Synthese (2009) 169:283–300
Knowledge, Rationality & Action 687–704
DOI 10.1007/s11229-009-9556-2

Inference and update

Fernando Raymundo Velázquez-Quesada

Received: 18 November 2008 / Accepted: 7 April 2009 / Published online: 29 April 2009

Abstract We look at two fundamental logical processes, often intertwined in planning and problem solving: inference and update. Inference is an internal process with which we uncover what is implicit in the information we already have. Update, on the other hand, is produced by external communication, usually in the form of announcements and in general in the form of observations, giving us information that might not have been available (even implicitly) before. Both processes have received attention from the logic community, usually separately. In this work, we develop a logical language that allows us to describe them together. We present syntax, semantics and a complete axiom system; we discuss similarities and differences with other approaches and mention how the work can be extended.

Keywords Inference · Update · Epistemic logic · Dynamic epistemic logic

1 Introduction

Consider the following situation, from van Benthem (2008a):

> You are in a restaurant with your parents, and you have ordered fish, meat, and vegetarian, for you, your father and your mother, respectively. A new waiter comes with the three dishes. What can he do to know which dish corresponds to which person?

F. R. Velázquez-Quesada (✉)
Institute for Logic, Language and Computation, Universiteit van Amsterdam, Plantage Muidergracht 24, 1018 TV, Amsterdam, The Netherlands
e-mail: F.R.VelazquezQuesada@uva.nl

The waiter can ask "*Who has the fish?*"; then he can ask "*Who has the meat?*". Now he does not have to ask anymore: "two questions plus one inference are all that is needed" (van Benthem 2008a).

The present work looks at these two fundamental logical processes, often intertwined in real-life activities. Inference is an *internal* process: the agent revises her own information in search of what can be derived from it. Update, on the other hand, is produced by *external* communication: the agent gets new information via observations. Both are logical processes, both describe dynamics of information, both are used in every day situations and still, they have been studied separately.

Inference has been the main subject of study of logic, allowing us to extract new information from what we already have. Among the most important branches, we can mention Hilbert-style proof systems, natural deduction and tableaux. Recent works, like Jago (2006b,a), have incorporated modal logics to the field, representing inference as a non-deterministic step-by-step process.

Update has been the main subject of *Dynamic Epistemic Logic* (DEL) (van Ditmarsch et al. 2007). Works like Plaza (1989) and Gerbrandy (1999) turned attention to the effect public announcements have on the knowledge of an agent. Many works have followed them, including the study of more complex actions (Baltag et al. 1999, van Ditmarsch 2000) and the effect of announcements on diverse concepts of information (the soft/hard facts of van Benthem (2007); the knowledge/belief/safe belief of Baltag and Smets (2008)).

In van Benthem (2008c), the author shows how these two phenomena fall within the scope of modern logic: "asking a question and giving an answer is just as 'logical' as drawing a conclusion!". Here, we propose a merging of the two traditions. We consider that both processes are important, but so is their interaction. In this work, we develop a logical language that allows us to express inference and update together.

We start in Sect. 2 by providing a general framework for representing *implicit* and *explicit* information; then, by asking for the adequate properties, we focus on the case where information is true, i.e., we deal with *knowledge*. Section 3 provides a representation of *truth-preserving inference*; moreover, we show how dynamics of the inference process itself can be represented. Section 4 introduces the other logical process: *update*. Then, we compare our proposal with other approaches (Sect. 5) and present a summary and further work (Sect. 6). We focus in the single-agent case, leaving group-information concepts for future work.

2 Implicit and explicit information

The *Epistemic Logic* (EL) framework with Kripke models (Hintikka 1962) is one of the most widely used for representing and reasoning about an agent's information. Nevertheless, it is not fine enough to represent the restaurant example above. Agents whose information is represented with this framework suffer from what Hintikka called the *logical omniscience* problem:[1] they are informed of all validities and their information is closed under truth-preserving inference.

[1] See Sim (1997) for a survey about the logical omniscience problem.

This feature, useful in some applications, is too much in some others. More importantly for us, it *hides* the inference process: when representing our example, the answer to the second question tells the waiter not only that your father will get the meat but also that your mother will get the vegetarian dish. In this case, the inference is short and very simple, but in general this is not the case: proving a theorem, for example, consists on successive applications of deductive inference steps to show that the conclusion indeed follows from the premises. Some theorems may be straightforward, but some are not, and the distinction does not correspond to immediate notions like the length of the proof.

