

Introspective forgetting

Hans van Ditmarsch*

Computing Science, University of Aberdeen, UK & Computer Science, University of Otago, New Zealand, hans@cs.otago.ac.nz

Andreas Herzig

Institut de Recherche en Informatique, Université Paul Sabatier, France, herzig@irit.fr

Jérôme Lang

LAMSADE, Université Paris-Dauphine, France, lang@lamsade.dauphine.fr

Pierre Marquis

Centre de Recherche en Informatique de Lens, Université d'Artois, France, marquis@cril.univ-artois.fr

April 6, 2009

Abstract. We model the forgetting of propositional variables in a modal logical context where agents become ignorant and are aware of each others' or their own resulting ignorance. The resulting logic is sound and complete. It can be compared to variable-forgetting as abstraction from information, wherein agents become unaware of certain variables: by employing elementary results for bisimulation, it follows that beliefs not involving the forgotten atom(s) remain true.

Keywords: modal logic, forgetting, abstraction, action logic, belief change

1. There are different ways of forgetting

Becoming unaware In the movie 'Men in Black', Will Smith makes you forget knowledge of extraterrestrials by flashing you with a light in the face. After that, you have forgotten the green ooze flowing out of mock-humans and such: you do not remember that you previously had these experiences. In other words, even though for some specific forgotten fact p , previously Kp or $K\neg p$, the flash victims are not currently aware of their ignorance about p . One solution is to model that the flash victims no longer have p available as a propositional variable in the logical language. This sort of forgetting (amnesia) is dual to awareness, as in 'the agents become aware of a new propositional variable p '.

Becoming ignorant A different sort of forgetting is when you forgot which of two keys fits your office door, because you have been away

* Hans van Ditmarsch is corresponding author. The work for this publication was mainly carried out while he was associated to: Institut de Recherche en Informatique, Université Paul Sabatier, France.



from town for a while. Is it the bigger or the smaller key? This is about forgetting the value of an atomic proposition p —such as “the bigger key fits the door.” Sometime in the past, you knew: Kp or $K\neg p$. But currently, you don’t know: $\neg Kp \wedge \neg K\neg p$. You are embarrassingly aware of your current ignorance: introspection is involved, and we even have $K(\neg Kp \wedge \neg K\neg p)$. We call this *introspective forgetting*.

Modelling introspective forgetting is central to our concerns. We also pay some attention to its relation with forgetting as becoming unaware. There are several further issues within our scope:

Forgetting values Did it ever happen to you that you met a person whose face you recognize but whose name you no longer remember? Surely! Or that you no longer know the pincode of your bankcard? Hopefully not. But such a thing is very conceivable. This sort of forgetting means that you forgot the value of a proposition, or the assignment of two values from different sets of objects to each other. An atomic proposition about your office keys is a feature with two values only, true and false. The (finitely) multiple-valued feature can be modelled by means of a *number* of atomic propositions. Forgetting of such multiple boolean variables is in our approach similar to forgetting a single boolean variable.

Multi-agent versions of forgetting Will Smith only had to flash a whole group once, not each of its members individually. So, in a multi-agent setting some aspects of collectively ‘becoming unaware’ can be modelled. A different but much more familiar phenomenon is that of an individual becoming unaware in a group: “You forgot my birthday, *again!*”

A group version for becoming ignorant should involve prior common knowledge followed by some sort of collective memory loss. This collective introspective character is not always easy to justify—but societal phenomena come to mind like collective denial of the Holocaust, and, for something completely different, collective belief among banks that there is an inexhaustive supply of credit. The case where an *individual* in a group becomes ignorant is more intuitive, because individuals interact with their environment: here you are standing in front of your office door again, now in company of four freshmen students, “Ohmigod, I forgot *again* which is my office key!”

Some issues are in our scope, but are left to further research:

Remembering prior knowledge Having forgotten which is your office key, you may also remember that you *knew* which key it was. You just

forgot. Previously Kp or $K\neg p$, and only now $\neg Kp$ and $\neg K\neg p$. But the memory lingers: you now also know that you used to know whether p . (A corresponding formal expression is only introduced in Section 6.)

Forgetting modal formulas I may have forgotten whether you knew about a specific review result for our jointly edited journal issue. In other words, previously $K_{me}K_{you}accept$ or $K_{me}K_{you}\neg accept$ but currently $\neg K_{me}K_{you}accept$ and $\neg K_{me}K_{you}\neg accept$. Some meaningful propositions that can be forgotten are themselves modal.

Forgetting events Forgetting that a certain *event* (action) took place in the past is different from forgetting a *proposition*. It models increased temporal uncertainty instead of increased epistemic uncertainty. The two are related, as epistemic uncertainty can follow from temporal uncertainty. E.g., if you forgot that you told me that p , you have also forgotten the consequences of telling me p . Therefore, you will have forgotten that I know p .

Belief contraction In belief revision, contraction of a set of beliefs consists of removing some of the current beliefs (Alchourrón et al., 1985). This therefore seems related to the forgetting of variables in formulas with explicit belief operators. Our modal logical context does not translate straightaway to belief contraction (see (Seegerberg, 1999) for an elegant approach on which we also base our exposition). Beyond that, there seem to be relations between forgetting and notions known as symmetric contraction and symmetric erasure, and the distinctions between belief revision and belief update (factual change), see e.g. (Gärdenfors, 1992, p.196), (Winslett, 1991), and (Liberatore and Schaerf, 1998). Regrettably we leave these matters to further research.

Finally, some issues are outside our scope:

Explicit awareness In a multi-agent setting, there are logics with explicit awareness operators (Fagin and Halpern, 1988), to describe that Will Smith knows that there are extraterrestrials, and also knows that you are unaware of extraterrestrials and of his knowledge of them. Such logics make a distinction between explicit and implicit knowledge, and the agents may know something implicitly without being aware of it: non-introspective knowledge. These fascinating approaches we consider outside our scope.

Cognitive modelling of forgetting Forgetting tends to be a gradual process. First you know something, then, slowly, it recedes from memory,

i.e. it drops below the level of awareness, and only then you become aware having forgotten it (e.g. when standing in front of your office door, realizing that you forgot which was your office key). This can subsequently even be followed by a process of conscious stepwise recovery of the lost information, by recalling the circumstances under which you still knew. In other words, you finally remember your knowledge. This cannot be modelled in a dynamic epistemic logic wherein agents' beliefs are deductively closed.

2. Motivation

A short history of forgetting in AI Forgetting as abstraction (a.k.a. projection, marginalization, ...), not necessarily in an epistemic setting, goes back a long way. Boole described variable abstraction in (Boole, 1854) wherein he called this the 'elimination of middle terms'. Lin and Reiter propose in '*Forget it!*' (Lin and Reiter, 1994) a way to abstract from ground atoms (that can be equated to propositional variables) in a set of first-order beliefs, employing the notion of *similarity of models for a theory except for such a ground atom*. They leave it open whether such forgetting is the result of an agent consciously updating a knowledge base after having learnt about *factual* change, or whether this is simple erosion of her working memory, purely *epistemic* change. Their work was built upon by Lang, Liberatore and Marquis with their in-depth study on the computational costs of transforming theories by variable forgetting (Lang et al., 2003), or rather the costs of determining the independence of parts of a theory from specific variables. Baral and Zhang joined this battleground for more explicit operators for knowledge and belief with their (Baral and Zhang, 2005), wherein the result of an agent forgetting a variable results in her (explicit) ignorance of that variable's value, and in (Zhang and Zhou, 2008), in progress, Zhang and Zhou make an original and interesting backtrack to the ideas of (Lin and Reiter, 1994) by suggesting *bisimulation invariance except for the forgotten variable*, in order to model forgetting. Forgetting has been generalized to logic programs in (Wang et al., 2005; Zhang et al., 2005; Eiter and Wang, 2006) and to description logics in (Zhao et al., 2007). Forgetting of (abstracting from) actions in planning has been investigated in (Erdem and Ferraris, 2007).

