Logics of rational interaction

Barteld Kooi Eric Pacuit

March 5, 2010

1 Introduction

There is a growing literature focused on using logical methods to reason about communities of agents engaged in some form of social interaction. Much of the work builds upon existing logical frameworks developed by philosophers and computer scientists incorporating insights and ideas from philosophy (especially epistemology and philosophy of action), game theory, decision theory and social choice theory. The result is a web of logical systems each addressing different aspects of rational agency and social interaction. Rather than providing an encyclopedic account of these different logical systems[1], this paper focuses on issues surrounding the modeling of informational attitudes in social interactive situations. The main objective is to introduce the two main approaches to modeling "rational interaction" and provide pointers to the current literature.

Of course, there is no single approach that can address *all* of the complex phenomena that arise when rational agents interact with one another and the environment. Thus it is important to understand how the different analyses from within and across the disciplines mentioned above can fit together. This suggests the following three general questions:

1. How can we *compare* different logical frameworks addressing similar aspects of rational agency and social interaction (i.e., how information evolves through social interaction)?

2. How should we *combine* logical systems which address *different* aspects of social interaction towards the goal of a comprehensive (formal) theory of rational agency?

3. How do the logical frameworks discussed in this literature contribute to the broader discussion of rational agency and social interaction within philosophy and the social sciences?

[1]The interested reader can consult Meyer and Veltman (2007); van der Hoek and Wooldridge (2003); van Benthem (2008) and references therein.

Certainly, the first two questions raise numerous methodological issues and technical problems. However, they also make explicit certain foundational and philosophical issues surrounding rational interaction (cf. van Benthem et al., 2008). In particular, viewing the various logical systems found in the literature as (sometimes competing) accounts of rational agency forces us to carefully examine what we even mean by a 'rational agent' (see van Benthem, 2005, for an extensive discussion). Of course, the nature of rationality and human agency is a central concern of many philosophers from Aristotle to David Hume to present-day philosophers (cf. Bratman, 2007; Searle, 1985; Hyman and Steward, 2004). The point here is that there are many different types of reasoning and dynamic processes that agents use when interacting with other agents. Comparing and combining the different logical systems forces us to consider how these different processes interact.

In this survey, the modeling of informational attitudes of a group of *rational* agents engaged in some form of social interaction (eg. having a conversation or playing a card game) takes center stage. Indeed, there are many logical systems today that describe how an agent's information changes over time. Sometimes the differences between two competing logical systems are technical in nature reflecting different conventions used by different research communities. And so, with a certain amount of technical work, such frameworks are seen to be equivalent up to model transformations (cf. Halpern, 1999; Lomuscio and Ryan, 1997; Pacuit, 2007; van Benthem et al., 2008). Other differences point to key conceptual issues about rational interaction. We will introduce the two main logical accounts of rational interaction and highlight such similarities and differences.

2 Reasoning about rational interaction

This section introduces two logical frameworks that describe the dynamics of information over time in a multiagent situation. The first is *epistemic temporal logic* (ETL, Fagin et al., 1995; Parikh and Ramanujam, 1985) which uses linear or branching time models with added epistemic structure induced by the agents' different capabilities for observing events. These models provide a "grand stage" where histories of some social interaction unfold constrained by a **protocol**. Here a **protocol** is intended to represent the rules or conventions that govern many of our social interactions. For example, in a conversation, it is typically not polite to "blurt everything out at the beginning", as we must speak in small chunks. Other natural conversational protocol rules include "do not repeat yourself", "let others speak in turn", and "be honest". Imposing such rules *restricts* the legitimate sequences of possible statements.

The other framework is *dynamic epistemic logic* (DEL, Gerbrandy, 1999;

Baltag et al., 1998a; van Ditmarsch et al., 2007) that describes social interactions in terms of epistemic **event models** (which may occur inside modalities of the language). Similar to the way Kripke structures are used to capture the information the agents' have about a *fixed* social situation[2], an **event model** describes the agents' information about which actual events are currently taking place. The temporal evolution of the situation is then computed from some initial epistemic model through a process of successive "product updates". Details of both frameworks are provided in the subsequent sections.

Often DEL and ETL are presented as *competing* ways of adding dynamics to multi-agent epistemic models. Based on (van Benthem et al., 2008; van Benthem, 2006; van Benthem and Pacuit, 2006), we will see how DEL and ETL should rather be viewed as *complementary* accounts of social interaction. The focus is on conceptual issues leaving some of the more technical details and proofs to the relevant papers. The following running example will help guide intuitions (also discussed in Pacuit and Parikh, 2006).

Example 2.1 *Suppose that Ann would like Bob to attend her talk; however, she only wants Bob to attend if he is interested in the subject of her talk, not because he is just being polite. There is a very simple procedure to solve Ann's problem: Have a (trusted) friend tell Bob the time and subject of her talk.*

Taking a cue from computer science, perhaps we can prove *that this simple procedure correctly solves Ann's problem. However, it is not so clear how to define a correct solution to Ann's problem. If Bob is actually present during Ann's talk, can we conclude that Ann's procedure succeeded? Not really. Bob may have figured out that Ann wanted him to attend, and so is there only out of politeness. Thus for Ann's procedure to succeed, she must achieve a certain "level of knowledge" (cf. Parikh, 2003) between her and Bob. Besides both Ann and Bob knowing about the talk and Ann knowing that Bob knows about*

Bob does not know *that Ann knows about the talk.*

This last point is important, since, if Bob knows that Ann knows that he knows about the talk, he may feel social pressure to attend[3]*. Thus, the procedure* to have a friend tell Bob about the talk, but not reveal that it is at Ann's suggestion*, will satisfy all the conditions. Telling Bob directly will satisfy the first three, but not the essential last condition.*

[2]A Kripke structure is a set of states with relations on this set for each agent. The states, or possible worlds, represent different ways the social situation could have evolved and the relations describe the agents' (current) information. See, for example, (Blackburn et al., 2002; Fagin et al., 1995) for details.

[3]Of course, this is not meant to be a complete analysis of "social politeness".

2.1 Epistemic Temporal Logic

Fix a finite set of agents $\mathcal{A}$ and a (possibly infinite) set of events[4] Σ. A **history** is a finite sequence of events[5] from Σ. We write Σ^* for the set of histories built from elements of Σ. For a history h, we write he for the history h followed by the event e. Given $h, h' \in \Sigma^*$, we write $h \preceq h'$ if h is a prefix of h', and $h \prec_e h'$ if $h' = he$ for some event e.

For example, consider the social interaction described in Example 2.1. There are three relevant participants: Ann (A), Bob (B) and Ann's friend (call him Charles (C)). What are the relevant primitive events? To keep things simple, assume that Ann's talk is either at 2PM or 3PM and initially none of the agents know this. Say, that Ann receives a message stating that her talk is at 2PM (denote this event — Ann receiving a private message saying that her talk is at 2PM — by $\mathsf{e}_{\mathsf{A}}^{\mathsf{2PM}}$). Now, after Ann receives the message that the talk is at 2PM, she proceeds to tell her trusted friend Charles that the talk is at 2PM (and that she wants him to inform Bob of the time of the talk without acknowledging that the information can from her — call this event $\mathsf{e}_{\mathsf{C}}^{\mathsf{A}}$), then Charles tells Bob this information (call this event $\mathsf{e}_{\mathsf{B}}^{\mathsf{C}}$). Thus, the history

$$\mathsf{e}_{\mathsf{A}}^{\mathsf{2PM}}\ \mathsf{e}_{\mathsf{C}}^{\mathsf{A}}\ \mathsf{e}_{\mathsf{B}}^{\mathsf{C}}$$

represents the sequence of events where "Ann receives a (private) message stating that the talk is at 2PM, Ann tells Charles the talk is at 2PM, then Charles tells Bob the talk is at 2PM". Of course, there are other events that are also relevant to this situation. For one thing, Ann could have received a message stating that her talk is at 3PM (denote this event by $\mathsf{e}_{\mathsf{A}}^{\mathsf{3PM}}$). This will be important to capture Bob's uncertainty about whether Ann knows that he knows about the talk. Furthermore, Charles may learn about the time of the talk independently of Ann (denote these two events by $\mathsf{e}_{\mathsf{C}}^{\mathsf{2PM}}$, $\mathsf{e}_{\mathsf{C}}^{\mathsf{3PM}}$). So, for example, the history

$$\mathsf{e}_{\mathsf{A}}^{\mathsf{2PM}}\ \mathsf{e}_{\mathsf{C}}^{\mathsf{2PM}}\ \mathsf{e}_{\mathsf{B}}^{\mathsf{C}}$$

represents the situation where Charles independently learns about the time of the talk and informs Bob.

[4]There is a large literature addressing the many subtleties surrounding the very notion of an *event* and when one event *causes* another event (see, for example, Cartwright, 2007). However, for this paper we take the notion of event as primitive. What is needed is that if an event takes place at some time t, then the fact that the event took place can be observed by a relevant set of agents at t. Compare this with the notion of an event from probability theory. If we assume that at each clock tick a coin is flipped exactly once, then "the coin landed heads" is a possible event. However, "the coin landed head more than tails" would not be an event, since it cannot be observed at any one moment. As we will see, the second statement will be considered a *property* of histories, or sequences of events.

[5]To be precise, elements of Σ should, perhaps, be thought of as event *types* whereas elements of a history are event *tokens*.

There are a number of simplifying assumptions that we adopt in this section. They are not crucial for the analysis of Example 2.1, but do simplify the some of the formal details. Since, histories are sequences of (discrete) events, we assume the existence of a global discrete clock (whether the agents have access to this clock is another issue that will be discussed shortly). The length of the history then represents the amount of time that has passed. Note that this implies that we are assuming a finite past with a possibly infinite future. Furthermore, we assume that at each clock tick, or moment, *some* event takes place (which need not be an event that any agent directly observes). Thus, we can include an event $\mathsf{e_t}$ (for 'clock tick') which can represent that "Charles does *not* tell Bob that the talk is at 2PM." So the history

$$\mathsf{e_A^{2PM}}\ \mathsf{e_C^A}\ \mathsf{e_t}$$

describes the sequence of events where, after learning about the time of the talk, Ann informs Charles, but Charles does *not* go on to tell Bob that the talk is at 2PM. Once a set of events Σ is fixed, the temporal evolution and moment-by-moment uncertainty of the agents can be described.

