p\}$, $var(\varphi \wedge \psi) = var(\varphi) \cup var(\psi)$, $var(\neg\varphi) = var(\varphi)$, $var(K_i\varphi) = var(\varphi) + i$, $var(\exists p\varphi) = var(\varphi) - p$, $var(\exists i\varphi) = var(\varphi) - i$, and $var(A\varphi) = var(\varphi)$.

Semantics: Let $M = (S, R, \mathcal{A}, V)$ be given. We remind the reader that the function $\mathcal{A}$ in the case of public global awareness is constant for all agents and for all states, and that we write $\mathcal{A}(S)$ for the set of relevant facts and visible agents. Below, let the models M' have structure $(S', R', \mathcal{A}', V')$.

$(M, s) \models p$	iff	$s \in V(p)$
$(M, s) \models \varphi \wedge \psi$	iff	$(M, s) \models \varphi$ and $(M, s) \models \psi$
$(M, s) \models \neg\varphi$	iff	$(M, s) \not\models \varphi$
$(M, s) \models K_i\varphi$	iff	for all t: $(s, t) \in R_i \Rightarrow (M, t) \models \varphi$
$(M, s) \models \exists p\varphi$	iff	there is a (M', s') such that $(M, s) \leftrightarrows_{\overline{p}} (M', s')$ & $(M', s') \models \varphi$
$(M, s) \models \exists i\varphi$	iff	there is a (M', s') such that $(M, s) \leftrightarrows_{\overline{i}} (M', s')$ & $(M', s') \models \varphi$
$(M, s) \models A\varphi$	iff	$var(\varphi) \subseteq \mathcal{A}(S)$

The set of validities (and the logic) is called $LPGA$ (Logic of Public Global Awareness). The semantics of all these operators, including $K_i\varphi$, $\exists p\varphi$, and $A\varphi$, is nearly perfectly standard—except for the additional bisimulation requirement with respect to the awareness function.

We now explain the awareness clauses in becoming aware of and forgetting about other agents. Consider $\dot{\exists} i\varphi$ (the agents forget about agent i) which stands for $AK_i\top \wedge \exists i(\varphi \wedge \neg AK_i\top)$. The requirement $AK_i\top$ states that the agents must currently be aware of that agent i for the forgetting to be able to take place. Agent label i has to occur somewhere in the formula bound by A. This cannot be, e.g., the formula AK_ip: if the agents are unaware of p, AK_ip would be false, even if they are aware of i! By choosing $AK_i\top$ this is avoided: using the notational abbreviation for $\top$ this stands for $AK_i\exists p(p \vee \neg p)$, and the set of free variables of this formula is $\{i\}$.

Example: Figure 1 models that the agents (namely agent i) become aware of q. Initially, the agent is only aware of p. We can now check in the semantics that all of the following hold throughout the initial model: $Ap, \neg Aq, \dot{\exists} q\dot{K}_i\neg(p \wedge q)$. The two models in the figure are bisimilar except for fact q. In Figure 2 we have that in the initial model, in the (left) state where p is false and relevant and q is true and irrelevant, $\exists j(K_j\neg p \rightarrow \neg K_jK_iK_j\neg p \wedge \neg K_j\neg K_iK_j\neg p)$ is true: after the agents become aware of j, then if that agent knows that p is false he is uncertain whether agent i knows that.

Axiomatization and theory: We have not axiomatized the logic yet. Among the more obvious principles are $\forall p\ p \leftrightarrow \bot$ and $\forall p\ q \leftrightarrow q$, for $q \neq p$. Bisimulation quantified logics are known to be hard to axiomatize. We think it is feasible to show that the logic $LPGA$ is decidable, via a translation into the μ-calculus, employing the techniques of [1].

IV. Individual global awareness

Language: The difference with the language $\mathcal{L}_0$ for public global awareness is that the operators $\exists$ and A are now relative to an agent. The language for individual awareness serves both the logic of individual global awareness $LIGA$ and the logic of individual local awareness $LILA$, to be introduced in the next section.