Forgetting propositional variables in propositional logic In (Lang et al., 2003), three ways of propositional variable forgetting in propositional logic are distinguished. Given a (propositional) formula φ and a propositional variable p , we write $Fg(\varphi, p)$ for the result (also of type formula)

of forgetting variable p in φ . Let $\varphi(\psi/p)$ be the replacement of all (possibly zero) occurrences of p in φ by ψ .

1. $Fg(\varphi, p)$ is the result of replacing p in φ by ‘true’ in disjunction with replacing p in φ by ‘false’: $Fg(\varphi, p) =_{\text{def}} \varphi(\top/p) \vee \varphi(\perp/p)$. (Or a preferred logically equivalent formula.)
2. Modulo logical equivalence, $Fg(\varphi, p)$ is the strongest formula not containing p that follows from φ ; i.e., $\varphi \models Fg(\varphi, p)$ and p not in $Fg(\varphi, p)$, and for all χ not containing p , if $\varphi \models \chi$ then $Fg(\varphi, p) \models \chi$.
3. Modulo logical equivalence, $Fg(\varphi, p)$ is the strongest formula that is true in all models that only (may) differ from a φ -model in the valuation of p . In propositional logic, a model is a valuation / interpretation V of propositional variables to true and false.

A result of (Lang et al., 2003) is that the three formulations are equivalent for propositional logic. For typical examples, $Fg(p \wedge q, p) = q$, and $Fg(p \leftrightarrow q, p) = \top$.

Forgetting propositional variables in modal formulas Consider again the formula $p \wedge q$. Such a formula is often seen as representing the beliefs or knowledge of an agent, as in the AI area known as ‘belief revision’ (Alchourrón et al., 1985). We can express this in the logic as $K(p \wedge q)$ or, equivalently, $Kp \wedge Kq$. (We ignore other naming conventions, such as writing $B(p \wedge q)$ —most of our results hold for *any* normal modal operator, although some require seriality of the underlying structures.) We have seen that $Fg(p \wedge q, p) = q$. Now we want to know the outcome of $Fg(Kp \wedge Kq, p)$.

Given $Kp \wedge Kq$, what should result when the agent forgets p ? This depends. If we interpret forgetting as *becoming unaware*, we would like the result of forgetting p to be Kq —in general, as in the propositional case, all occurrences of p should be absent from the formula resulting from forgetting, and the truth of formulas in which p does not occur should be preserved after forgetting. On the other hand, studies such as Segerberg’s (Segerberg, 1999) that compare (non-modal) belief revision to (modal) dynamic epistemics have absence of a formula in the belief set correspond to explicit ignorance about it in the modal representation. This means that we interpret forgetting as *becoming ignorant*. The result of forgetting p in $Kp \wedge Kq$ should then be $Kq \wedge \neg Kp \wedge \neg K\neg p$ —‘the agent knows q and is uncertain about p ’. Also adding introspection of knowledge, delivers us the required postcondition $Kq \wedge K\neg Kp \wedge K\neg K\neg p$ —‘the agent knows q and is aware of his uncertainty about p ’. Also in this case, we require that the truth of

formulas in which p does not occur should be preserved after forgetting. More fine-tuned choices are conceivable, in particular when comparing these two options to a large corpus of work on belief contraction in AGM-style belief revision, but as mentioned in the introduction we leave that to further research and instead focus here on a logic wherein (i) the truth of formulas in which p does not occur is preserved after forgetting, and wherein (ii) in the case of introspective forgetting, the postcondition $K\neg Kp \wedge K\neg K\neg p$ is realized as well.

A technique that fails Let us apply the three constructions in (Lang et al., 2003) to the forgetting of modal formulas, simultaneously shifting from propositional models (valuations) to Kripke structures. In construction 1, $Fg(\varphi, p)$ is the result of replacing p in φ by ‘true’ in disjunction with replacing p in φ by ‘false’. As p then no longer occurs in the formulas, that cannot be the way to model introspective forgetting. It is also problematic to model unawareness forgetting. This (only) becomes apparent when applying forgetting to ignorance formulas. For example, if the agent is already ignorant of p , we expect unawareness forgetting to result in a state of information where the agent still has consistent beliefs but only trivial ones, such as $\neg K\perp$. But unfortunately, applying construction 1 to $\neg Kp \wedge \neg K\neg p$ we get a contradiction \perp if we also assume a standard property of knowledge (‘iff’ stands for logical equivalence):

$$Fg(\neg Kp \wedge \neg K\neg p, p)$$

is by definition

$$(\neg K\top \wedge \neg K\neg\top) \vee (\neg K\perp \wedge \neg K\neg\perp)$$

iff (assuming the property of belief and knowledge that $K\perp \leftrightarrow \perp$)

$$(\neg\top \wedge \neg\perp) \vee (\neg\perp \wedge \neg\top)$$

iff

$$\perp$$

In construction 2, $Fg(\varphi, p)$ is the strongest formula (modulo logical equivalence) not containing p that follows from φ . This also does not result in introspective forgetting. It appears to apply to awareness forgetting, as e.g. the strongest ψ following from $\neg Kp \wedge \neg K\neg p$ not containing p is $\neg K\perp$. Under conditions of normal belief this is equivalent to \top .

In construction 3, applied to a Kripke model, we look for the strongest formula that is true in all models that only (may) differ from a model for $\neg Kp \wedge \neg K\neg p$ in the valuation of p . Applying this to Kripke models, varying the valuation of p in the model means changing the subset of the domain where p holds. This seems a candidate for introspective

forgetting, as we continue to be allowed to talk about p . But now observe that, whatever the original formula φ was, none of Kp , $K\neg p$, $\neg Kp$, $\neg K\neg p$ satisfies the requirement. For example, $\neg Kp$ does not satisfy it, because one way of changing the valuation of p is making it true everywhere, so that Kp then holds. Etc. This leaves us only with trivial formulas such as $Kp \vee \neg Kp$. As with construction 2, we end up having just (a formula logically equivalent to) $\neg K\perp$, again.

Let us not continue in this direction. Despite incidental results, such as (Herzig et al., 2003) and unpublished work in progress by the authors, no satisfactory solution has been proposed to transform any given modal formula φ into another modal formula $Fg(\varphi, p)$ in which information about p is forgotten (whether of the ‘unawareness’ kind or of the ‘introspective’ kind). This motivated us to model forgetting as an event in a dynamic epistemic logic. First, we focus on introspective forgetting. The relation to unawareness forgetting is discussed later, in Section 4.

Introspective forgetting as a dynamic modal operator Let us model forgetting an atomic proposition p as an *action* or *event* $Fg(p)^n$ (the symbol n will be explained later). We do this in a propositional logic expanded with an epistemic modal operator K and dynamic modal operators such as $[Fg(p)^n]$. Formula $[Fg(p)^n]\varphi$ means that after the agent forgets his knowledge about p , φ is true. We call $[Fg(p)^n]$ a *dynamic* modal operator because it is interpreted as a state transformation, more particularly: by changing an information state that is represented by a pointed Kripke model (M, s) into another information state (M', s') .