Definition 2.2 (ETL Frames) Let Σ be a set of events. A **protocol** is a set $\mathsf{H} \subseteq \Sigma^*$ closed under non-empty prefixes. An **ETL frame** is a tuple $\langle \Sigma, \mathsf{H}, \{\sim_i\}_{i \in \mathcal{A}} \rangle$ with H a protocol, and for each $i \in \mathcal{A}$, a binary relation $\sim_i$ on[6] H. ◁

An ETL frame describes how the agents' *hard* information[7] evolves over time in some social situation. The protocol describes (among other things) the temporal structure, with h' such that $h \prec_e h'$ representing the point in time after e has happened in h. The relations $\sim_i$ represent the uncertainty of the agents about how the current history has evolved. Thus, $h \sim_i h'$ means that from agent i's point of view, the history h' looks the same as the history h.

Note that the protocol in an ETL frame captures not only the temporal structure of the social situation being modeled but also assumptions about the nature of the participants. For example, the following is a possible protocol built from the events described above:

---

[6]Although we will not do so here, typically it is assumed that $\sim_i$ is an equivalence relation.

[7]As opposed to *soft* information which may revised. See (van Benthem, 2007) for a general discussion of hard and soft information.

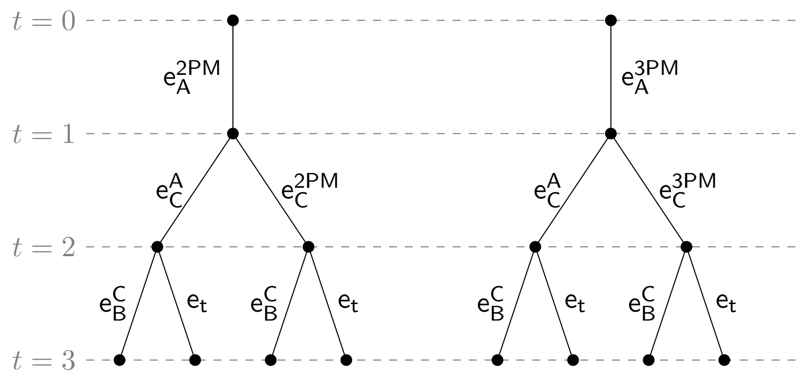

While this protocol does describe possible ways the situation described in Example 2.1 could evolve, it does not account for the *motivation* of the agents. For example, the history

$$e_A^{3PM}\ e_C^A\ e_B^C$$

describes the sequence of events where Ann learns the talk is at 3PM but tells Charles (who goes on to inform Bob) that the talk is at *2PM*. Of course, given the assumption that Ann *wants* Bob to attend her talk, this should not be part of (Ann's) protocol. Similarly, since we assume Charles is trustworthy, we should not include any histories where e_t follows the event e_C^A. Taking into account these underlying assumptions about the motivations (eg. Ann wants Bob to attend the talk) and dispositions (eg. Charles tells the truth and lives up to his promises) of the agents we can drop a number of histories from the protocol shown above. Note that we keep the history

$$e_A^{2PM}\ e_C^{2PM}\ e_t$$

in the protocol, since if Charles learns independently about the time of the talk, then he is under no obligation to inform Bob. In the picture below, we also add some of the uncertainty relations for Ann and Bob (to keep the picture simple, we do not draw the full ETL frame). The solid line represents Bob's uncertainty while the dashed line represents Ann's uncertainty. The main assumption is that Bob can only observe the event (e_B^C). So, for example, the histories $h = e_A^{2PM}\ e_C^A\ e_B^C$ and $h' = e_A^{2PM}\ e_C^{2PM}\ e_B^C$ look the same to Bob (i.e., $h \sim_B h'$).[8]

[8]Note that we do not include any reflexive arrows in the picture in order to keep things simple.

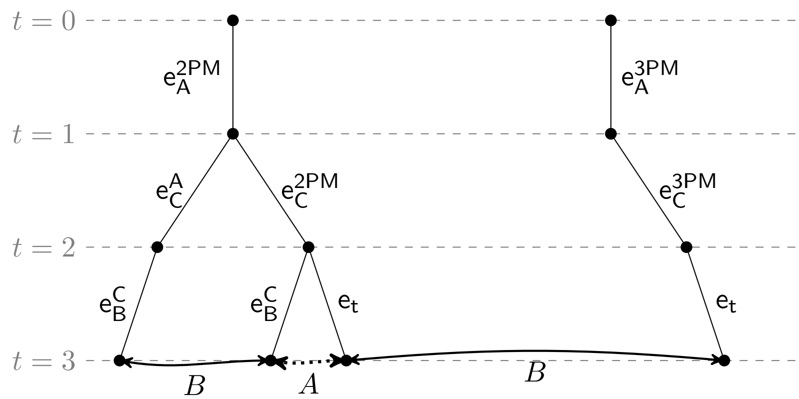

Assumptions about the underlying protocol in an ETL frame corresponds to "fixing the playground" where the agents will interact. As we have seen, the protocol not only describes the temporal structure of the situation being modeled, but also any *causal* relationships between events (eg., sending a message must always proceed receiving that message) plus the motivations and dispositions of the participants (eg., liars send messages that they *know* — or believe — to be false). Thus the "knowledge" of agent i at a history h in some ETL frame is derived from both i's observational powers (via the $\sim_i$ relation) and i's information about the (fixed) protocol.

Remark 2.3 (Three Equivalent Approaches) *There are at least two further approaches to uncertainty in the literature. The first, discussed by Parikh and Ramanujam (1985), explicitly describes the agents' "observational" power. That is, each agent i has a set E_i of events she* can *observe*[9]*. For simplicity, we assume $E_i \subseteq \Sigma$ but this is not necessary. A* **local view** *function is a map $\lambda_i : \mathsf{H} \to E_i^*$. Given a finite history $h \in \mathsf{H}$, the intended interpretation of $\lambda_i(h)$ is "the sequence of events observed by agent i at h". The second approach comes from Fagin et al. (1995). Each agent has a set L_i of* **local states** *(if necessary, one can also assume a set L_e of environment states). Events e are tuples of local states (one for each agent) $\langle l_1, \ldots, l_n \rangle$ where for each $i = 1, \ldots, n$, $l_i \in L_i$. Then two finite histories h and h' are i-equivalent provided the local state of the last of event on h and h' is the same for agent i. From a technical point of view, the three approaches (uncertainty relations, local view functions and local states) to modeling uncertainty are equivalent (Pacuit, 2007; van Benthem and Pacuit, 2006, provide the relevant discussions).*

Although, syntactic issues do not play an important role in this paper, we give the bare necessities to facilitate a comparison between ETL and DEL. Dif-

[9]This may be different from what the agent *does* observe in a given situation.

ferent modal languages describe ETL frames (see, for example, Hodkinson and Reynolds, 2006; Fagin et al., 1995), with 'branching' or 'linear' variants. Let At be a countable set of atomic propositions. The language $\mathcal{L}_{ETL}$ is generated by the following grammar:

$$P \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid \langle e\rangle\varphi$$

where $i \in \mathcal{A}$, $e \in \Sigma$ and $P \in \mathsf{At}$. The usual boolean connectives $(\vee, \rightarrow, \leftrightarrow)$ and the dual modal operators (L_i, $[e]$) are defined as usual. The pure epistemic language, denoted $\mathcal{L}_{EL}$, is the fragment of $\mathcal{L}_{ETL}$ with only epistemic modalities (which we will refer to both as the "language of epistemic logic" and the "epistemic fragment" of $\mathcal{L}_{ETL}$ or the language $\mathcal{L}_{DEL}$ defined below). The intended interpretation of '$K_i\varphi$' is "according to agent i's current information, φ is true." The intended interpretation of '$\langle e\rangle\varphi$' is "after event e (does) take place, φ is true." Formulas are interpreted at histories in an **ETL model**:

Definition 2.4 (ETL Model) An **ETL model** is a tuple $\langle \Sigma, \mathsf{H}, \{\sim_i\}_{i\in\mathcal{A}}, V\rangle$ with $\langle \Sigma, \mathsf{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ an ETL frame and V a valuation function ($V : \mathsf{At} \rightarrow 2^{\mathsf{H}}$). ◁

Definition 2.5 (Truth of $\mathcal{L}_{ETL}$ Formulas) Let $\mathcal{H} = \langle \Sigma, \mathsf{H}, \{\sim_i\}_{i\in\mathcal{A}}, V\rangle$ be an ETL model. The truth of a formula φ at a history $h \in \mathsf{H}$, denoted $\mathcal{H}, h \models \varphi$, is defined inductively as follows:

1. $\mathcal{H}, h \models P$ iff $h \in V(P)$
2. $\mathcal{H}, h \models \neg\varphi$ iff $\mathcal{H}, h \not\models \varphi$
3. $\mathcal{H}, h \models \varphi \wedge \psi$ iff $\mathcal{H}, h \models \varphi$ and $\mathcal{H}, h \models \psi$
4. $\mathcal{H}, h \models K_i\varphi$ iff for each $h' \in \mathsf{H}$, if $h \sim_i h'$ then $\mathcal{H}, h' \models \varphi$
5. $\mathcal{H}, h \models \langle e\rangle\varphi$ iff there exists $h' \in \mathsf{H}$ such that $h \prec_e h'$ and $\mathcal{H}, h' \models \varphi$ ◁

It is often natural to extend the language $\mathcal{L}_{ETL}$ with group knowledge operators (e.g., common or distributed knowledge) and more expressive temporal operators (e.g., arbitrary future or past modalities).