As argued in van Benthem (2006), we can give the modal operator another interpretation, reading the formula $\Box\varphi$ as "*the agent is implicitly informed about* φ". We follow that idea, and we extend EL to also represent *explicit* information; moreover, we also provide a mechanism with which the agent can increase it. The work of this section resembles previous literature; a comparison appears in Sect. 5.

2.1 Formulas, rules and the explicit/implicit information framework

In our framework, the agent's explicit information is given by a set of formulas, and it can be increased by the use of syntactic rules. We start by defining the language to represent explicit information and by indicating what a rule in that language is.

Definition 2.1 (*Formulas and rules*). Let $\mathtt{P}$ be a set of atomic propositions. The *internal language* $\mathcal{I}$ is given by the propositional language over $\mathtt{P}$. A *rule* based on $\mathcal{I}$ is a pair (Γ, γ) (also represented as $\Gamma \Rightarrow \gamma$) where Γ is a finite set of formulas and γ is a formula, all of them in $\mathcal{I}$. Given a rule $\rho = (\Gamma, \gamma)$, we call Γ the *set of premises of* ρ ($\text{prem}(\rho)$) and γ the *conclusion of* ρ ($\text{conc}(\rho)$). We denote by $\mathcal{R}$ the set of rules based on formulas of $\mathcal{I}$.

Our internal language allows the agent to have explicit information about facts but not about her own information. This is indeed a limitation, but it allows us to define one of the two processes we are interested in: *update* (Sect. 4). In Sect. 6 we briefly discuss the reasons for this limitation, leaving a deep analysis for further work. Note also that we have defined the premises of a rule as a *set*, and not as a more general notion like an *ordered sequence* or a *multi-set*.[2] This makes it closer to the classical notion of *truth-preserving inference*, which we explore in Sect. 3.

Our language extends that of EL by adding two kinds of formulas: one for expressing the agent's explicit information ($I\,\gamma$) and another expressing the rules she can apply ($L\,\rho$).

Definition 2.2 (*Explicit/implicit information language* $\mathcal{EI}$). Let $\mathtt{P}$ be a set of atomic propositions. Formulas of the *explicit/implicit information language* $\mathcal{EI}$ are given by

[2] As a referee pointed out, with such a generalized definition, one can also analyze inference in "resource-conscious" sub-structural logics in which order and multiplicity matters, like Linear Logic or Categorial Grammar (see Moortgat 1997). This interesting extension fits well with the idea of our awareness analysis, but we must leave it to further work.

$$\varphi ::= \top \mid p \mid I\,\gamma \mid L\,\rho \mid \neg\varphi \mid \varphi \vee \psi \mid \Diamond\varphi$$

with $p \in \mathtt{P}$, $\gamma \in \mathcal{I}$ and $\rho \in \mathcal{R}$. Formulas of the form $\Diamond\varphi$ are read as *implicitly, the agent considers* φ *possible*. Other boolean connectives ($\wedge$, $\rightarrow$ and $\leftrightarrow$) as well as the modal operator $\Box$ are defined as usual.

The semantic model extends a Kripke model by assigning two new sets to each possible world: one indicating the *formulas* the agent is explicitly informed about, and other indicating the *rules* she can apply.

Definition 2.3 (*Explicit/implicit information model*) Let $\mathtt{P}$ be a set of atomic propositions. An *explicit/implicit information model* is a tuple $M = \langle W, R, V, Y, Z\rangle$ where $\langle W, R, V\rangle$ is a standard Kripke model (W the non-empty set of worlds, $R \subseteq W \times W$ the *accessibility relation*, $V : W \rightarrow \wp(\mathtt{P})$ the *atomic valuation function*) and

- $Y : W \rightarrow \wp(\mathcal{I})$ is the *information set function*, satisfying ***coherence for formulas*** (if $\gamma \in Y(w)$ and Rwu, then $\gamma \in Y(u)$).
- $Z : W \rightarrow \wp(\mathcal{R})$ is the *rule set function* satisfying ***coherence for rules*** (if $\rho \in Z(w)$ and Rwu, then $\rho \in Z(u)$).