A *precondition* for event $Fg(p)^n$ seems prior knowledge of the value of p : $Kp \vee K\neg p$. How can you forget something unless you know it in the first place? To make our approach comparable to variable forgetting in the ‘abstracting-from-information’-sense, we do not require prior knowledge as a precondition for forgetting. That makes forgetting an operation that can always be executed:

\top is the precondition for $Fg(p)^n$.

The *postcondition* for event $Fg(p)^n$ is ignorance of the value of p :

$[Fg(p)^n](\neg Kp \wedge \neg K\neg p)$ is valid.

as well as knowledge of that ignorance:

$[Fg(p)^n]K(\neg Kp \wedge \neg K\neg p)$ is valid.

Seen as a binary relation between pointed Kripke models, the interpretation of our forgetting operator is both functional and total. This

means that we can describe its semantics as:

$$(M, s) \models [Fg(p)^n]\varphi \text{ iff } (M', s') \models \varphi,$$

where (M', s') is the pointed Kripke model resulting from forgetting p in (M, s) . The relation to a (supposed) formula-transforming operation $Fg(\psi, p)$ is as follows:

$$\psi \rightarrow [Fg(p)^n]Fg(\psi, p).$$

In other words: if proposition ψ is true, then after forgetting about p , the proposition is now true that is constructed from ψ relative to the forgetting of p . Put more into terms of construction 2 above, the relation becomes:

- 2'. Modulo logical equivalence, $Fg(\varphi, p)$ is the strongest formula that follows from φ after executing event $Fg(p)^n$.

We emphasize that this construction of $Fg(\varphi, p)$ remains hypothetical and is deferred to further research.

Forgetting or no-forgetting? The principle of no-forgetting or perfect recall states that an agent remembers the sequence of actions leading to her current information state at all times (Fagin et al., 1995, p.128–130) from which follows for the setting of (synchronous) dynamic epistemic logics, that if the agent can distinguish two events from another, she will be able to distinguish them from one another no matter what further events take place in either state. In other words, she will not ‘forget’ a distinction once she can make it, thus the term ‘no forgetting’. When an agent forgets factual information, this is independent from the truth of those facts in the actual world. Therefore, this concerns epistemic change and not ontic change (a.k.a. factual change). Dynamic epistemic logics for epistemic change that satisfy no-forgetting, also satisfy preservation of positive knowledge: once Kp , always Kp . As we are proposing a modelling of ‘forgetting’ in a dynamic epistemic logic, the question therefore comes up how this is compatible with this principle of no-forgetting. This can be explained as follows:

Suppose the agent knows that p . This means that in all states accessible from the actual state, p must be true. After the agent forgets p , some state must be accessible from the actual state where p is false: from the perspective of the agent, this does not involve factual change compared to a previous state where she knew p , she merely cannot distinguish it from a previous state where she knew $\neg p$. But as no $\neg p$ -state was accessible to the agent before forgetting, and as the new situation evolves from the actual state, this makes it inevitable that

the dynamic forgetting operator that we propose must involve factual change somewhere, otherwise we cannot simulate this new perspective of the agent.

Our *simulation* of forgetting therefore involves *changing the value of p in the actual or other states, in a way known to be unobservable by the agent*, and we use a known logic combining ontic and epistemic change for that. This involves executing an event in the actual state that is, from the perspective of the agent, non-deterministic. We therefore continue to adhere to the principle of ‘no-forgetting’ in its traditional sense that the agent continues to be able to make distinctions she already could make before; it is merely that she is unable to perceive which of several events involving factual change takes place. And that results, after all, in the agent forgetting the value of p in the actual state.¹ We continue with the technical realization of this proposal.

3. Introspective forgetting

We present a version of the logic for a single agent and for forgetting a single variable. This is only done to simplify the exposition. All results trivially generalize to multiple agents and to forgetting multiple values simultaneously. Multi-agent forgetting is handled in Section 5 on page 17. Multi-variable forgetting is defined in Definition 9 in Section 3.3, below.

3.1. LANGUAGE, STRUCTURES AND SEMANTICS

Given is a set P of propositional variables.

Definition 1 (Language and structures) Our language \mathcal{L} is

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K\varphi \mid [Fg(p)^j]\varphi$$

where $p \in P$, $j = 0, 1, n$, and our structures are *pointed Kripke models* $((S, R, V), s)$, with $R \subseteq (S \times S)$, $V : P \rightarrow \mathcal{P}(S)$, and $s \in S$. \dashv

If $P' \subseteq P$, then $\mathcal{L}(P')$ is the language restricted to P' . The diamond versions of our modal operators are defined as $\hat{K}\varphi =_{\text{def}} \neg K\neg\varphi$ and $\langle Fg(p)^j \rangle\varphi =_{\text{def}} \neg[Fg(p)^j]\neg\varphi$. For $[Fg(p)^n]\varphi$ we read “After the agent

¹ As an aside, note that our solution to make positive knowledge disappear is different from how belief revision is modelled in dynamic epistemic (doxastic) logic. Prior belief in p that is revised with $\neg p$ and results in belief in $\neg p$ is standardly modelled by considering this a ‘soft’ or defeasible form of belief, i.e., not knowledge, and implemented by changing a preference relation between states (van Ditmarsch, 2005; Baltag and Smets, 2006).

forgets that p , φ holds.” The other two dynamic modal operations, $[Fg(p)^0]\varphi$ and $[Fg(p)^1]\varphi$, only play a technical role; still, we can paraphrase them as follows: $[Fg(p)^0]\varphi$ expresses ‘after making p false, φ ’, and $[Fg(p)^1]\varphi$ expresses ‘after making p true, φ ’. The structures are typically $S5$ to model knowledge and $KD45$ to model belief—but this is not a requirement. We employ a notion of structural similarity between Kripke models that guarantees their logical equivalence (in other words: they describe the same set of beliefs), namely *bisimilarity*; for a definition see (Blackburn *et al.*, 2001).

Definition 2 (Semantics) Assume an epistemic model $M = (S, R, V)$.

$$\begin{array}{ll}
(M, s) \models p & \text{iff } s \in V(p) \\
(M, s) \models \neg\varphi & \text{iff } (M, s) \not\models \varphi \\
(M, s) \models \varphi \wedge \psi & \text{iff } (M, s) \models \varphi \text{ and } (M, s) \models \psi \\
(M, s) \models K\varphi & \text{iff for all } t \in S : (s, t) \in R \text{ implies } (M, t) \models \varphi \\
(M, s) \models [Fg(p)^j]\varphi & \text{iff } (M', (s, j)) \models \varphi \quad \text{for } j = 0, 1, n
\end{array}$$

where $M' = (S', R', V')$ such that: $S' = S \times \{0, 1, n\}$; $((s, i), (t, j)) \in R'$ iff $(s, t) \in R$ and $i, j \in \{0, 1, n\}$; $V'(p) = \{(s, 1) \mid s \in S\} \cup (V(p) \times \{n\})$ and for $q \neq p$, $V'(q) = V(q) \times \{0, 1, n\}$. The set of validities is called *FG*. \dashv

The model M' defined here is in fact the restricted modal product of the Kripke model M and of what is known as an action model or *event model* (Baltag *et al.*, 1998). This event model is of the particular form that also allows factual change—summarily mentioned in (Baltag *et al.*, 1998), but worked out in detail in (van Benthem *et al.*, 2006; van Ditmarsch and Kooi, 2008). See the Appendix for details.