2.2 Dynamic Epistemic Logic

An alternative account of interactive dynamics was elaborated by Gerbrandy (1999); Baltag et al. (1998a); van Benthem (2006); van Benthem et al. (2006) and others. From an initial epistemic model, temporal structure evolves as explicitly triggered by complex informative events.

Definition 2.6 (Epistemic Model) Let $\mathcal{A}$ be a finite set of agents and At a set of atomic propositions. An **epistemic model** is a tuple $\langle W, \{R_i\}_{i\in\mathcal{A}}, V\rangle$ where W is a non-empty set, for each $i \in \mathcal{A}$, R_i is a relation[10] on W ($R_i \subseteq W \times W$) and V a valuation function ($V : \mathsf{At} \to 2^W$). We call the set W the domain of $\mathcal{M}$, denoted by $D(\mathcal{M})$. A pair $\mathcal{M}, w$ where $\mathcal{M}$ is an epistemic model and $w \in D(\mathcal{M})$ is called a **pointed epistemic model**. ◁

We can interpret the epistemic language, $\mathcal{L}_{EL}$, defined above at states in an epistemic model. Truth is defined as usual (see Blackburn et al., 2002, for details). We only recall the definition of the knowledge operators:

$$\mathcal{M}, w \models K_i\varphi \text{ iff for each } w' \in W, \text{ if } wR_iw' \text{ then } \mathcal{M}, w' \models \varphi$$

Returning to our running example (Example 2.1), initially we assume that none of the agents knows the time of Ann's talk. Let P be the proposition "Ann's talk is at 2PM." Then this initial model can be pictured as follows: there are two states w and v with P true at w ($w \in V(P)$). The agent's uncertainty relations is the universal relation (since all agents have the same information, we do not label the arrows). Note that the convention followed in this section is that a solid line around a state means that state is the *actual* or current state (i.e., where the formulas are to be evaluated):

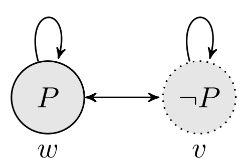

Whereas an ETL frame describes the agents' information at all moments, **event models** are used to build new epistemic models as needed.

Definition 2.7 (Event Model) An **event model** is a tuple $\langle S, \{\longrightarrow_i\}_{i\in\mathcal{A}}, \mathsf{pre}\rangle$, where S is a nonempty set of **primitive events**, for each $i \in \mathcal{A}$, $\longrightarrow_i \subseteq S \times S$ and $\mathsf{pre} : S \to \mathcal{L}_{EL}$ is the **pre-condition function**. The set S in an event model $\mathcal{E}$ is called the domain of $\mathcal{E}$, denoted $D(\mathcal{E})$. ◁

Given two primitive events e and f, $e \longrightarrow_i f$ means that "according to agent i, event e looks like event f." Event models then describe an "epistemic event". In Example 2.1 the first event is Ann receiving a private message that the talk is at 2PM. This can be described by a simple event model: there are two primitive events e and f. The precondition of e is P ($\mathsf{pre}(e) = P$) and the precondition of f is $\top$ (i.e., f is the "skip event").

[10] Again, the R_i are often taken to be equivalence relations on W - but we do not commit.

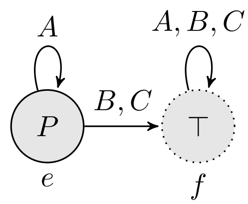

Thus, initially Ann observes the actual event e (and so, learning that P is true) while Bob an Charles observe a skip event (and so, their information does not change). What is the effect of this event on the initial model pictured above? Intuitively, it is not hard to see that after the initial event, Ann knows that P is true while Bob and Charles are still ignorant of P *and the fact that Ann knows* P. That is, combining the initial epistemic model with the above event model should yield the following epistemic model (for simplicity we only draw Ann and Bob's uncertainty relations):

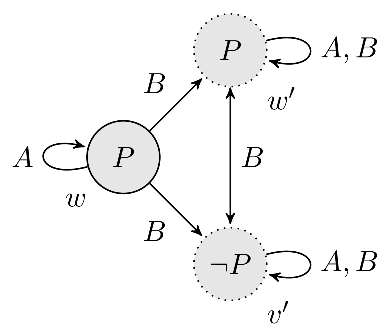

The following definition gives a general procedure for constructing a new epistemic model from a given epistemic model and an event model.

Definition 2.8 (Product Update) The **product update** $\mathcal{M} \otimes \mathcal{E}$ of an epistemic model $\mathcal{M} = \langle W, \{R_i\}_{i\in\mathcal{A}}, V\rangle$ and event model $\mathcal{E} = \langle S, \{\longrightarrow_i\}_{i\in\mathcal{A}}, \mathsf{pre}\rangle$ is the epistemic model $\langle W', R_i', V'\rangle$ with

1. $W' = \{(w,e) \mid w \in W, e \in S$ and $\mathcal{M}, w \models \mathsf{pre}(e)\}$,
2. $(w,e)R_i'(w',e')$ iff wR_iw' in $\mathcal{M}$ and $e \longrightarrow_i e'$ in $\mathcal{E}$, and
3. For all $P \in \mathsf{At}$, $(s,e) \in V'(P)$ iff $s \in V(P)$ $\triangleleft$

We illustrate this construction using our running example. The main event in Example 2.1 is "Charles telling Bob (without Ann present) that Ann's talk is at 2PM". This can be described using the following event model (again only the Ann and Bob relations will be drawn): Ann is aware of the actual event taking place while Bob thinks the event is a private message to himself.

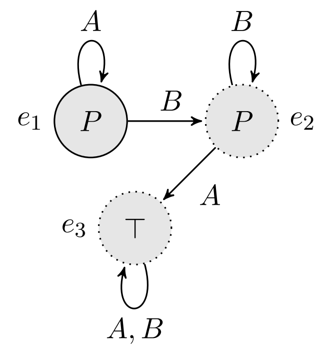

As in the previous section, there are implicit assumptions here about the motivations and dispositions of the agents. Thus, even though Ann is not present during the actual event[11], she *trusts* that Charles will honestly tell Bob that the talk is at 2PM (without revealing he received the information from her). This explains why in the above event model, $e_1 \longrightarrow_A e_1$. Starting from a slightly modified epistemic model from the one given above (where Bob now knows that Ann knows *whether* the talk is at 2PM), using Definition 2.8, we can calculate the effect of the above event model as follows (again focusing only on Ann and Bob's information):

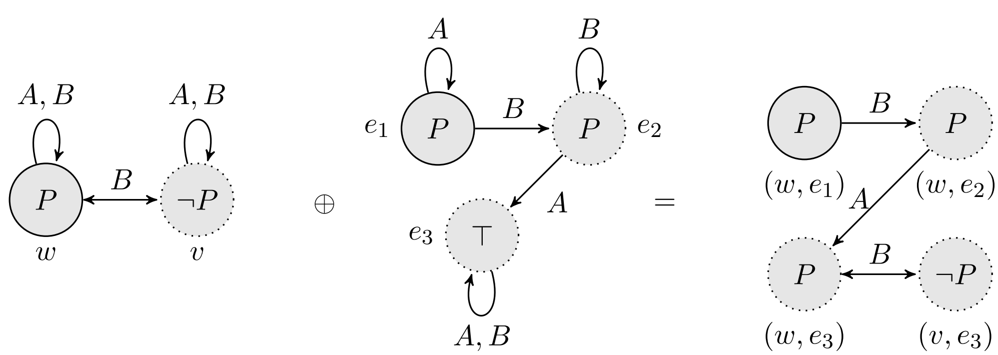

Note that, in the epistemic model on the right, for simplicity, the reflexive arrows are not drawn.

Finally, a few comments about syntactic issues. The language $\mathcal{L}_{DEL}$ extends $\mathcal{L}_{EL}$ with operators $\langle \mathcal{E}, e \rangle$ for each pair of event models $\mathcal{E}$ and event e in the domain of $\mathcal{E}$. Truth for $\mathcal{L}_{DEL}$ is defined as usual. We only define the typical DEL modalities:

$$\mathcal{M}, w \models \langle \mathcal{E}, e \rangle \varphi \text{ iff } \mathcal{M}, w \models \mathsf{pre}(e) \text{ and } \mathcal{M} \otimes \mathcal{E}, (w, e) \models \varphi$$

[11] Of course, we must assume that she knows precisely *when* Charles will meet with Bob.

Example 2.9 (Public Announcement Logic) The **public announcement** of a formula $\varphi \in \mathcal{L}_{EL}$ is the event model $\mathcal{E}_\varphi = \langle \{e\}, \{\longrightarrow_i\}_{i\in\mathcal{A}}, \mathsf{pre}\rangle$ where for each $i \in \mathcal{A}$, $e \longrightarrow_i e$ and $\mathsf{pre}(e) = \varphi$ (see Plaza, 2007; Gerbrandy, 1999). As the reader is invited to verify, the product update of an epistemic model $\mathcal{M}$ with a public announcement model $\mathcal{E}_\varphi$ is the submodel of $\mathcal{M}$ containing all the states that satisfy φ. In this case, the DEL modality $\langle \mathcal{E}_\varphi, e\rangle$ will be denoted $\langle \varphi \rangle$. Henceforth, $\mathcal{L}_{PAL}$ will denote this language.

2.3 Comparing DEL and ETL

Both ETL and DEL are logical frameworks that describe the flow of information in a social interactive situation. For instance the *broadcasts* studied by van der Meyden (1996) and Lomuscio et al. (2000) are essentially the *public announcements* of Example 2.9. So, it is natural to ask how these two frameworks are related (cf. question 1 from the Introduction). Different logical frameworks, such as DEL and ETL, can be compared along many different dimensions. One key way to compare two different logical frameworks focuses on their *expressivity.* In order to show that one logic is at least as expressive as another logic, there are two main tasks to be carried out:

1. One has to establish a relation between the models of the two logics so that if we are given a model from the one logic, we can construct a corresponding model for the other logic;

2. One has to provide a formal translation so that if we are given a formula in the one formal language, we can produce a formula in the other with the same meaning.