The language $\mathcal{L}$ of individual awareness is defined as

$$\varphi ::= p \mid \varphi \wedge \varphi \mid \neg\varphi \mid K_i\varphi \mid \exists_i p\varphi \mid \exists_i i\varphi \mid A_i\varphi$$

where $p \in P$ and $i \in N$. The abbreviations for explicit knowledge and awareness now are

$$\begin{array}{rcl} \dot{K}_i\varphi & = & A_i\varphi \wedge K_i\varphi \\ \dot{\exists}_i p\varphi & = & \neg A_ip \wedge \exists_i p(\varphi \wedge A_i\varphi) \\ \dot{\exists}_i j\varphi & = & \neg A_iK_j\top \wedge \exists_i j(\varphi \wedge A_iK_j\top) \\ \dot{\exists}_i p\varphi & = & A_ip \wedge \exists_i p(\varphi \wedge \neg A_ip) \\ \dot{\exists}_i j\varphi & = & A_iK_j\top \wedge \exists_i j(\varphi \wedge \neg A_iK_j\top) \end{array}$$

Formula $\dot{\exists}_i p\varphi$ stands for 'after (some way in which) agent i becomes aware of atom p, φ', and $A_i\varphi$ stands for 'agent i is aware of φ', etc. The free variables of a formula in $\mathcal{L}$ are defined as before with different clauses $var(\exists_i p\varphi) = var(\varphi) + i - p$, $var(\exists_i j\varphi) = var(\varphi) + i - j$, and $var(A_i\varphi) = var(\varphi) + i$.

Semantics: For each agent, its awareness of other agents and facts is the same throughout a model, and as already introduced we use shorthand $\mathcal{A}_i(S)$ for the subset of the relevant facts and visible agents *for agent* i. One could consider further constraints such as *self-awareness*, 'for all $i \in N$, $i \in \mathcal{A}_i(S)$', but we do not require that, as we prefer to keep our approach as general as possible. The crucial clauses are as follows (forgetting is similar).

$(M, s) \models \exists_i p\varphi$	iff	there is a (M', s') s.t. $(M, s) \leftrightarrows_{\overline{i}} (M', s')$, $(M, s) \leftrightarrows_{\overline{p}} (M', s')$ and $(M', s') \models \varphi$
$(M, s) \models \exists_i j\varphi$	iff	there is a (M', s') s.t. $(M, s) \leftrightarrows_{\overline{i}} (M', s')$, $(M, s) \leftrightarrows_{\overline{j}} (M', s')$ and $(M', s') \models \varphi$
$(M, s) \models A_i\varphi$	iff	$var(\varphi) \subseteq \mathcal{A}_i(S)$

The set of validities (and the logic) is called $LIGA$ (Logic of Individual Global Awareness). The semantics for $\dot{\exists}_i p\varphi$ amounts to the requirement that in (M, s) (there is a way such that), after agent i becomes aware of p, φ is true, if and only if φ *remains* true in (M', s') for all agents except i and for all atoms except p. Note that this is a stronger requirement than $(M, s) \leftrightarrows_{\overline{j}, \overline{p}} (M', s')$! Suppose the latter were the case, and let $q \neq p$. M' may now differ for agent i in the value of q—and that would be awkward if i were already aware of q... But if it is required that $(M, s) \leftrightarrows_{\overline{i}} (M', s')$ and that $(M, s) \leftrightarrows_{\overline{p}} (M', s')$ this sort of eventuality is ruled out.

There are open questions on the relation between the logics $LPGA$ and $LIGA$. For example, it is unclear if $A\varphi$ (public awareness!) can be expressed in $\mathcal{L}$, as the obvious infinitary conjunction $\bigwedge_{i \in N} A_i\varphi$ is not a formula in the language. As the current agents may always become aware of yet another agent, we cannot restrict the set of all agents to a finite set. We do not know of a logic with our intuitive requirement that the set of agents must be infinite. We have not yet explored the axiomatization of this logic $LIGA$ of individual global awareness—it suffers from the same backdraws for bisimulation quantifiers as $LPGA$.

V. Individual local awareness

We now move to the most complex stage of awareness change. A first observation is that we can keep the *same* language $\mathcal{L}$ and even the *same* semantics for the operators as in $LIGA$. The difference is that it applies to a larger class of models, therefore the change of awareness allowed in the bisimulation variation 'except for variable p' need no longer be global, for all states in the model, but may now be 'almost' local: in the actual state only. Almost: in accordance with the conventions proposed in Section II, 'no uncertain awareness' is an invariant, so, e.g. when the accessibility is an equivalence relation, we only may change the awareness of p in all states of the agent's actual equivalence class (or, therefore, across a union of such classes). So we *have* a logic $LILA$ for individual local awareness change. As also said in Section II, the variation allowed by the part $(M, s) \leftrightarrows_{\overline{p}} (M', s')$ is now too restrictive, and we need to employ more in full the freedom for local variation for *other* agents. This can be expressed with our alternative notion, of *awareness bisimulation*. The corresponding basic construct for becoming aware is $\exists_i^A p\varphi$, with an upper index to distinguish it from the previous $\exists_i p\varphi$, where the A expresses that it is interpreted using $\mathfrak{R}^A$. Its semantics is:

$(M, s) \models \exists_i^A p\varphi$ iff there is a (M', s') such that
$(M, s) \leftrightarrows^A (M', s')$ and $(M', s')^{\mathcal{A}_i + p} \models \varphi$

This says that (there is a way in which) the agent i becomes aware of atom p in the current state if there is a model similar to the current one in all its observable aspects except that fact p is added to the awareness set for that agent in the actual state s. If 'no uncertain awareness' holds, p should *also* be added to the awareness set for that agent in all states accessible for that agent from actual state s, a simple adaptation of the semantics. We now have that (details omitted) *awareness bisimilar structures satisfy the same explicit knowledge*: if $(M, s) \leftrightarrows^A (M', s')$, then $(M, s) \models \dot{K}_i\varphi$ iff $(M, s) \models \dot{K}_i\varphi$ for all $\varphi \in \mathcal{L}$, and we conjecture that *given individual global awareness, $\exists_i p$ and $\exists_i^A p$ are (explicitly) indistinguishable*: if $(M, s) \models \neg A_i p \wedge A_i \exists_i p\varphi$, then $(M, s) \models \exists_i p\varphi$ iff $(M, s) \models \exists_i^A p\varphi$.

Acknowledgment

This paper explores the approach suggested in Section 6 of [15], entitled 'Becoming aware of factual information'. We kindly acknowledge interactions during LOFT 2008 with Oliver Board, Jan Broersen, Brian Hill, Joris Hulstijn, and Giacomo Sillari. In particular we owe Jan Broersen for sharing with us his similar ideas on becoming aware of other agents, and for stimulating discussions. We thank the WLIAMAS anonymous referees. Hans van Ditmarsch is also affiliated to the University of Otago, New Zealand.

References

[1] G. d'Agostino and M. Hollenberg, "Logical questions concerning the μ-calculus: Interpolation, Lyndon and Los-Tarski," *Journal of Symbolic Logic*, vol. 65, no. 1, pp. 310–332, 2000.

[2] R. Fagin and J. Halpern, "Belief, awareness, and limited reasoning," *Artificial Intelligence*, vol. 34, no. 1, pp. 39–76, 1988.

[3] S. Modica and A. Rustichini, "Unawareness and partitional information structures," *Games and Economic Behavior*, vol. 27, pp. 265–298, 1999.

[4] A. Heifetz, M. Meier, and B. Schipper, "Interactive unawareness," *Journal of Economic Theory*, vol. 130, pp. 78–94, 2006.

[5] T. Agotnes and N. Alechina, "Full and relative awareness: a decidable logic for reasoning about knowledge of unawareness," in *Proceedings of TARK XI*. New York, NY, USA: ACM, 2007, pp. 6–14.

[6] F. Velazquez-Quesada, "Inference and update," University of Amsterdam, Tech. Rep., 2008, iLLC Research Report PP-2008-48.

[7] A. Visser, "Bisimulations, model descriptions and propositional quantifiers," 1996, logic Group Preprint Series 161, Department of Philosophy, Utrecht University.

[8] M. Hollenberg, "Logic and bisimulation," Ph.D. dissertation, University of Utrecht, 1998.

[9] K. Fine, "Propositional quantifiers in modal logic," *Theoria*, vol. 36(3), pp. 336–346, 1970.

[10] G. Sillari, "Quantified logic of awareness and impossible possible worlds," *Review of Symbolic Logic*, vol. 1(4), pp. 514–529, 2008.

[11] B. Hill, "Awareness dynamics," *Journal of Philosophical Logic*, 2009, to appear.

[12] O. Board and K.-S. Chung, "Object-based unawareness," 2009, unpublished manuscript, prior version presented at LOFT 2006.

[13] H. van Ditmarsch, W. van der Hoek, and B. Kooi, *Dynamic Epistemic Logic*, ser. Synthese Library. Springer, 2007, vol. 337.

[14] A. Baltag, L. Moss, and S. Solecki, "The logic of public announcements, common knowledge, and private suspicions," in *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, I. Gilboa, Ed., 1998, pp. 43–56.

[15] H. van Ditmarsch and T. French, "Simulation and information," in *Proceedings of KRAMAS, Sydney*, ser. LNAI, J. Broersen and J.-J. Meyer, Eds. Springer, 2009, to appear. Also presented at LOFT 2008, Amsterdam.