The construction of M' amounts to taking three copies of the model M , making p true everywhere in the first (the 1-states), doing nothing in the second (the n -states, where n stands for nothing), making p false everywhere in the third (the 0-states), and making corresponding states indistinguishable for the agent, i.e., given $s \in S$, $(s, 0)$, $(s, 1)$ and (s, n) are indistinguishable for the agent.

Example Assume the standard $S5$ notion of knowledge. Suppose the agent knows p but does not know q (and where in fact q is true), and where the agent forgets that p . The execution of $Fg(p)^n$ is pictured as follows. In the resulting Kripke model, the agent no longer knows p , and remains uncertain about q . We visualize $S5$ models by linking states that are indistinguishable for an agent. Reflexivity and transitivity are assumed. The actual state is underlined.

$$\begin{array}{ccc}
pq & \longrightarrow & p\neg q \\
\Downarrow & & \Downarrow \\
Fg(p)^n & \Longrightarrow & \begin{array}{ccc}
\neg pq & \longrightarrow & \neg p\neg q \\
| & & | \\
pq & \longrightarrow & p\neg q \\
| & & | \\
pq & \longrightarrow & p\neg q
\end{array}
\end{array}$$

3.2. AXIOMATIZATION

Definition 3 contains a complete axiomatization **FG** for the logic FG . To obtain the axiomatization, we used that the three dynamic operators $[Fg(p)^j]$ stand for the execution of the different points of an event model, so that we can apply the reduction axioms for event models in (Baltag et al., 1998) and (van Ditmarsch and Kooi, 2008). (See the Appendix for details.) Completeness follows immediately, as the logic **FG** is merely event model logic for one particular event. If one adds **KD45** or **S5** axioms to the logic, one gets the axiomatization for forgetting given normal belief or knowledge conditions, respectively.

Definition 3 (Axiomatization FG) The axiomatization **FG** consists of all rules and axioms for the system **K**, necessitation for the three dynamic modal operators for Fg , plus the following axioms involving interaction between Fg and other logical operators:

$$\begin{array}{llll}
[Fg(p)^n]p & \leftrightarrow & p & \\
[Fg(p)^0]p & \leftrightarrow & \perp & \\
[Fg(p)^1]p & \leftrightarrow & \top & \\
[Fg(p)^j]q & \leftrightarrow & q & \text{for } q \neq p \\
[Fg(p)^j]\neg\varphi & \leftrightarrow & \neg[Fg(p)^j]\varphi & \\
[Fg(p)^j](\varphi \wedge \psi) & \leftrightarrow & [Fg(p)^j]\varphi \wedge [Fg(p)^j]\psi & \\
[Fg(p)^j]K\varphi & \leftrightarrow & \bigwedge_{i=0,1,n} K[Fg(p)^i]\varphi & \dashv
\end{array}$$

Theorem 4 Axiomatization **FG** is sound and complete. \dashv

Let us explain the meaning of the axioms. Axiom $[Fg(p)^n]p \leftrightarrow p$ says that the truth value of p does not change as a result of forgetting. On the other hand, $[Fg(p)^0]p \leftrightarrow \perp$ says that it is not the case that p is true after the event ‘make p false’. Etc. Axiom $[Fg(p)^j]q \leftrightarrow q$ says that the truth of atoms other than p remains unchanged. The remaining axioms express that forgetting commutes with negation and distributes over conjunction, and, finally, $[Fg(p)^j]K\varphi \leftrightarrow \bigwedge_{i=0,1,n} K[Fg(p)^i]\varphi$ says that the agent knows φ after forgetting p , if and only if she knows that after making p false, making p true, or not changing the value of p , φ is true.

3.3. SEMANTIC RESULTS

The axioms can be combined to derive a number of useful principles of forgetting.

Proposition 5 (Preservation of factual information)

Let ψ be a boolean formula. Then $\psi \leftrightarrow [Fg(p)^n]\psi$ is valid. ⊢

Proof. By induction on booleans. The case for p is the axiom above. Negation: $[Fg(p)^n]\neg\varphi$ iff (see above) $\neg[Fg(p)^n]\varphi$ iff (IH) $\neg\varphi$. Conjunction: $[Fg(p)^n](\varphi \wedge \psi)$ iff ($[Fg(p)^n]\varphi$ and $[Fg(p)^n]\psi$) iff (IH) (φ and ψ) iff $\varphi \wedge \psi$.

As long as R is serial, so that not just $K\top \leftrightarrow \top$ but also $K\perp \leftrightarrow \perp$ are valid, ignorance is indeed obtained:

Proposition 6 On the class of serial Kripke models, $[Fg(p)^n](\neg Kp \wedge \neg K\neg p)$ is valid. ⊢

Proof. The validity follows directly from the semantics: given any state, after execution of $Fg(p)^n$ a state where p has become true is indistinguishable from one where p has become false. It is also instructive to show an argument employing the axiomatization **FG**. First, we have that $[Fg(p)^n](\neg Kp \wedge \neg K\neg p)$ iff $[Fg(p)^n]\neg Kp \wedge [Fg(p)^n]\neg K\neg p$. For the left conjunct of the latter, we have $[Fg(p)^n]\neg Kp$ iff $\neg[Fg(p)^n]Kp$, which is equivalent to $\neg(K[Fg(p)^0]p \wedge K[Fg(p)^1]p \wedge K[Fg(p)^n]p)$, i.e. $\neg K[Fg(p)^0]p \vee \neg K[Fg(p)^1]p \vee \neg K[Fg(p)^n]p$. This is equivalent to $\neg K\perp \vee \neg K\top \vee \neg Kp$. Given seriality, $\neg K\perp$ is equivalent to \top which makes the disjunction true. The right conjunct is similar.

Corollary 7 $[Fg(p)^n]K(\neg Kp \wedge \neg K\neg p)$ is valid. ⊢

Proof. The preceding Proposition also holds for $[Fg(p)^0]$ and $[Fg(p)^1]$. The corollary now follows from the equivalence of $[Fg(p)^n]K(\neg Kp \wedge \neg K\neg p)$ to $\bigwedge_{j=0,1,n} K[Fg(p)^j](\neg Kp \wedge \neg K\neg p)$.

Proposition 8 (Interaction between variables) Valid are:

1. $[Fg(p)^n][Fg(p)^n]\varphi \leftrightarrow [Fg(p)^n]\varphi$
2. $[Fg(p)^n][Fg(q)^n]\varphi \leftrightarrow [Fg(q)^n][Fg(p)^n]\varphi$ ⊢

Proof. We use that bisimilar structures with respect to the epistemic language satisfy the same formulas in event model logic (Baltag

et al., 1998). Given some model M with domain S wherein the forgetting takes place, and a state $s \in S$, the validity of the first item follows from the bisimulation $\mathfrak{R} : ((s, j), i) \leftrightarrow (s, i)$ (where $i, j = 0, 1, n$), and the validity of the second item follows from the bisimulation $\mathfrak{R} : ((s, i), j) \leftrightarrow ((s, j), i)$ (where $i, j = 0, 1, n$).

The idempotency of forgetting is not surprising. After having become ignorant about p , repeating that procedure has no effect. The second item in the Proposition expresses that the order of forgetting variables is irrelevant. That allows us to define multi-variable forgetting by abbreviation:

Definition 9 (Forgetting multiple variables)

$[Fg(\{q\})^n]\varphi =_{\text{def}} [Fg(q)^n]\varphi$ and $[Fg(P' \cup \{q\})^n]\varphi =_{\text{def}} [Fg(P')^n][Fg(q)^n]\varphi.$ \dashv

3.4. VARIANTS

In this subsection we present two alternative ways to define forgetting with a dynamic modal operator.