Connections[12] between DEL and ETL along these lines have been worked out in detail by van Benthem and Pacuit (2006) and van Benthem et al. (2008).

The key observation is that by repeatedly updating an epistemic model with event models, the machinery of DEL (i.e., Definition 2.8) in effect creates ETL models. Note that an ETL model contains not only a description of how the agents' information changes over time, but also "protocol information" describing *when* each event *can* be performed[13]. Details of this comparison can be found in (van Benthem et al., 2008). Instead we identify the properties present in all DEL-generated ETL models. These properties have been discussed elsewhere (cf. Fagin et al., 1995; Bonanno, 2004), but can also be seen as coming out of the definition of product update (Definition 2.8).

[12]The first formal connection was established by Gerbrandy (1999, Section 5.3).

[13]The *preconditions of DEL* also encode protocol information of a 'local' character, and hence they can do some of the work of global protocols, as has been pointed out by van Benthem (2006).

Definition 2.10 (Synchronicity, Perfect Recall, Uniform No Miracles) Let $\mathcal{H} = \langle \Sigma, \mathsf{H}, \{\sim_i\}_{i \in \mathcal{A}}, V \rangle$ be an ETL model. $\mathcal{H}$ satisfies:

- **Sychronicity** iff for all $h, h' \in \mathsf{H}$, if $h \sim_i h'$ then $\mathsf{len}(h) = \mathsf{len}(h')$ ($\mathsf{len}(h)$ is the number of events in h).

- **Perfect Recall** iff for all $h, h' \in \mathsf{H}$, $e, e' \in \Sigma$ with $he, h'e' \in \mathsf{H}$, if $he \sim_i h'e'$, then $h \sim_i h'$

- **Uniform No Miracles** iff for all $h, h' \in \mathsf{H}$, $e, e' \in \Sigma$ with $he, h'e' \in \mathsf{H}$, if there are $h'', h''' \in \mathsf{H}$ with $h''e, h'''e' \in \mathsf{H}$ such that $h''e \sim_i h'''e'$ and $h \sim_i h'$, then $he \sim_i h'e'$. $\triangleleft$

Note that Definition 2.10 are properties of ETL *frames*. Already with these properties we can say something about how to relate the two frameworks. Suppose that $\mathcal{H}$ is an ETL frame satisfying the properties in Definition 2.10. We can easily read off an epistemic *frame* (i.e., a set of states W and relations R_i for each agent $i \in \mathcal{A}$ on W) to serve as the initial model (let the histories of length 1 be the states and simply copy the uncertainty relations). Furthermore, we can define a "DEL-like" protocol $\mathsf{P}_{\mathcal{H}}$ consisting of sequences of event models where the precondition function assigns to the primitive events *sets of finite histories*. Intuitively, if e is a primitive event (i.e., a state in an event model), then $\mathsf{pre}(e)$ is the set of histories where e can "be performed". Thus, we have a comparison of the two frameworks at the level of frames provided we work with a modified definition of an event model. However, the comparison is between *models*, so we need additional properties. In particular, at each level of the ETL model we will need to specify a *formula* of $\mathcal{L}_{EL}$ as a pre-condition for each primitive event e (cf. Definition 2.8). As usual, this requires that the set of histories preceding an event e be *bisimulation-closed* (cf. Blackburn et al., 2002, for a discussion of the notion of bisimulation). One final assumption that propositional variables do not change their truth value along a fixed history is needed since we are assuming that product update does not change the ground facts (although see the discussion in the next section about *factual change*). Consult (van Benthem et al., 2008, Theorem 1) for the details of the proof that any ETL model with the properties discussed above is generated from an initial epistemic model by a *DEL protocol* (i.e., a sequence of event models).

This technical result and discussion illustrates how DEL product update (Definition 2.8) may be used to generate interesting ETL frames and describes the observational powers of the agents presupposed in the DEL setting. Of course, this is not the only way to compare DEL and ETL. We can also we can also draw distinctions and comparisons by focusing on technical properties such as axiomatization and/or complexity results.

2.3.1 Axiomatizations

Axiomatizations in both DEL and ETL frameworks have been extensively studied. Both take as a starting point standard axiomatizations of epistemic logic (cf. Fagin et al., 1995; Blackburn et al., 2002). This short section reports on some of these results and highlights some of the important technical issues.

A sound and complete axiomatization of a number of different classes of ETL frames under the assumptions discussed in the previous section can be found in Halpern et al. (2004). Without assumptions about the observational powers of the agents (cf. Definition 2.10), such axiomatizations involve a straightforward *fusion* of appropriate axiomatizations of epistemic logic and temporal logic (See Kurucz, 2006, Section 3.2, for an extended discussion of this). It becomes much more interesting when there are assumptions connecting knowledge and time. For example, assuming an ETL frame satisfies perfect recall validates the following axiom scheme:

$$K_i\langle e\rangle\varphi \rightarrow \langle e\rangle K_i\varphi$$

For if agent i knows (at the current moment) that φ will be true at the next moment (i.e., after event e) then, since i has perfect recall, i cannot lose this piece of information. Therefore, at the next moment (after event e) agent i will know φ.[14]

There are three parameters that govern axiomatization results in the ETL framework. The first is the expressiveness of the language (i.e., does the language include a common knowledge operator? an *arbitrary future* operator? a past operator?). The second is structural conditions on the ETL frames (i.e., is the ETL frame a single tree with a unique root? finitely branching?). Finally, the third parameter is the assumptions made about the observational powers of the agents (i.e., do the agents have perfect recall? do the agents agree on the time? do the agents satisfy the properties from Definition 2.10?). At one extreme, with at least two agents and languages containing common knowledge operators and arbitrary future operators, the validity problem over classes of ETL frames that satisfy perfect recall is Π^1_1-*complete* (see Halpern and Vardi, 1989; van Benthem and Pacuit, 2006, for proofs). Nonetheless, many classes of ETL frames (under different combinations of assumptions about the observational power of the agents) in a variety of modal languages (typically without a

[14]Interestingly, van der Meyden (1994) showed in languages with an "until" operator ($\varphi U\psi$ meaning there is a point in the future satisfying ψ and that φ is true at every moment until that point) adding only this axiom to an epistemic and temporal logic is *not* complete for ETL frames with perfect recall. What is needed is the more complex axiom scheme: $K_i\varphi_1 \wedge N(K_i\varphi_2 \wedge \neg K_i\varphi_3) \rightarrow L_i((K_i\varphi_1)\ U[(K_i\varphi_2)U\neg\varphi_3])$, where '$N$' is the *next-time operator* — after any event e (cf. Halpern et al., 2004).

common knowledge operator or in a restricted temporal language) can be found in (Halpern et al., 2004; French et al., 2004; van der Meyden and Wong, 2003). Despite many different axiomatization and non-axiomatization results, it is fair to say that no general picture has yet emerged (although see the discussion by van Benthem and Pacuit (2006) and Kurucz (2006) for some first steps in this direction).

In contrast, the so-called *reduction axioms* have proven an invaluable method for providing sound and complete axiomatizations in the DEL framework. They were first used by Plaza to prove completeness for *public announcement logic* (see Example 2.9). This is the logic where the only event models are those where there is one primitive event, and the uncertainty relation for all agents is the universal relation. A public announcement can then be referred to simply by its precondition, resulting in formulas of the form $\langle\varphi\rangle\psi$. The following are reduction axioms for PAL:

$$\langle\varphi\rangle p \leftrightarrow (\varphi \wedge p)$$
$$\langle\varphi\rangle\neg\psi \leftrightarrow (\varphi \wedge \neg\langle\varphi\rangle\psi)$$
$$\langle\varphi\rangle(\psi \wedge \chi) \leftrightarrow (\langle\varphi\rangle\psi \wedge \langle\varphi\rangle\chi)$$
$$\langle\varphi\rangle K_i\psi \leftrightarrow (\varphi \wedge K_i\langle\varphi\rangle\psi)$$

These are reduction axioms in the sense that going from left to right either the number of announcement operators is reduced or the complexity of the formulas within the scope of announcement operators is reduced. In the first axiom we see that an announcement has been eliminated. In the second axiom we see that the announcement operator and the negation have switched place. In the third we see an announcement of a conjunction on the left and a conjunction of announcements on the right. In the fourth axiom we also see that the announcement and the epistemic operator have switched place. The reduction axioms for event models in general are a straightforward generalization of the axioms above.

The reduction axioms for PAL provide an insightful syntactic analysis of announcements which complements the semantic analysis. In a sense, the reduction axioms describe the effect of an announcement in terms of what is true before the announcement. By relating pre- and postconditions for each logical operator, the reduction axioms completely characterize the announcement operator.

In the completeness proof for PAL the reduction axioms play an essential role. Given a formula containing an announcement operator, one can completely eliminate the announcement by repeatedly applying the reduction axioms. In this way one produces a formula of epistemic logic. By adding the appropriate reduction axioms to a complete axiomatization for epistemic logic, it is straightforward to show the resulting proof system is complete in the following manner. Suppose a formula φ is a semantic tautology. By applying the reduction axioms one obtains a *provably* equivalent formula φ'. This is a semantic tautology in the language of epistemic logic. By completeness of the proof system for epistemic logic, there must be a proof of φ', and since φ and φ' are provably equivalent one can construct a proof of φ. This technique for proving completeness is considered so elegant that many have adopted it (Plaza, 2007; Gerbrandy, 1999; Baltag et al., 1998b; Herzig et al., 2000; Kooi, 2003a; Renardel de Lavalette, 2004; Kooi and van Benthem, 2004; van Eijck, 2004b,a; Ruan, 2004; van Benthem et al., 2006; van Benthem, 2007; van Benthem and Liu, 2007; Kooi, 2007; van Benthem and Ikegami, 2008).