We denote by **EI** the class of all explicit/implicit information models. Note how, just as in the definition of the premises of a rule, the agent's information about formulas and rules is also given by a set.

Our restrictions reflect the following idea. The sets $Y(w)$ and $Z(w)$ represent the formulas and rules the agent is explicitly informed about; if while staying in w the agent considers u possible, it is natural to ask for u to preserve the agent's explicit information at w.

Definition 2.4 (*Semantics for* $\mathcal{EI}$) Given a model $M = \langle W, R, V, Y, Z\rangle$ in **EI** and a world $w \in W$, the semantics for $\top$ and disjunctions is as usual. For the rest,

$$\begin{array}{llll} (M,w) \Vdash p & \textit{iff} \;\; p \in V(w) & (M,w) \Vdash I\,\gamma & \textit{iff} \;\; \gamma \in Y(w) \\ (M,w) \Vdash \neg\varphi & \textit{iff} \;\; (M,w) \not\Vdash \varphi & (M,w) \Vdash L\,\rho & \textit{iff} \;\; \rho \in Z(w) \\ (M,w) \Vdash \Diamond\varphi & \textit{iff} \;\; \exists u\, s.t. Rwu \,\&\, (M,u) \Vdash \varphi & & \end{array}$$

Note how the operators I and L just look into the correspondent sets.

We provide a syntactic characterization of formulas of $\mathcal{EI}$ that are valid in the class of models **EI**.

Theorem 1 (Sound and complete logic for $\mathcal{EI}$ w.r.t. **EI**). *The logic* EI *(Table 1) is sound and strongly complete for the language* $\mathcal{EI}$ *with respect to models in* **EI**.

Proof Soundness follows from axioms being valid and rules being validity-preserving. Completeness follows by a standard modal canonical model construction with information set and rule set functions given by $Y^{\mathsf{EI}}(w) := \{\gamma \in \mathcal{I} \mid I\,\gamma \in w\}$ and $Z^{\mathsf{EI}}(w) := \{\rho \in \mathcal{R} \mid L\,\rho \in w\}$. It is easy to show that these satisfy the crucial coherence properties. □

Table 1 Axioms and inference rules for the logic EI

(*P*)	All propositional tautologies	$(Coh_{\mathcal{I}}) \vdash I\,\gamma \rightarrow \Box I\,\gamma$
(*K*)	$\vdash \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\phi)$	$(Coh_{\mathcal{R}}) \vdash L\,\rho \rightarrow \Box L\,\rho$
(*Dual*)	$\vdash \Diamond\varphi \leftrightarrow \neg\Box\neg\varphi$	
(*MP*)	From $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$ infer $\vdash \psi$	(*Gen*) From $\vdash \varphi$ infer $\vdash \Box\varphi$

Axioms of Table 1 say, in particular, that the agent's implicit information is closed under modus ponens (axiom *K*), and that her explicit information satisfy both coherence properties ($Coh_{\mathcal{I}}$ and $Coh_{\mathcal{R}}$). Note how the agent's explicit information does not suffer from the *logical omniscience* problem: the validity of γ does not imply the validity of $I\,\gamma$, and $I\,(\gamma \rightarrow \delta) \rightarrow (I\,\gamma \rightarrow I\,\delta)$ is not valid.

2.2 When information is true

Models in the class **EI** can represent information that is not true: the actual world does not have to be in those considered possible by the agent, and formulas (rules) in information (rule) sets do not have to be true (truth-preserving) at the corresponding world. By asking for the adequate properties, we can represent diverse kinds of information concepts, like *safe belief*, *belief* or *knowledge*. Here, we will focus on the case of *true* information, i.e., *knowledge*.[3]

Among models in **EI**, we distinguish those where implicit and explicit information are true and the rules are truth-preserving. For implicit information, we consider equivalence accessibility relations, as it is usually done in EL.[4] For explicit information, we ask for every formula in an information set to be true in the corresponding world, and for rules we define a translation TR that maps each rule to an implication of the form $\mathrm{TR}(\rho) := \bigwedge \mathrm{prem}(\rho) \rightarrow \mathrm{conc}(\rho)$, and we ask for this translation to be true in the correspondent world.