Swapping variables An alternative but indistinguishable modelling of forgetting is by making every state in the model indistinguishable from one wherein the value of p has been swapped / switched: if true, it became false, and if false it became true. This forgetting operator $\text{Fg}(p)^n$ comes with the following semantics—instead of three dynamic modal operators two are now sufficient in the logical language, the other one we name $\text{Fg}(p)^{\bar{n}}$. Executing this operator amounts to making two copies of the given model (not three, as for $Fg(p)^n$). Assume an epistemic model $M = (S, R, V)$. Then:

$$(M, s) \models [\text{Fg}(p)^j]\varphi \text{ iff } (M', (s, j)) \models \varphi \quad \text{for } j = n, \bar{n}$$

where $M' = (S', R', V')$ such that: $S' = S \times \{n, \bar{n}\}$; $((s, i), (t, j)) \in R'$ iff $(s, t) \in R$ and $i, j \in \{n, \bar{n}\}$; $V'(p) = (V(p) \times \{n\}) \cup (S \setminus V(p) \times \{\bar{n}\})$, and for $q \neq p$, $V'(q) = V(q) \times \{n, \bar{n}\}$.

We can show that the two languages are equally expressive by way of a bisimulation. Given an epistemic model M , let M' be as above and let M'' be the model produced by execution of our standard forgetting operator $Fg(p)^n$. Given some state $s \in S$, the following defines a (total) bisimulation $\mathfrak{R} : M' \leftrightarrow M''$.

- $\mathfrak{R} : (s, n) \mapsto (s, n)$
- If $s \in V(p)$ then $\mathfrak{R} : (s, 0) \mapsto (s, \bar{n})$ and $\mathfrak{R} : (s, 1) \mapsto (s, n)$

- If $s \notin V(p)$ then $\mathfrak{R} : (s, 0) \mapsto (s, n)$ and $\mathfrak{R} : (s, 1) \mapsto (s, \bar{n})$

The result follows, because bisimulation implies logical equivalence.

Precondition of prior knowledge If we *were* to require a precondition of knowledge about p prior to forgetting p , i.e. precondition $Kp \vee K\neg p$, the forgetting operation becomes partially executable, namely subject to that condition, and the principles of forgetting are as follows.

Proposition 10 (Principles of forgetting with prior knowledge)

$$\begin{array}{ll}
[Fg(p)^n]p & \leftrightarrow (Kp \vee K\neg p) \rightarrow p \\
[Fg(p)^0]p & \leftrightarrow \neg(Kp \vee K\neg p) \\
[Fg(p)^1]p & \leftrightarrow Kp \vee K\neg p \\
[Fg(p)^j]q & \leftrightarrow (Kp \vee K\neg p) \rightarrow q \quad \text{for } q \neq p, j = 0, 1, n \\
[Fg(p)^j](\varphi \wedge \psi) & \leftrightarrow [Fg(p)^j]\varphi \wedge [Fg(p)^j]\psi \quad \text{for } j = 0, 1, n \\
[Fg(p)^j]K\varphi & \leftrightarrow (Kp \vee K\neg p) \rightarrow \bigwedge_{i=0,1,n} K[Fg(p)^i]\varphi
\end{array}$$

–

We now have that one cannot forget something *twice* (at least, without having *learnt* about it in between); after having forgotten p , the precondition $Kp \vee K\neg p$ is no longer satisfied, therefore, the second execution cannot take place.

Proposition 11 (Forgetting twice) $[Fg(p)^n][Fg(p)^n]\perp$ is valid. –

4. Forgetting as becoming unaware

In this section we investigate the relation between the notion of introspective forgetting already defined, and a notion of unawareness forgetting. This provides a very direct proof that the truth of formulas not involving p is preserved after forgetting p , whether introspectively or not.

4.1. UNAWARENESS FORGETTING

Scrambling the valuation of the forgotten atom The third perspective on propositional forgetting in propositional logic, construction 3 on page 5, can be generalized to modal logic by ‘releasing the value of p ’ in a given epistemic model M and then looking for the consequences. This is similar to a proposal in (Lin and Reiter, 1994). Consider again the structure $pq \text{---} p\neg q$ encoding that the agent knows p but is ignorant about q . There are four models agreeing with $pq \text{---} p\neg q$ on everything

except maybe the value of p . What holds after forgetting p (unawareness forgetting) is what holds in all four. This comes with the following transition:

$$pq \text{ — } p\neg q \quad \Longrightarrow \quad \begin{array}{l} \neg pq \text{ — } \neg p\neg q \\ \neg pq \text{ — } p\neg q \\ pq \text{ — } \neg p\neg q \\ pq \text{ — } p\neg q \end{array}$$

The four models on the right still have ignorance about q in common, but only two of them satisfy ignorance about p . The only meaningful formulas about p that are true in all models, such as $\neg Kq \wedge Kp \vee \neg Kp$, are equivalent to formulas not involving p , in this case, to $\neg Kq$. This encodes *unawareness* of p , but not *ignorance* about p .

Consider the even more abstract perspective of *bisimulation quantification*: apart from the four above, $pq \text{ — } p\neg q$ is also similar except for the value of p to other structures, e.g. to $pq \text{ — } p\neg q \text{ — } \neg pq$, three indistinguishable states, which satisfies that $K(p \vee q)$ is true. Because if we abstract from the value of p , then $q \text{ — } \neg q$ is bisimilar to $q \text{ — } \neg q \text{ — } q$. In other words, not just the valuation of p may vary ‘at random’ but also the epistemic uncertainty about its value. This prepares the ground for the next paragraph.

Bisimulation quantification Becoming unaware of an atom p can be modelled as universal bisimulation quantification over p (Visser, 1996; French, 2006) namely as follows. The notation $(M', s') \Leftrightarrow_{P-p}(M, s)$ means that epistemic state (M', s') is bisimilar to epistemic state (M, s) with respect to the set of all atoms *except* p (write $P - p$ for $P \setminus \{p\}$).

Definition 12 (Unawareness forgetting (Zhang and Zhou, 2008))

Unawareness forgetting of atom p is defined as $\forall p\varphi$ (‘after forgetting p , φ ’), where $(M, s) \models \forall p\varphi$ iff for all (M', s') such that $(M', s') \Leftrightarrow_{P-p}(M, s) : (M', s') \models \varphi$. ⊣

This definition satisfies the requirement for unawareness forgetting that the truth of all formulas not involving p should be preserved: bisimilar structures except for p agree on the truth of all such formulas. This original definition is found in (Zhang and Zhou, 2008), and was (probably later, but independently) suggested in (van Ditmarsch and French, 2009). According to (Zhang and Zhou, 2008), we have that $\psi \rightarrow \forall p\varphi$ is valid iff (in their notation) $KForget(\{\psi\}, p) \rightarrow \varphi$ is valid—this is similar to the intended generalization to modal logic of the notion $Fg(\psi, p)$ in (Lang et al., 2003). In (Zhang and Zhou, 2008) it is also observed that by constraining the ψ -models to those also satisfying $\neg Kp \wedge \neg K\neg p$

(some sort of belief expansion following the belief contraction induced by going to all $(P-p)$ -bisimilar models), we can enforce ignorance about p as in (Baral and Zhang, 2005). This relates introspective forgetting to unawareness forgetting. This is possible, and correct, but different from what we will do now to relate the two.