Reduction axioms are not only useful in providing a syntactic analysis of updates and for proving completeness, they also show that the language containing the update is just as expressive as the language without it. So the results mentioned above are also expressivity results showing the the language of PAL is no more expressive than the language of epistemic logic. Yet Lutz (2006) has shown that in the case of PAL at least, the language is more succinct than the language of epistemic logic (there is a formula scheme in PAL such that every equivalent formula scheme in epistemic logic is exponentially longer). This suggests that PAL describes announcements at an appropriate level of abstraction.

When a logical language becomes strictly more expressive by adding dynamic operators, reduction axioms are not available. Adding public announcement operators to epistemic logic with common knowledge is such a case. It was shown by Baltag et al. (1998b) that the language of epistemic logic with common knowledge and public announcements is more expressive than epistemic logic with common knowledge. Therefore a reduction axiom for formulas of the form $[\varphi]C_{\mathcal{B}}\psi$ does not exist. Baltag, Moss and Solecki also showed that adding private announcements to epistemic logic with common knowledge adds expressivity. Renne (2007b) showed that the expressivity of these two logics is incomparable. In cases where adding dynamic operators strictly increases the expressivity of the language a completeness proof using reduction axioms is not available and a complete proof system is harder to obtain.

3 Extensions, connections and applications

The previous section introduced two different logical frameworks that describe how an agent's information evolves through observation when interacting with other agents. The results discussed in Section 2.3 provide a concrete answer to question 1 from the Introduction (how should we compare two logical frameworks addressing the same aspect of rational agency). But what about the other two questions? Here, especially regarding question 3 (how the logics of rational interaction contribute to broader discussions on rational agency), we cannot point to any concrete results as answers to these questions. Rather, this section turns to several extensions of these logics of rational interaction, as well as connections with other fields, and some applications.

To keep this survey at a manageable length we will not be able to provide anything approaching a complete survey of all extensions and applications of ETL and DEL. See (Fagin et al., 1995) for a textbook presentation of a number of extensions and applications of ETL, and (van Ditmarsch et al., 2007; van Benthem, 2008) for applications and extensions of DEL. The topics discussed below where chosen because they are representative of current research directions and issues addressed in this volume. We start by briefly discussing a few other logics frameworks that can broadly be categorized as "logics of rational interaction".

Propositional dynamic logic The language of DEL is set up similarly to the languages of *propositional dynamic logic* (PDL). Two distinct classes of formulas and programs (i.e. updates) are defined. Therefore it would be natural to include the Kleene star for iteration as well, as it allows one to express things such as no matter how matter how many times it is announced that φ, it will not become common knowledge that ψ as $[\varphi^*]\neg C_{\mathcal{B}}\psi$. Miller and Moss (2005) showed that the satisfiability problem for PAL with iteration is undecidable[15]. Hence, DEL with iteration has the same problem.

Still, DEL can be embedded in PDL, i.e. for each formula in DEL, there is a formula in PDL which is equivalent to it (van Eijck, 2004b). In van Eijck's approach, PDL formulas are read epistemically, for instance $[i]\varphi$ is read as agent i knows that φ. Another link between DEL and PDL is developed by Aucher and Herzig (2007) where $[e]\varphi$ is read as after event e it is the case that φ, i.e. e is taken to be an event from an event model. Separate modalities for agents are added to PDL (just as was done by van Benthem, 2001). With the addition of a converse operator, this logic can express properties of event models.

Belief revision The ground breaking paper by Alchourrón et al. (1985) put information change prominently on the agenda of philosophical logic. Their approach, abbreviated as AGM, focuses on what to do when receiving (and accepting) information not in accordance with the agent's theory of the current state of affairs. This led to a stream of publications in an area nowadays called *belief revision.*

[15] In fact, they show that the validity problem for public announcement logic with iteration is *highly undecidable* (Π^1_1-complete). In light of the translation between the DEL framework and the ETL framework discussed in Section 2.3, this is related to classic results of Halpern and Vardi (1989) showing that the validity problem for ETL frames that satisfy perfect recall and no miracles in certain modal languages is Π^1_1-complete. See (van Benthem et al., 2008, Section 6.1) for an extended discussion of this relationship.

Indeed, there are by now many different approaches to modeling how agents (should) change their beliefs in the presence of new (and trusted) information. Shoham and Leyton-Brown (2009, Section 14.2) discuss many of these different approaches (including *nonmonotonic* consequence relations, default logics and probabilistic frameworks). Rather than discussing this expanding literature, we point to contributions that are most relevant to the logics we discussed in Section 2. These can be roughly divided into two categories. The first are ETL-style logics that describe how an agent's beliefs change through time (see, for example, Friedman and Halpern (1997, 1999); Bonanno (2007) and references therein). The second category can be described as dynamic modal logics of belief revision. Building on a suggestion of van Benthem (1989), de Rijke (1994) took some first steps to develop a dynamic modal logic of belief revision. This led to the development of *dynamic doxastic logic* (see Segerberg's contribution to this volume and (Lindström and Segerberg, 2007) for an overview of this approach). Recent work has provided a multi-agent perspective with a number of DEL-style logics of belief revision[16] (see, for example, Aucher, 2003; van Ditmarsch, 2005; Cantwell, 2006; van Benthem, 2007; Baltag and Smets, 2008). Building on the results discussed in Section 2.3, van Benthem and Dégremont (2009) formally compare the ETL-style and DEL-style logics of belief change.

Probability logic Probability theory provides a quantitative analysis of information. Rather than a proposition being known or unknown, its degree of certainty is represented by a number. For instance, the chance that the queen of hearts is drawn from a shuffled ordinary deck of cards is 1/52. The are many connections between probability and logic, including epistemic logic (see (Halpern, 2003) for a textbook presentation). The connection with logics discussed in this survey becomes apparent by noting that a Bayesian update resembles the public announcement of Example 2.9. It is therefore is quite natural to combine probability logics and dynamic epistemic logics (cf. van Benthem, 2003; Kooi, 2003b; Aucher, 2007; van Benthem et al., 2008; Sack, 2008a).

Situation calculus Reasoning about actions is important area of research in artificial intelligence and the *situation calculus* of McCarthy and Hayes (1969) is one of the most influential approaches (see (Reiter, 2001) for a textbook on the subject). The situation calculus is a fragment of second order logic that can describe many situations and how situations change due to actions. Typical examples involve robots moving blocks. Comparisons with ETL-style logics is relatively straightforward since the situation calculus can express most epistemic

[16]Closely related are the dynamic logics of preferences discussed by van Benthem and Liu (2007).

and temporal modalities. The comparison with DEL is more subtle. The link between the two formalisms was established by van Ditmarsch et al. (2007), who use situation calculus and DEL to approach the *frame problem*[17].

The comparisons between the different logical frameworks discussed above and in Section 2.3 suggest a number of *extensions* to the basic DEL and ETL frameworks. For example, the results of (van Benthem et al., 2008) suggest adding temporal operators, such as a past-time operator, to DEL (cf. Sack, 2008b; Hoshi and Yap, 2009). Again we do not have the space to cover all extensions to the logics of rational interaction (see van Benthem, 2008, for an extended discussion), so we focus on a few key research avenues.

Factual change Although DEL is mainly used to model information change due to communication, comparisons with ETL and the situation calculus suggest that it may be convenient to model situations where the bare facts of the world do change. This was already foreshadowed in the CWI technical report version of the paper by Baltag et al. (1998b). The first DEL with factual change was proposed by Bleeker and van Eijck (2000), where multiple propositional letters can change simultaneously. Baltag (2002) considers DEL with 'flip' actions, which changes the extension of a propositional letter p to its complement. Renardel de Lavalette (2004) uses operators $p := \varphi$ which changes the extension of p to the extension of φ and applies the same idea to agents where $i := \pi$ changes the accessibility relation of i to that of π. Van Ditmarsch, van der Hoek, and Kooi 2005 provide a logic with such operators and public announcements. Van Eijck 2004a showed that DEL with simultaneous factual changes in event models can be reduced to PDL. Factual change has been further studied in (van Benthem et al., 2006; Herzig and de Lima, 2006; Kooi, 2007; van Ditmarsch and Kooi, 2008).

Logics of rational agency The logics discussed in this survey focus primarily on information change. But logics have also been developed to reason more broadly about rational agency. Indeed, there are now many different "logics of rational agency" (see van der Hoek and Wooldridge, 2003; Meyer and Veltman, 2007; Horty, 2001, for a discussion and pointers to the relevant literature) that not only focus on describing various informational and/or motivational attitudes but also explicating their relationships. An overarching theme in many of these papers is that during a social interaction, an agent's "knowledge" and "beliefs" both influence *and* shape the *social* events. The following example (taken from Pacuit et al., 2006) illustrates this point.

[17] Both (Reiter, 2001) and (McCarthy and Hayes, 1969) discuss this classic problem of AI.

Example 1: Uma is a physician whose neighbour is ill. Uma does not know and has not been informed. Uma has no obligation (as yet) to treat the neighbour.

Example 2: Uma is a physician whose neighbour Sam is ill. The neighbour's daughter Ann comes to Uma's house and tells her. Now Uma does have an obligation to treat Sam, or perhaps call in an ambulance or a specialist.

Example 3: Mary is a patient in St. Gibson's hospital. Mary is having a heart attack. The caveat which applied in case 1) does not apply here. The hospital cannot plead ignorance, but rather it has an obligation to *be aware* of Mary's condition at all times and to provide emergency treatment as appropriate.

In all the cases we mentioned above, the issue of an obligation arises. This obligation is circumstantial in the sense that in other situations, the obligation might not apply. If Sam is ill, Uma needs to know that he is ill, and the nature of the illness, but not where Sam went to school. Thus an agent's obligations are often dependent on what the agent knows, and indeed one cannot reasonably be expected to respond to a problem if one is not aware of its existence. This, in turn, creates a secondary obligation on Ann to inform Uma that her father is ill.