Definition 2.5 (*The class* $\mathbf{EI}_K$) We denote by $\mathbf{EI}_K$ the class of models in **EI** satisfying ***equivalence*** (*R* is an equivalence relation), ***truth for formulas*** (for every world w, if $\gamma \in Y(w)$, then $(M, w) \Vdash \gamma$) and ***truth for rules*** (for every world w, if $\rho \in Z(w)$, then $(M, w) \Vdash \mathrm{TR}(\rho)$).

From now on, we will use the term "information" for the general case of information (that is, models in **EI**) and we will use the term "knowledge" for *true* information (that is, models in $\mathbf{EI}_K$).

Theorem 2 (Sound and complete logic for $\mathcal{EI}$ w.r.t. $\mathbf{EI}_K$). *The logic* EI_K*, extending* EI *with axioms of Table 2, is sound and strongly complete for the language* $\mathcal{EI}$ *with respect to models in* $\mathbf{EI}_K$.

Proof Soundness is again simple. Completeness is proved by showing that the canonical model for EI_K satisfies *equivalence* (from axioms *T*, *4*, *5*), *truth for formulas* (from $Tth_{\mathcal{I}}$) and *truth for rules* (from $Tth_{\mathcal{R}}$). □

[3] For literature about information that can be true or false, we refer to Dretske (1981) and Floridi (2005).

[4] Given our understanding of *knowledge*, we actually just need for the relation to be reflexive. This is enough to make the information *true*.

Table 2 Extra axioms for the logic EI_K

$(T) \vdash \Box\varphi \to \varphi$	$(Tth_{\mathcal{I}}) \vdash I\,\gamma \to \gamma$
$(4) \vdash \Box\varphi \to \Box\Box\varphi$	$(Tth_{\mathcal{R}}) \vdash L\,\rho \to \mathrm{TR}(\rho)$
$(5) \vdash \neg\Box\varphi \to \Box\neg\Box\varphi$	

Axioms *T*, *4* and *5* express properties of implicit knowledge: it is *true* (*T*) and it has the *positive and negative introspection* properties (*4* and *5*); axioms $Tth_{\mathcal{I}}$ and $Tth_{\mathcal{R}}$ indicate that the agent's explicit knowledge about formulas and rules is also *true*. Note that from $Coh_{\mathcal{I}}(I\,\gamma \to \Box I\,\gamma)$ and $Tth_{\mathcal{I}}(I\,\gamma \to \gamma)$ we get $I\,\gamma \to \Box\gamma$: whatever is part of the agent's explicit knowledge belongs to her implicit knowledge too.

Now we turn our attention to *dynamics* of explicit and implicit information. In the following sections, we extend the framework to describe inference and update.

3 Inference

The agent can extend her explicit information by using rules. Intuitively, a rule (Γ, γ) indicates that if every $\delta \in \Gamma$ is true, so is γ. However, so far, we have not stated any restriction on how the agent can use a rule. She can use it to get the conclusion without having all the premises, or even deriving the premises whenever she has the conclusion. In the previous section we focused on true-information models; in the same spirit, this section deals with truth-preserving inference.

3.1 A particular case: truth-preserving inference

The inference process adds formulas to the information set. In order to preserve truth, we restrict the way in which the rule can be applied.

Definition 3.1 (*Deduction operation*). Let $M = \langle W, R, V, Y, Z\rangle$ be a model in **EI**, and let σ be a rule in $\mathcal{R}$. The model $M_\sigma = \langle W, R, V, Y', Z\rangle$ differs from M just in the information set function, which is given by

$$Y'(w) := \begin{cases} Y(w) \cup \{\mathrm{conc}(\sigma)\} & \text{if } \mathrm{prem}(\sigma) \subseteq Y(w) \text{ and } \sigma \in Z(w) \\ Y(w) & \text{otherwise} \end{cases}$$

The conclusion of the rule is added to a world just when all the premises and the rule are already present. This allows us to prove that, in particular, the deduction operation preserves models in $\mathbf{EI}_K$.