4.2. UNAWARENESS FORGETTING AND INTROSPECTIVE FORGETTING

Given some epistemic model M , the model M' wherein p is introspectively forgotten according to $Fg(p)^n$ is bisimilar to M except for p .

$$M \Leftrightarrow_{P-p} M'$$

The total bisimulation is defined as $\mathfrak{R} : s \mapsto (s, j)$ for $j = 0, 1, n$. This implies that:

Theorem 13 If $\psi \in \mathcal{L}(P-p)$ then $\psi \leftrightarrow [Fg(p)^n]\psi$ is valid. \dashv

It also follows that for formulas not containing the forgotten variable, unawareness forgetting and introspective forgetting amount to the same.

Corollary 14 If $\psi \in \mathcal{L}(P-p)$ then ψ is logically equivalent to $\forall p\psi$, and also logically equivalent to $[Fg(p)^n]\psi$. \dashv

Alternatively we can define an introspective forgetting operator by way of bisimulation quantification. We then have to make corresponding points in all $P-p$ -bisimilar models indistinguishable for the agent. This can be achieved, somewhat artificially, as follows:

Given an epistemic model (M, s) , consider the direct sum of $\{(M', s') \mid (M', s') \Leftrightarrow_{P-p}(M, s)\}$; assume some index set I such that each such (M', s') has a different index $i \in I$ and for convenience assume index 0 for (M, s) . The model \mathfrak{M} is this direct sum plus additional pairs in the accessibility relation $R^{\mathfrak{M}}$ as follows. Suppose $\mathfrak{R}' : (M, s) \Leftrightarrow_{P-p} (M', s')$ and $\mathfrak{R}'' : (M, s) \Leftrightarrow_{P-p} (M'', s'')$, and let i be the index of (M', s') and j be the index of (M'', s'') . If for some $t \in M$, $stateb' \in M'$, $t'' \in M''$: $(t, t') \in \mathfrak{R}'$ and $(t, t'') \in \mathfrak{R}''$, then $add((t', i), (t'', j))$ to $R^{\mathfrak{M}}$. Clearly, $\mathfrak{M} \models \neg Kp \wedge \neg K\neg p$. Now define

Definition 15 $(M, s) \models \bar{\forall}p\varphi$ iff $(\mathfrak{M}, (s, 0)) \models \varphi$, with \mathfrak{M} as above. \dashv

Then $\bar{\forall}$ defines a notion of introspective forgetting, and $\bar{\forall}p(\neg Kp \wedge \neg K\neg p)$ is valid. We even have that, if M' is the model resulting from executing $Fg(p)^n$:

$$M \Leftrightarrow_{P-p} \mathfrak{M} \Leftrightarrow_{P-p} M'$$

So, we can after all reclaim introspective forgetting using a bisimulation quantification. We prefer $Fg(p)^n$, as this is a constructive notion.

5. Multiple agents

Given a parameter set of agents A , we adjust the language to

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi \mid [Fg_B(p)^j]\varphi,$$

where $p \in P$, $a \in A$, $j = 0, 1, n$, and $B \subseteq A$. The formula $[Fg_B(p)^n]\varphi$ stands for ‘the agents in B (collectively) forget φ ’. In the Kripke models we replace the accessibility relation $R \subseteq (S \times S)$ by an accessibility *function* $R : A \rightarrow \mathcal{P}(S \times S)$. We write R_a for $R(a)$. Accessibility relation R_a interprets operator K_a , describing the knowledge of agent a . The clauses in the semantics for knowledge and for forgetting are now different.

Definition 16 (Semantics of multi-agent forgetting) Assume a multi-agent epistemic model $M = (S, R, V)$ where R is an accessibility *function*.

$$\begin{aligned} (M, s) \models K_a\varphi & \quad \text{iff for all } t \in S : (s, t) \in R_a \text{ implies } (M, t) \models \varphi \\ (M, s) \models [Fg_B(p)^j]\varphi & \quad \text{iff } (M', (s, j)) \models \varphi \quad \text{for } j = 0, 1, n \end{aligned}$$

where $M' = (S', R', V')$ such that

- $S' = S \times \{0, 1, n\}$;
- for all $a \in B$, $((s, i), (t, j)) \in R'_a$ iff $(s, t) \in R_a$ and $i, j \in \{0, 1, n\}$;
- for all $a \in A \setminus B$, $((s, i), (t, i)) \in R'_a$ iff $(s, t) \in R_a$ and $i \in \{0, 1, n\}$;
- $V'(p) = \{(s, 1) \mid s \in S\} \cup (V(p) \times \{n\})$ and for $q \neq p$, $V'(q) = V(q) \times \{0, 1, n\}$.

As before, executing forgetting can be seen as making three copies of the current model, one wherein nothing happens, one wherein p becomes false and one wherein p becomes true. Corresponding states in these copies cannot be distinguished from one another by the forgetting agents (the agents in group B), but they can be distinguished by the remaining agents (and this is common knowledge). \dashv

Definition 17 (Axiomatization \mathbf{FG}_A for multi-agent forgetting)

The axiomatization \mathbf{FG}_A is as \mathbf{FG} (replacing K by K_a and $Fg(p)^j$ by

$Fg_B(p)^j$ everywhere), except for the axiom on the interaction between forgetting and knowledge, for which there are now two different cases:

$$\begin{aligned} [Fg_B(p)^j]K_a\varphi &\leftrightarrow \bigwedge_{i=0,1,n} K_a[Fg_B(p)^i]\varphi && \text{if } a \in B \\ [Fg_B(p)^j]K_a\varphi &\leftrightarrow K_a[Fg_B(p)^j]\varphi && \text{if } a \in A \setminus B \quad \dashv \end{aligned}$$

There must be two cases, because we need to distinguish the effects of forgetting on the agents involved in that, from the effects of forgetting on the remaining agents. If they already knew p , they will still know p after forgetting! A simple bisimulation suffices to demonstrate the validity of the following interaction for different groups of agents.

Proposition 18 (Multi-agent forgetting)

$$[Fg_B(p)][Fg_{B'}(p)]\varphi \leftrightarrow [Fg_{B'}(p)][Fg_B(p)]\varphi \leftrightarrow [Fg_{B \cup B'}(p)]\varphi \quad \dashv$$

And we emphasize that all other results for introspective forgetting remain applicable to the multi-agent case. So, multiple variable forgetting for multiple agents is again commutative and definable, factual truths are still preserved, and so on.

Two particular cases of ‘group forgetting’ seem of special interest—both already came to the fore in the introductory section. The case $Fg_A(p)$ where B is the set of all agents A models public introspective forgetting, as in the public belief in unlimited credit. The case $Fg_a(p)$ where $B = \{a\}$ models individual introspective forgetting (namely of agent a) in the presence of A , as in the example of the lecturer exclaiming that he forgot again which one was his office key, in the presence of his students.

In the multi-agent case, unawareness forgetting definable by a bisimulation quantifier only applies to public forgetting: forgetting by the entire group of agents. This applies to the Men in Black example, where all collectively forget their memory of extraterrestrials. Adjustments to subgroup unawareness forgetting can also be made, but are beyond our current scope.