Based on the logical framework discussed in Section 2.1 and Horty (2001), Pacuit et al. (2006) develop a logical framework that formalizes the reasoning of Uma and Ann in the above examples. It is argued that this reasoning is shaped by the assumption that Uma and Ann's preferences are aligned (i.e., both want Sam to get better). For example, Ann will not be under any obligation to tell Uma that her father is ill, if Ann justifiably believes that Uma would not treat her father even if she knew of his illness. Thus, in order for Ann to *know* that she has an obligation to tell Uma about her father's illness, Ann must *know* that "Uma will, in fact, treat her father (in a reasonable amount of time) upon learning of his illness". More formally, in all the histories that Ann currently considers possible, the event where her father is treated for his illness is always preceded by the event where she tells Uma about his illness. That is, the histories where Uma learns of Sam's illness but does not treat him are not part of the protocol. Similar reasoning is needed for Uma to derive that she has an obligation to treat Sam. Obviously, if Uma has a good reason to believe that Ann always lies about her father being ill, then she is under no obligation to treat Sam. See (Pacuit et al., 2006) for a formal treatment of these examples.

Inference logic Besides information about the world and the discourse information, there is a third kind of information that plays a role in interaction, namely information derived from (logical) *inference*. What conclusions is one allowed to draw from a set of premises, and how is the process of inference carried out? A number of logical frameworks have been developed that explicitly reason about such inferential steps (Duc, 1997, 2001; Jago, 2006). Frameworks that extend ETL style logics include (Ågotnes and Alechina, 2007; Alechina et al., 2009). Combinations of "inference logics" and DEL have been put forward by van Benthem (2008) and Velazquez-Quesada (2009).

Justification logic Justification logic is an epistemic logic where explicit reasons for knowledge are represented. A formula $t : \varphi$ is intended to mean "the agent knows φ for reason t". It was introduced by Artemov and Nogina (2005) based on their work on explicit provability logic. Renne (2007a) added public announcements to this logic and proved a number of expressivity results. See also Renne's contribution to this volume for an extended discussion.

Finally, we conclude this section with a brief discussion of a number of key applications of the logics of rational interaction discussed in Section 2.

Puzzles and paradoxes The development of DEL in particular was fueled by a number of puzzles and paradoxes. These did not only function as an inspiration, but also as a touchstone for DEL. Both Plaza and Gerbrandy analyzed the Muddy Children puzzle using PAL. Plaza also treats the Sum and Product puzzle. Another example of a puzzle where a specification of the solution in DEL offers a method of evaluating solutions suggested in the literature is the Russian cards problems (van Ditmarsch, 2003).

Although some of these puzzles are also found in recreational mathematics, some have serious philosophical repercussions. The hangman paradox, or unexpected examination paradox was first analyzed using PAL by Gerbrandy (1999), 2007. A judge sentences a prisoner to death and says that he will be hanged next week but that the day of the execution will come as a surprise. The prisoner then reasons as follows. If the execution were on Friday, then I would know on Thursday evening that this is so, and the day of the execution would not be a surprise. Therefore the execution cannot take place on Friday. So, Thursday is the last possible day for the execution. By the same reasoning as before the prisoner concludes that the execution cannot take place on Thursday either, and so he continues eliminating all days of the week. The prisoner cheerfully infers that the execution cannot take place at all. To his great surprise he is executed on Tuesday.

The central point of Gerbrandy's analysis is that the announcement of the judge maybe an *unsuccessful update*. That is, a formula that becomes false by its announcement. This phenomenon also occurs in Update Semantics, when an update system does not satisfy the condition of *idempotence* (cf. Veltman (1996)).

A literary example of this phenomenon is found in the fairy tale *Rumpelstilzchen* (Grimm and Grimm, 1857), where a goblin who sings the following song (our translation below[18]):

> Heute back ich, morgen brau ich,
> übermorgen hol ich der Königin ihr Kind;
> ach, wie gut dass niemand weiß,
> dass ich Rumpelstilzchen heiß!
>
> Today I bake, tomorrow I brew,
> The day after tomorrow I will fetch the queen's child;
> Oh, it's good that nobody knows,
> that I'm called Rumpelstilzchen.

In the fairy tale his song is overheard and therefore it is no longer true that nobody knows the goblin's name. Thus uttering a true statement, can make that statement itself false! In PAL such statements are called unsuccessful updates. A successful update is a formula φ such that $[\varphi]\varphi$ is a tautology. An update is unsuccessful if it is not successful. The announcement of the judge is an unsuccessful update, i.e. the judge may ruin the surprise by saying that the day of the execution will come as a surprise. van Ditmarsch and Kooi (2006) discuss this phenomenon in a number of contexts.

The formula $p \wedge \neg K_a p$ is a typical example of an unsuccessful update, which play a role in the *Fitch paradox* or *knowability paradox*. If one accepts that all truths are knowable, then if p is true but unknown it should be knowable that p is true but unknown. This leads to a contradiction. Using DEL the paradox was analyzed by van Benthem (2004). This led to the development of arbitrary public announcement logic, where formulas $\Diamond\varphi$ occur, which are read as there is some announcement such that afterwards φ is true (Balbiani et al., 2007).

Game theory Any (formal) model that addresses issues of (practical) rationality needs to account for the possibility of conflicting *goals* of the different agents. Starting from the work of Ramsey (1926); de Finetti (1937); von Neumann and Morgenstern (1944) and Savage (1954), the mathematical analyses provided by decision and game theorists have generated many important insights about such *strategic* interactive situations. Indeed, in their classic text, von Neumann and Morgenstern explain that they want "to find the mathematically complete principles which define 'rational behavior' for the participants" (von Neumann and Morgenstern, 1944, p. 31). Nonetheless, many foundational questions remain

[18]Regrettably the English translations we consulted do not contain this phenomenon.

open. These questions are not mathematical in nature but involve the meaning of the fundamental concepts employed in the mathematical analyses.

Building on seminal work by John Harsanyi (1967) on incomplete information games[19] and Robert Aumann (1999, 1976) introducing common knowledge to game theory, many researchers have forcefully argued that the basic mathematical model of a "game-theoretic situation" should be extended with an explicit representation of the players' relevant informational attitudes[20] (following Harsanyi (1967), this parameter is called a player's *type*. See, for example, (Brandenburger, 2007; Bonanno and Battigalli, 1999) for an extended discussion). A central concern in this literature is the players' attitude towards statements about the *rationality of the other players* and whether such statements can be revised during the course of a (dynamic) game[21]. Although there is considerable disagreement over the precise formulation, it is generally assumed that such statements about the rationality of the other players are more *entrenched* than, for example, higher-order beliefs about the *types* of the other players.

One lesson to take away from this discussion is that game-theoretic analyses of multiagent strategic situations should be embedded in a larger framework that describes how the players' (hard and soft) information evolves over time. The logical systems discussed in this paper focus on precisely this issue (cf. Section 2). Thus, these frameworks complement the game theoretic models described above by focusing on how a player's type may evolve over time and how a player may change types during the course of a game. Much more can be said on this general topic, but we will not go into this here (see van der Hoek and Pauly, 2006, for discussion along these lines and pointers to the relevant literature).

Security One of the recent application areas of logics of rational interaction is security, especially authentication and privacy. Both DEL and ETL frameworks have been used to verify that security protocols meet their specification (Bleeker and van Eijck, 2000; Hommersom et al., 2004; Dechesne and Wang, 2007; Halpern and Pucella, 2003; Ramanujam and Suresh, 2005; van der Meyden and Wilke, 2007). A topic of special interest is so-called *zero-knowledge protocols*. These are security protocols where security does not depend on the bounded computational

[19]That is, situations in which the *structure* of the game is not common knowledge. For example, games where players may be uncertain about their own available actions and preferences and/or the available actions and preferences of the other players. This should be contrasted with *imperfect information* games where players may receive different information during the course of the game. See (Myerson, 2004) for a recent discussion of Harsanyi's classic paper.

[20]Typically this means the players *first-order* beliefs about the available choices of the other players, the players beliefs about the other players beliefs about these first-order beliefs, and so on *ad infinitum*.

[21]Or, in the case of a one-shot strategic game, whether such statements can be revised during the players' initial period of deliberation.

resources of the participants in the protocol. An example is the solution of the *Russian cards problem* (see van Ditmarsch, 2003). This problem has been modeled in DEL and formal checked by the model checker DEMO, developed by van Eijck (2005), (van Ditmarsch et al., 2006). Other typical problems found in the literature on security have been analyzed, including the dining cryptographers problem (van Eijck and Orzan, 2005; van der Meyden, 2007).

4 Conclusion: towards a unified account of rational interaction

There is a multitude of logics that aim to model different aspects of rational interaction. Often, for one and the same aspect there are numerous approaches. Do these alternative approaches represent radically different conceptual frameworks, or are the same concepts represented in different guises? In our case, how should we compare two different frameworks that model how an agent's information changes through interaction with other agents and the environment. This was exactly what we described in Section 2 for ETL and DEL. So these observations point to one "coherent" account of rational interaction. Yet this is not the whole story of rational interaction.

Agents are faced with many diverse tasks as they interact with the environment and one another. At certain moments, agents must *react* to the (perhaps surprising) events they observe while at other moments they must be *proactive* and choose to perform a specific action. One central underlying assumption is that rational agents obtain what they want via the implementation of (successful) *plans* (cf. Bratman, 1987). And this implementation often requires, among other things, representation of various informational attitudes of the other agents involved in the social interaction. As illustrated by the discussion of Example 2.1 in Section 2, in social situations there are many (sometimes competing) *sources* of information for these attitudes: for example, the type of "communicatory event" (public announcement, private announcement, etc), the disposition of the other participants (liars, truth-tellers, etc.) and other implicit assumptions about the protocol information (reducing the number of possible histories). This naturally leads to different notions of "knowledge" and "belief" that drive social interaction.