Proposition 1 *Let σ be a rule. If M is a model in $\mathbf{EI}_K$, so is M_σ.*

Proof Equivalence and both properties of rules are immediate since neither the accessibility relation nor the rule set function are modified. The properties of formulas can be verified easily. □

The language $\mathcal{EID}$ extends $\mathcal{EI}$ by closing it under deduction modalities $\langle \mathrm{D}_\sigma\rangle$ for σ a rule: if φ is a formula in $\mathcal{EID}$, so is $\langle \mathrm{D}_\sigma\rangle\,\varphi$. These new formulas are read as *there is a deductive inference with σ after which φ is the case*. Define the abbreviation

$\text{Pre}_\sigma \equiv I\ \text{prem}(\sigma) \wedge L\,\sigma$ where, given Γ a finite set of formulas in $\mathcal{I}$, we write $I\,\Gamma$ for $\bigwedge_{\gamma\in\Gamma} I\,\gamma$. The semantics for deduction formulas is given as follows.

Definition 3.2 Let M be a model in **EI**, and take a world w in it.

$$(M, w) \Vdash \langle \text{D}_\sigma \rangle\, \varphi \quad \textit{iff} \quad (M, w) \Vdash \text{Pre}_\sigma \text{ and } (M_\sigma, w) \Vdash \varphi$$

that is, $\langle \text{D}_\sigma \rangle\, \varphi$ holds at w *iff* at w the agent explicitly has σ and its premises, and after applying the rule σ, the formula φ holds. The formula $[\text{D}_\sigma]\,\varphi$ is given by $[\text{D}_\sigma]\,\varphi \leftrightarrow \neg\langle \text{D}_\sigma \rangle\, \neg\varphi$, as usual.

For an axiom system, Proposition 1 tells us that the operation preserves models in $\mathbf{EI}_K$, so we can rely on the logic EI_K. We provide *reduction axioms*, expressing how deduction operations affect the truth-value of formulas of the language. This is a standard DEL technique, and we refer to van Benthem and Kooi (2004) for a deep explanation.

Theorem 3 (Sound and complete logic for $\mathcal{EID}$ w.r.t. $\mathbf{EI}_K$). *The logic EI_{KD}, extending EI_K with axioms and rules of Table 3, is sound and strongly complete for the language $\mathcal{EID}$ with respect to models in $\mathbf{EI}_K$.*

Proof Soundness is just as before. Strong completeness comes from the fact that, by a repetitive application of the new axioms, any deduction formula can be reduced to a formula in $\mathcal{EI}$, for which EI_K is strongly complete with respect to $\mathbf{EI}_K$. □

The novel axioms of Table 3 are those expressing how information and rule sets are affected by deduction. From them, we can derive formulas like (*1*) $\langle \text{D}_\sigma \rangle \top \rightarrow \text{Pre}_\sigma$, (*2*) $I\,\gamma \rightarrow [\text{D}_\sigma]\, I\,\gamma$, (*3*) $[\text{D}_\sigma]\, I\ \text{conc}(\sigma)$ and (*4*) $\langle \text{D}_\sigma \rangle\, I\,\gamma \rightarrow I\,\gamma$ (for $\gamma \neq \text{conc}(\sigma)$), indicating that in order to apply a rule we need the premises and the rule (*1*), that after applying a rule we preserve previous explicit knowledge (*2*) and that explicit knowledge is increased only by the rule's conclusion (*3*) and (*4*).

3.2 Dynamics of deduction

Just as the agent's explicit knowledge changes, her inferential abilities can also change. This may be because she is informed about a new rule (as with the *updates* of Sect. 4), but it may be also because she *builds* new rules from the ones she already has. For example, from the rules $\{p\} \Rightarrow q$ and $\{q\} \Rightarrow r$, we can derive the rule $\{p\} \Rightarrow r$. It takes one step to derive it, but it will save intermediate steps later.