6. Further research

Remembering prior knowledge For the agent to recall prior knowledge we have to be able to refer to past events. Let $Fg(p)^-$ be the converse of $Fg(p)$ (e.g. in the sense of (Aucher and Herzig, 2007; Yap, 2006; Sack, 2007)). Awareness of present ignorance and prior knowledge about p can now be formalized as

$$K(\neg Kp \wedge \neg K\neg p \wedge \langle Fg(p)^- \rangle (Kp \vee K\neg p))$$

We now need a structure allowing us to interpret such converse events. This is not possible in pointed Kripke models, but it can be elegantly done employing what is known as the ‘forest’ produced by the initial Kripke model and all possible sequences of all $Fg(p)$ events (for all atoms), see (Parikh and Ramanujam, 1985; van Benthem et al., 2007; Pacuit, 2007; van Ditmarsch et al., 2007; Yap, 2006; Sack, 2007). We now add assignments to the language, as in the underlying proposal, and additionally add theories for event models using converse actions (Aucher and Herzig, 2007; van Benthem et al., 2007). Thus we get a complete axiomatization, though not a reduction result to epistemic formulas (converse events cannot be eliminated from the logical language by equivalences). It seems that a reduction result can *still* not be obtained if we require prior knowledge about p as preconditions in the event model for forgetting.

Combining learning and forgetting One might wish to combine forgetting with other dynamic operations such as learning (by public announcements). We simply add an inductive construct $[\varphi]\psi$ to the language, which stands for ‘after announcement of φ , ψ holds’ (Plaza, 1989). The resulting logic is again equally expressive as epistemic logic: just add the rewrite rules involving announcements. One can formulate principles such as, for forgetting with prior knowledge about p as precondition: $\varphi \leftrightarrow [Fg(p)]([\![p]\!] \varphi \vee [\![\neg p]\!] \varphi)$, where $[p]$ and $[\![\neg p]\!]$ are the dynamic operators corresponding to ‘revealing the truth about p ’, i.e., the announcement of p and the announcement of $\neg p$, respectively; a.k.a. the sensing action of p . In other words: if you forget about p and afterwards learn about p again, nothing has changed. The principle does not hold without the precondition of prior knowledge, as $Fg(p)$ is also executable when the agent is ignorant about p . The principle is reminiscent of the *recovery principle* in AGM belief revision (Gärdenfors, 1988; Gärdenfors, 1992).

Forgetting modal formulas How to model the forgetting of modal formulas is unclear. We recall the introductory example wherein “I may have forgotten whether you knew about a specific review result for our jointly edited journal issue”: previously $K_{me}K_{you}accept$ or $K_{me}K_{you}\neg accept$ but currently $\neg K_{me}K_{you}accept$ and $\neg K_{me}K_{you}\neg accept$. Even in a propositional logical setting, the forgetting of non-variables is entirely unclear. A semantically motivated solution for the forgetting of modal formulas would require to adjust the part of the model required to make the formula true that is to be forgotten – but to do so in a minimal way not affecting the truth of other formulas. For such a kind

of forgetting, the simple ‘triplication’ of all states in the model is no longer sufficient.

Forgetting of events This amounts to introducing *temporal uncertainty* in the model, e.g., when given a history of five prior events, the agent has forgotten that event number three took place (including, therefore, all its epistemic postconditions). This can be done by introducing histories of events to structures, or moving to a temporal epistemic perspective using ‘forests’, as above, see (van Benthem *et al.*, 2007). It is clear how this has to be done, and the results should prove interesting.

Complexities It would be worthwhile to generalize the results in (Lang *et al.*, 2003) to the modal case.

Acknowledgements

Previous versions of this work were presented at the ESSLLI 2008 workshop Logic and Intelligent Interaction, in Hamburg, Germany (see <http://ai.stanford.edu/~epacuit/LaII/>), and at the Advances in Artificial Intelligence conference (AI08), in Auckland, New Zealand (see (van Ditmarsch *et al.*, 2008)). The underlying version contains: a more extensive overview comparing it to known results in propositional variable forgetting, more proof details, and full details and results on multiple-variable and multi-agent forgetting. The authors would like to thank the four anonymous AI08 referees, and also the participants of the ESSLLI 2008 Hamburg workshop Logic and Intelligent Interaction, in particular Audrey Yap who observed that temporal uncertainty may also result in epistemic uncertainty that amounts to forgetting. Hans van Ditmarsch acknowledges support of the Netherlands Institute of Advanced Study where he was Lorentz Fellow in 2008. Hans van Ditmarsch had numerous discussions with Philippe Balbiani on the subject of forgetting, he suggested yet other original perspectives that have not been discussed in this paper. Finally we would like to thank the two anonymous reviewers of the KRA journal. Both provided valuable comments. One of them wrote a review of the same length as the contribution. This was incredibly helpful, and also in other respects we acknowledge the generosity of that reviewer.

References

- Alchourrón, C., P. Gärdenfors, and D. Makinson: 1985, ‘On the logic of theory change: partial meet contraction and revision functions’. *Journal of Symbolic Logic* **50**, 510–530.
- Aucher, G. and A. Herzig: 2007, ‘From DEL to EDL : Exploring the Power of Converse Events’. In: K. Mellouli (ed.): *ECSQARU*, Vol. 4724 of *Lecture Notes in Computer Science*. pp. 199–209.
- Baltag, A., L. Moss, and S. Solecki: 1998, ‘The Logic of Public Announcements, Common Knowledge, and Private Suspicions’. In: I. Gilboa (ed.): *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*. pp. 43–56.
- Baltag, A. and S. Smets: 2006, ‘Dynamic Belief Revision over Multi-Agent Plausibility Models’. Proceedings of LOFT 2006 (7th Conference on Logic and the Foundations of Game and Decision Theory).
- Baral, C. and Y. Zhang: 2005, ‘Knowledge updates: semantics and complexity issues’. *Artificial Intelligence* **164**(1-2), 209–243.
- Blackburn, P., M. de Rijke, and Y. Venema: 2001, *Modal Logic*. Cambridge: Cambridge University Press. Cambridge Tracts in Theoretical Computer Science 53.
- Boole, G.: 1854, *An investigation of the laws of thought*. Basingstoke: Macmillan Publishers.
- Eiter, T. and K. Wang: 2006, ‘Forgetting and Conflict Resolving in Disjunctive Logic Programming’. In: *Proceedings of AAAI*.
- Erdem, E. and P. Ferraris: 2007, ‘Forgetting Actions in Domain Descriptions’. In: *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*. pp. 409–414.
- Fagin, R. and J. Halpern: 1988, ‘Belief, awareness, and limited reasoning’. *Artificial Intelligence* **34**(1), 39–76.
- Fagin, R., J. Halpern, Y. Moses, and M. Vardi: 1995, *Reasoning about Knowledge*. Cambridge MA: MIT Press.
- French, T.: 2006, ‘Bisimulation quantifiers for modal logic’. Ph.D. thesis, University of Western Australia.
- Gärdenfors, P.: 1988, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge, MA: Bradford Books, MIT Press.
- Gärdenfors, P.: 1992, *Belief Revision*. Cambridge University Press. Cambridge tracts in theoretical computer science 29.
- Harel, D., D. Kozen, and J. Tiuryn: 2000, *Dynamic Logic*. Cambridge MA: MIT Press. Foundations of Computing Series.
- Herzig, A., J. Lang, and P. Marquis: 2003, ‘Action representation and partially observable planning using epistemic logic’. In: *Proceedings of IJCAI*.
- Lang, J., P. Liberatore, and P. Marquis: 2003, ‘Propositional Independence: Formula-Variable Independence and Forgetting’. *J. Artif. Intell. Res. (JAIR)* **18**, 391–443.
- Liberatore, P. and M. Schaerf: 1998, ‘Arbitration (or How to Merge Knowledge Bases)’. *IEEE Trans. on Knowl. and Data Eng.* **10**(1), 76–90.
- Lin, F. and R. Reiter: 1994, ‘Forget It!’. In: *Proceedings of AAAI Fall Symposium on Relevance, New Orleans*. pp. 154–159.
- Pacuit, E.: 2007, ‘Some comments on history-based structures’. To appear in *Journal of Applied Logic*.
- Parikh, R. and R. Ramanujam: 1985, ‘Distributed Processing and the Logic of Knowledge’. In: *Logic of Programs*, Vol. 193 of *Lecture Notes in Computer*