The conclusion is that a comprehensive account of rational interaction cannot be isolated from other aspects of rational agency and social interaction. This paper presented some recent work which points to such a comprehensive account. Once the technical results of Section 2.3 are in place, the two major current views of how information evolves through social interaction can be seen as *complementary*. This opens the door to merging these two perspectives (cf. van Benthem et al., 2008, Section 4) which will, in turn, lead to a more diversified account of the reasoning and dynamic processes that govern social interaction.

References

Ågotnes, T. and N. Alechina (2007). The dynamics of syntactic knowledge. *Journal of Logic and Computation 17*(1), 83–116.

Alchourrón, C. E., 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.

Alechina, N., B. Logan, H. N. Nguyen, and A. Rakib1 (2009). Verifying time, memory and communication bounds in systems of reasoning agents. *Synthese (Knowledge, Rationality & Action)* (forthcoming).

Artemov, S. and E. Nogina (2005). Introducing justification into epistemic logic. *Journal of Logic and Computation 15*(6), 1059–1073.

Aucher, G. (2003). A combined system for update logic and belief revision. Master's thesis, ILLC, University of Amsterdam, Amsterdam.

Aucher, G. (2007). Interpreting an action from what we perceive and what we expect. *Journal of Applied Non-Classical Logics 17*(1), 9–38.

Aucher, G. and A. Herzig (2007). From DEL to EDL: exploring the power of converse events. In K. Mellouli (Ed.), *Proceedings of the 9th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, Volume 4724 of *Lecture Notes in Computer Science*, Berlin, pp. 199–209. Springer-Verlag.

Aumann, R. (1976). Agreeing to disagree. *Annals of Statistics 4*, 1236 — 1239.

Aumann, R. (1999). Interactive epistemology I: Knowledge. *International Journal of Game Theory 28*, 263–300.

Balbiani, P., A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. de Lima (2007). What can we achieve by arbitrary announcements? – a dynamic take on Fitch's knowability. In D. Samet (Ed.), *Proceedings of TARK 2007*, pp. 42–51.

Baltag, A. (2002). A logic for suspicious players: epistemic actions and belief-updates in games. *Bulletin of Economic Research 54*(1), 1–45.

Baltag, A., L. Moss, and S. Solecki (1998a). The logic of common knowledge, public announcements 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., L. S. Moss, and S. Solecki (1998b). 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 (2008). A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge (Eds.), *Proceedings of LOFT 2007*, Volume 3 of *Texts in Logic and Games*, Amsterdam. Amsterdam University Press. to appear.

Blackburn, P., M. de Rijke, and Y. Venema (2002). *Modal Logic.* Cambridge: Cambridge University Press.

Bleeker, A. and J. van Eijck (2000). The epistemics of encryption. Technical Report INS-R0019, CWI, Amsterdam. Available from `http://db.cwi.nl/rapporten/`.

Bonanno, G. (2004). Memory and perfect recall in extensive games. *Games and Economic Behaviour 47*, 237 – 256.

Bonanno, G. (2007). Axiomatic characterization of AGM belief revision in a temporal logic. *Artificial Intelligence 171*, 144 – 160.

Bonanno, G. and P. Battigalli (1999, June). Recent results on belief, knowledge and the epistemic foundations of game theory. *Research in Economics 53*(2), 149–225.

Brandenburger, A. (2007). The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory 35*, 465–492.

Bratman, M. (1987). *Intention, Plans and Practical Reason.* London: Harvard University Press.

Bratman, M. (2007). *Structures of Agency.* Oxford University Press.

Cantwell, J. (2006). A formal model of multi-agent belief-interaction. *Journal of Logic, Language and Information 15*(4), 303–329.

Cartwright, N. (2007). *Hunting Causes and Using Them: Approaches in Philosophy and Economics.* Cambridge University Press.

de Finetti, B. (1937). La prevision: Ses lois logiques, ses sources subjectives. In *Annales de l'Institut Henri Poincare 7*, pp. 1–68. Paris. Translated into English by Henry E. Kyburg Jr., Foresight: Its Logical Laws, its Subjective Sources. In Henry E. Kyburg Jr. and Howard E. Smokler (1964, Eds.), Studies in Subjective Probability, 53-118, Wiley, New York.

de Rijke, M. (1994). Meeting some neighbours. In J. van Eijck and A. Visser (Eds.), *Logic and information flow*, Cambridge, Massachusetts, pp. 170–195. The MIT Press.

Dechesne, F. and Y. Wang (2007). Dynamic epistemic verification of security protocols: framework and case study. In J. van Benthem, S. Ju, and F. Veltman (Eds.), *A meeting of the minds, Proceedings of the Workshop on Logic, Rationality and Interaction, Beijing, 2007*, Volume 8 of *Texts in Computer Science*, London, pp. 129–143. College Publications.

Duc, H. N. (1997). Reasoning about rational, but not logically omniscient, agents. *Journal of Logic and Computation 7*(5), 633–648.

Duc, H. N. (2001). *Resource-Bounded Reasoning about Knowledge*. Ph. D. thesis, Universität Leipzig.

Fagin, R., J. Halpern, Y. Moses, and M. Vardi (1995). *Reasoning about Knowledge*. Boston: The MIT Press.

French, T., R. van der Meyden, and M. Reynolds (2004). Axioms for logics of knowledge and past time: Synchrony and unique initial states. In *Advances in Modal Logic*, pp. 53–72.

Friedman, N. and J. Halpern (1997). Modeling belief in dynamic systems. part I: Foundations. *Artificial Intelligence 95*(2), 257 – 316.

Friedman, N. and J. Halpern (1999). Modeling belief in dynamic systems. part II: Revision and update. *Journal of AI Research 10*, 117 – 167.

Gerbrandy, J. (1999). *Bisimulations on Planet Kripke*. Ph. D. thesis, Institute for Logic, Language and Computation (DS-1999-01).

Gerbrandy, J. D. (2007). The surprise examination in dynamic epistemic logic. *Synthese 155*(1), 21–33.

Grimm, J. and W. Grimm (1857). *Kinder- und Hausmärchen*. Göttingen.

Halpern, J. (1999). Set-theoretic completeness for epistemic and conditional logic. *Annals of Mathematics and Artificial Intelligence 26*, 1 – 27.

Halpern, J. and R. Pucella (2003). Modeling adversaries in a logic for security protocol analysis. In *Formal Aspects of Security*.

Halpern, J., R. van der Meyden, and M. Vardi (2004). Complete axiomatizations for reasoning about knowledge and time. *SIAM Journal of Computing 33*(2), 674 – 703.

Halpern, J. and M. Vardi (1989). The complexity of reasoning about knowledge and time. *J. Computer and System Sciences 38*, 195 – 237.

Halpern, J. Y. (2003). *Reasoning about Uncertainty.* Cambridge, Massachusetts: The MIT Press.

Harsanyi, J. (1967). Games with incomplete informations played by 'bayesian' players. *Management Science 14*, 159–182, 320–334, 486–502.

Herzig, A. and T. de Lima (2006). Epistemic actions and ontic actions: a unified logical framework. In J. Simão Sichman, H. Coelho, and S. Oliveira Rezende (Eds.), *Advances in Artificial Intelligence - IBERAMIA-SBIA 2006*, Volume 4140 of *Lecture Notes in Computer Science*, Berlin, pp. 409–418. Springer-Verlag.

Herzig, A., J. Lang, and T. Polacsek (2000). A modal logic for epistemic tests. In W. Horn (Ed.), *Proceedings of ECAI 2000*, pp. 553–557.

Hodkinson, I. and M. Reynolds (2006). Temporal logic. In P. Blackburn, J. van Benthem, and F. Wolter (Eds.), *Handbook of Modal Logic*, Volume 3 of *Studies in Logic*, pp. 655 – 720. Elsevier.

Hommersom, A., J.-J. C. Meyer, and E. de Vink (2004). Update semantics of security protocols. *Synthese (Knowledge, Rationality & Action) 142*, 229–267.

Horty, J. (2001). *Agency and Deontic Logic.* Oxford University Press.

Hoshi, T. and A. Yap (2009). Dynamic epistemic logic with branching temporal structures. *Synthese (Knowledge, Rationality & Action)*, forthcoming.

Hyman, J. and H. Steward (Eds.) (2004). *Agency and Action.* Cambridge University Press.

Jago, M. (2006). *Logics for Resource-Bounded Agents.* Ph. D. thesis, University of Nottingham.

Kooi, B. (2003a). *Knowledge, Chance, and Change.* Ph. D. thesis, University of Groningen. ILLC Dissertation Series DS-2003-01.

Kooi, B. (2003b). Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information 12*(4), 381–408.

Kooi, B. (2007). Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics 17*(2), 231–253.

Kooi, B. and J. van Benthem (2004). Reduction axioms for epistemic actions. In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing (Eds.), *AiML-2004: Advances in Modal Logic*, Department of Computer Science, University of Manchester, Technical report series, UMCS-04-9-1, pp. 197–211.

Kurucz, A. (2006). Combining modal logics. In P. Blackburn, J. van Benthem, and F. Wolter (Eds.), *Handbook of Modal Logic*, Volume 3 of *Studies in Logic*, pp. 869 – 924. Elsevier.

Lindström, S. and K. Segerberg (2007). Modal logic and philosophy. In P. Blackburn, J. van Benthem, and F. Wolter (Eds.), *Handbook of Modal Logic*, Volume 3 of *Studies in Logic and Practical Reasoning*, pp. 1149–1214. Amsterdam: Elsevier.

Lomuscio, A. and M. Ryan (1997). On the relation between interpreted systems and kripke models. In *Proceedings of the AI97 Workshop on Theoretical and Practical Foundation of Intelligent Agents and Agent-Oriented Systems*, Volume LNCS 1441.

Lomuscio, A. R., R. van der Meyden, and M. Ryan (2000). Knowledge in multiagent systems: initial configurations and broadcasts. *ACM Transactions on Computational Logic 1*(2), 247–284.

Lutz, C. (2006). Complexity and succinctness of public announcement logic. In *Proceedings AAMAS-06: Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, Hakodate, Japan.

McCarthy, J. and P. J. Hayes (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence 4*, 463 – 502.

Meyer, J.-J. and F. Veltman (2007). Intelligent agents and common sense reasoning. In P. Blackburn, J. van Benthem, and F. Wolter (Eds.), *Handbook of Modal Logic*. Elsevier.

Miller, J. S. and L. S. Moss (2005). The undecidability of iterated modal relativization. *Studia Logica 79*(3), 373–407.

Myerson, R. (2004). Harsanyi's games with incomplete information. *Management Science 50*(12), 1818–1824.

Pacuit, E. (2007). Some comments on history based structures. *J. Applied Logic 5*(4), 613–624.

Pacuit, E. and R. Parikh (2006). Social interaction, knowledge and social software. In *Interactive Computation: The New Paradigm*, pp. 441 – 462. Springer.

Pacuit, E., R. Parikh, and E. Cogan (2006). The logic of knowledge based obligation. *Knowledge, Rationality and Action: A Subjournal of Synthese 149*(2), 311 – 341.

Parikh, R. (2003). Levels of knowledge, games, and group action. *Research in Economics 57*, 267 — 281.

Parikh, R. and R. Ramanujam (1985). Distributed processes and the logic of knowledge. In *Logic of Programs*, Volume 193 of *Lecture Notes in Computer Science*, pp. 256 – 268. Springer.

Plaza, J. (2007). Logics of public communications. *Synthese 158*(2), 165–179. This paper was originally published as Plaza, J. A. (1989). Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z.W. Ras (Eds.), *Proceedings of the fourth international symposium on methodologies for intelligent systems: Poster session program* (pp. 201–216). Publisher: Oak Ridge National Laboratory, ORNL/DSRD-24.

Ramanujam, R. and S. Suresh (2005). Deciding knowledge properties of security protocols. In *Proceedings of Theoretical Aspects of Rationality and Knowledge*, pp. 219–235.

Ramsey, F. (1926). Truth and probability. In R. Braithwaite (Ed.), *The Foundations of Mathematics and other Logical Essays*. Routledge.

Reiter, R. (2001). *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. Cambridge, Massachusetts: The MIT Press.

Renardel de Lavalette, G. R. (2004). Changing modalities. *Journal of Logic and Computation 14*(2), 253–278.

Renne, B. (2007a). Public communication in justification logic. Technical Report TR-2005025, CUNY Ph.D. Program in Computer Science.

Renne, B. (2007b). The relative expressivity of public and private communication in BMS logic. In J. van Benthem, S. Ju, and F. Veltman (Eds.), *A Meeting of the Minds: Proceedings of the Workshop on Logic, Rationality, and Interaction, Beijing, 2007*, Volume 8 of *Texts in Computer Science*, London, pp. 213–229. College Publications.

Ruan, J. (2004). Exploring the update universe. Master's thesis, ILLC, Amsterdam.

Sack, J. (2008a). Extending probabilistic dynamic epistemic logic. In *Proceedings of the Workshop on Logic and Intelligent Interaction.*

Sack, J. (2008b). Temporal languages for epistemic programs. *Journal of Logic, Language and Information 17*(2), 183–216.

Savage, L. (1954). *The Foundations of Statistics.* New York: Dover Publication.

Searle, J. (1985). *The Construction of Social Reality.* The Free Press.

Shoham, Y. and K. Leyton-Brown (2009). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.* Cambridge University Press.

van Benthem, J. (1989). Semantic parallels in natural language and computation. In H.-D. Ebbinghaus, J. Fernandez-Prida, M. Garrido, D. Lascar, and M. R. Artalejo (Eds.), *Logic Colloquium '87*, Amsterdam. North-Holland.

van Benthem, J. (2001). Games in dynamic-epistemic logic. *Bulletin of Economic Research 53*(4), 219–248.

van Benthem, J. (2003). Conditional probability meets update logic. *Journal of Logic, Language and Information 12*(4), 409–421.

van Benthem, J. (2004). What one may come to know. *Analysis 64*(2), 95–105.

van Benthem, J. (2005). Rational animals: What is 'KRA'? invited lecture Malaga ESSLLI Summer School 2006.

van Benthem, J. (2006). One is a lonely number: on the logic of communication. In Z. Chatzidakis, P. Koepke, and W. Pohlers (Eds.), *Logic Colloquium '02*, Volume 27 of *Lecture Notes in Logic*, pp. 96 – 129. ASL & A.K. Peters.

van Benthem, J. (2007). Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics 17*(2), 129–155.

van Benthem, J. (2008). Logical dynamics of information and interaction. Book manuscript.

van Benthem, J. (2008). Merging observation and access in dynamic logic. manuscript.

van Benthem, J. and C. Dégremont (2009). Building bridges between dynamic and temporal doxastic logics. In *Proceedings of LOFT 8.*

van Benthem, J., J. Gerbrandy, T. Hoshi, and E. Pacuit (2008). Merging frameworks of interaction. *Journal of Philosophical Logic*, Forthcoming.

van Benthem, J., J. Gerbrandy, and B. Kooi (2008). Dynamic update with probabilities. Manuscript.

van Benthem, J. and D. Ikegami (2008). Modal fixed-point logic and changing models. In A. Avron, N. Dershowitz, and A. Rabinovich (Eds.), *Pillars of Computer Science, Essays Dedicated to Boris (Boaz) Trakhtenbrot on the Occasion of His 85th Birthday*, Volume 4800 of *Lecture Notes in Computer Science*, pp. 146–165. Berlin: Springer-Verlag.

van Benthem, J. and F. Liu (2007). Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics 17*(2), 157–182.

van Benthem, J. and E. Pacuit (2006). The tree of knowledge in action: Towards a common perspective. In G. Governatori, I. Hodkinson, and Y. Venema (Eds.), *Proceedings of Advances in Modal Logic Volume 6*, pp. 87 – 106. King's College Press.

van Benthem, J., J. van Eijck, and B. Kooi (2006). Logics of communication and change. *Information and Computation 204*(11), 1620 – 1662.

van der Hoek, W. and M. Pauly (2006). Modal logic for games and information. In P. Blackburn, J. van Benthem, and F. Wolter (Eds.), *Handbook of Modal Logic*, Volume 3 of *Studies in Logic*, pp. 1077 – 1148. Elsevier.

van der Hoek, W. and M. Wooldridge (2003). Towards a logic of rational agency. *Logic Journal of the IGPL 11*(2), 135 – 160.

van der Meyden, R. (1994, July). Axioms for knowledge and time in distributed systems with perfect recall. In *Proc. IEEE Symposium on Logic in Computer Science*, pp. 448–457.

van der Meyden, R. (1996). Finite state implementations of knowledge-based programs. In V. Chandru and V. Vinay (Eds.), *Proceedings of the Conference on the Foundations of Sofware Technology and Theoretical Computer Science*, Volume 1180 of *Lecture Notes in Computer Science*, Berlin, pp. 31–50. Springer-Verlag.

van der Meyden, R. (2007). Two applications of epistemic logic in computer security. In *Logic at the Crossroads*, pp. 207–222.

van der Meyden, R. and T. Wilke (2007). Preservation of epistemic properties in security protocol implementations. In *Proceedings of Theoretical Aspects of Rationality and Knowledge*.

van der Meyden, R. and K. Wong (2003, October). Complete axiomatizations for reasoning about knowledge and branching time. *Studia Logica 75*(1), 93 – 123.

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*, Amsterdam, pp. 87–117. Amsterdam University Press.

van Ditmarsch, H., W. van der Hoek, and B. Kooi (2007). *Dynamic Epistemic Logic*. Synthese Library. Springer.

van Ditmarsch, H. P. (2003). The Russian cards problem. *Studia Logica 75*, 31–62.

van Ditmarsch, H. P. (2005). Prolegomena to dynamic logic for belief revision. *Synthese (Knowledge, Rationality & Action) 147*, 229–275.

van Ditmarsch, H. P., A. Herzig, and T. de Lima (2007). Optimal regression for reasoning about knowledge and actions. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, pp. 1070–1075. AAAI Press.

van Ditmarsch, H. P. and B. Kooi (2006). The secret of my success. *Synthese 151*(2), 201–232.

van Ditmarsch, H. P., W. van der Hoek, and B. P. Kooi (2005). Dynamic epistemic logic with assignment. In F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh, and M. Wooldridge (Eds.), *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)*, Volume 1, New York, pp. 141–148. ACM Inc.

van Ditmarsch, H. P., W. van der Hoek, R. van der Meyden, and J. Ruan (2006). Model checking russian cards. *Electronic Notes in Theoretical Computer Science 149*, 790–795.

van Eijck, J. (2004a). Guarded actions. Technical Report SEN-E0425, CWI, Amsterdam. Available from `http://db.cwi.nl/rapporten/`.

van Eijck, J. (2004b). Reducing dynamic epistemic logic to PDL by program transformation. Technical Report SEN-E0423, CWI, Amsterdam. Available from `http://db.cwi.nl/rapporten/`.

van Eijck, J. (2005). DEMO: a system for dynamic epistemic modeling. `http://homepages.cwi.nl/~jve/demo`.

van Eijck, J. and S. M. Orzan (2005). Modelling the epistemics of communication with functional programming. In M. van Eekelen (Ed.), *Sixth Symposium on Trends in Functional Programming TFP'05*, pp. 44–59.

Velazquez-Quesada, F. R. (2009). Inference and update. *Synthese (Knowledge, Rationality & Action)*, forthcoming. manuscript.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic 25*(3), 221–261.

von Neumann, J. and O. Morgenstern (1944). *A Theory of Games and Economic Behaviour*. Princeton University Press: Princeton, NJ.