- Science*. pp. 256–268. A newer version appeared in *Journal of Logic, Language and Information*, vol. 12, 2003, pp. 453–467.
- Plaza, J.: 1989, ‘Logics of Public Communications’. In: M. Emrich, M. Pfeifer, M. Hadzikadic, and Z. Ras (eds.): *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems: Poster Session Program*. pp. 201–216. ORNL/DSRD-24.
- Sack, Y.: 2007, ‘Adding Temporal Logic to Dynamic Epistemic Logic’. Ph.D. thesis, Indiana University, Bloomington, USA.
- Segerberg, K.: 1999, ‘Two traditions in the logic of belief: bringing them together’. In: H. Ohlbach and U. Reyle (eds.): *Logic, Language, and Reasoning*. Dordrecht, pp. 135–147.
- van Benthem, J., J. Gerbrandy, and E. Pacuit: 2007, ‘Merging frameworks for interaction: DEL and ETL’. In: D. Samet (ed.): *Proceedings of TARK 2007*. pp. 72–81.
- van Benthem, J., J. van Eijck, and B. Kooi: 2006, ‘Logics of Communication and Change’. *Information and Computation* **204(11)**, 1620–1662.
- van Ditmarsch, H.: 2005, ‘Prolegomena to dynamic logic for belief revision’. *Synthese (Knowledge, Rationality & Action)* **147**, 229–275.
- van Ditmarsch, H. and T. French: 2009, ‘Simulation and information’. In: J. Broersen and J.-J. Meyer (eds.): *Proceedings of KRAMAS, Sydney*. To appear. Also presented at LOFT 2008, Amsterdam.
- van Ditmarsch, H., A. Herzig, J. Lang, and P. Marquis: 2008, ‘Introspective forgetting’. In: W. Wobcke and M. Zhang (eds.): *AI 2008: Advances in Artificial Intelligence*. pp. 18–29. LNAI 5360.
- van Ditmarsch, H. and B. Kooi: 2008, ‘Semantic Results for Ontic and Epistemic Change’. In: G. Bonanno, W. van der Hoek, and M. Wooldridge (eds.): *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games 3. Amsterdam: Amsterdam University Press, pp. 87–117.
- van Ditmarsch, H., J. Ruan, and W. van der Hoek: 2007, ‘Model checking dynamic epistemics in branching time’. In: *(Informal) Proceedings of FAMAS 2007, Durham UK*.
- Visser, A.: 1996, ‘Bisimulations, Model Descriptions and Propositional Quantifiers’. Logic Group Preprint Series 161, Department of Philosophy, Utrecht University.
- Wang, K., A. Sattar, and K. Su: 2005, ‘A Theory of Forgetting in Logic Programming’. In: *AAAI*. pp. 682–688.
- Winslett, M.: 1991, *Updating Logical Databases*, Cambridge Tracts in Theoretical Computer Science (No. 9). Cambridge University Press.
- Yap, A.: 2006, ‘Product Update and Looking Backward’. Technical report, University of Amsterdam. ILLC Research Report PP-2006-39.
- Zhang, Y., N. Foo, and K. Wang: 2005, ‘Solving Logic Program Conflict through Strong and Weak Forgetting’. In: *Proceedings of IJCAI*. pp. 627–634.
- Zhang, Y. and Y. Zhou: 2008, ‘Knowledge Forgetting: Properties and Applications’. Work in progress under submission.
- Zhao, Y., K. Wang, R. Topor, J. Pan, and F. Giunchiglia: 2007, ‘Semantic Cooperation and Knowledge Reuse by Using Autonomous Ontologies’. In: *ISWC/ASWC*. pp. 666–679.

Appendix: Forgetting as an event model

The dynamic operator $[Fg(p)^n]$ is relative to the state transformer $Fg(p)^n$ that is a pointed *event model*. The pointed Kripke models are static structures, encoding knowledge and belief, and the event models are dynamic structures, encoding *change of* knowledge and belief. Formally, *event models* (a.k.a. action models) are structures $M = (S, R, \text{pre}, \text{post})$, where $\text{pre} : S \rightarrow \mathcal{L}$ assigns to each event $s \in S$ a *precondition* and where $\text{post} : S \rightarrow (P \rightarrow \mathcal{L})$ assigns to each event a *postcondition*, a.k.a. assignment, for each atom of a *finite* subset of all atoms. The remaining atoms do not change value. If $s \in S$, then (M, s) is a pointed event model. For such event models see (van Benthem et al., 2006; van Ditmarsch and Kooi, 2008). Dynamic operators expressing event model execution (*semantics*) can be seen as part of the logical language (*syntax*), similar to how this is done for automata-PDL (Harel et al., 2000).

Forgetting $Fg(p)^n$ is the pointed event model (read this as $(Fg(p), n)$) that expresses that the agent cannot distinguish between nothing happening or one of two assignments having taken place: p becomes true, or p becomes false. It consists of three events, that are always executable: their precondition is \top . The non-pointed version of forgetting is $Fg(p)$.

Definition 19 (Forgetting) $Fg(p)$ is the event model $((S, R, \text{pre}, \text{post}), S')$ where $S = \{0, 1\}$, $R = S \times S$, $\text{pre}(0) = \top$ and $\text{pre}(1) = \top$, $\text{post}(0)(p) = \perp$ and $\text{post}(1)(p) = \top$ (and $\text{post}(i)(q) = q$ for all $q \neq p$, $j = 0, 1$), and $S' = S$. \dashv

The model M' resulting from executing forgetting $Fg(p)^n$ in a given model (M, s) is in fact the restricted modal product $M \otimes Fg(p)$, with actual state (s, n) . See (Baltag et al., 1998; van Ditmarsch and Kooi, 2008) for details on its computation.

The axioms listed in **FG** are simple instantiations of the reduction principles for event model logic. For example, the axiom to reduce an atom following an event model is

$$[M, s]p \leftrightarrow (\text{pre}(s) \rightarrow \text{post}(s)(p))$$

For the specific events $Fg(p)^j$ (i.e., $(Fg(p), j)$) for $j = 0, 1, n$ we therefore get:

$$[Fg(p)^n]p \text{ iff } (\text{pre}(n) \rightarrow \text{post}(n)(p)) \text{ iff } (\top \rightarrow p) \text{ iff } p$$

$$[Fg(p)^0]p \text{ iff } (\text{pre}(0) \rightarrow \text{post}(0)(p)) \text{ iff } (\top \rightarrow \perp) \text{ iff } \perp$$

$$[Fg(p)^1]p \text{ iff } (\text{pre}(1) \rightarrow \text{post}(1)(p)) \text{ iff } (\top \rightarrow \top) \text{ iff } \top$$

$$[Fg(p)^j]q \text{ iff } (\text{pre}(j) \rightarrow \text{post}(j)(q)) \text{ iff } (\top \rightarrow q) \text{ iff } q$$

The cases for the other axioms are even more straightforward.