Table 3 Axioms and rules for deduction operation formulas

$\vdash \langle \text{D}_\sigma \rangle \top \leftrightarrow \text{Pre}_\sigma$	$\vdash \langle \text{D}_\sigma \rangle\, p \leftrightarrow (\text{Pre}_\sigma \wedge p)$
$\vdash \langle \text{D}_\sigma \rangle \neg\varphi \leftrightarrow (\text{Pre}_\sigma \wedge \neg\langle \text{D}_\sigma \rangle\, \varphi)$	$\vdash \langle \text{D}_\sigma \rangle\, (\varphi \vee \psi) \leftrightarrow (\langle \text{D}_\sigma \rangle\, \varphi \vee \langle \text{D}_\sigma \rangle\, \psi)$
$\vdash \langle \text{D}_\sigma \rangle \Diamond\varphi \leftrightarrow (\text{Pre}_\sigma \wedge \Diamond\langle \text{D}_\sigma \rangle\, \varphi)$	
$\vdash \langle \text{D}_\sigma \rangle\, I\ \text{conc}(\sigma) \leftrightarrow \text{Pre}_\sigma$	$\vdash \langle \text{D}_\sigma \rangle\, L\,\rho \leftrightarrow (\text{Pre}_\sigma \wedge L\,\rho)$
$\vdash \langle \text{D}_\sigma \rangle\, I\,\gamma \leftrightarrow (\text{Pre}_\sigma \wedge I\,\gamma)$ for $\gamma \neq \text{conc}(\sigma)$	
From $\vdash \varphi$, infer $\vdash [\text{D}_\sigma]\,\varphi$	

In fact the example, a kind of *transitivity*, represents the application of *Cut* over the mentioned rules. In general, inference relations can be characterized by *structural rules*, indicating how to derive new rules from the ones already present. In the case of deduction, we have

Reflexivity: $\dfrac{}{\varphi \Rightarrow \varphi}$ *Contraction:* $\dfrac{\psi, \chi, \xi, \chi, \phi \Rightarrow \varphi}{\psi, \chi, \xi, \phi \Rightarrow \varphi}$

Permutation: $\dfrac{\psi, \chi, \xi, \phi \Rightarrow \varphi}{\psi, \xi, \chi, \phi \Rightarrow \varphi}$ *Monotonicity:* $\dfrac{\psi, \phi \Rightarrow \varphi}{\psi, \chi, \phi \Rightarrow \varphi}$

Cut: $\dfrac{\chi \Rightarrow \xi \quad \psi, \xi, \phi \Rightarrow \varphi}{\psi, \chi, \phi \Rightarrow \varphi}$

Each time a structural rule is applied, we get a rule that can be added to the rule set. Note that the application of *Contraction* or *Permutation* does not yield a new rule, since we are already considering the premises of a rule as a set.[5] On the other hand, *Reflexivity*, *Monotonicity* and *Cut* can produce rules that were not present before.

Definition 3.3 (*Structural operations*). Let $M = \langle W, R, V, Y, Z \rangle$ be a model in **EI**. The structural operations $(\cdot)_{\mathbf{Ref}(\delta)}$, $(\cdot)_{\mathbf{Mon}(\delta,\varsigma)}$ and $(\cdot)_{\mathbf{Cut}(\varsigma_1,\varsigma_2)}$, return a model that differs from M just in the rule set function.

Reflexivity Let δ be a formula in $\mathcal{I}$ and consider the rule $\varsigma_\delta = (\{\delta\}, \delta)$. The new rule set function is given by $Z'(w) := Z(w) \cup \{\varsigma_\delta\}$.

Monotonicity Let δ be a formula in $\mathcal{I}$ and let ς be a rule over $\mathcal{I}$. Consider the rule $\varsigma' = (\text{prem}(\varsigma) \cup \{\delta\}, \text{conc}(\varsigma))$, extending ς by adding δ to its premises. The new rule set function is given by

$$Z'(w) := \begin{cases} Z(w) \cup \{\varsigma'\} & \text{if } \varsigma \in Z(w) \\ Z(w) & \text{otherwise} \end{cases}$$

Cut Let ς_1, ς_2 be rules over $\mathcal{I}$ such that the conclusion of ς_1 is contained in the premises of ς_2. Consider the rule ς' with $(\text{prem}(\varsigma_2) - \{\text{conc}(\varsigma_1)\}) \cup \text{prem}(\varsigma_1)$ as premises and $\text{conc}(\varsigma_2)$ as conclusion. The new rule set function is given by

$$Z'(w) := \begin{cases} Z(w) \cup \{\varsigma'\} & \text{if } \{\varsigma_1, \varsigma_2\} \subseteq Z(w) \\ Z(w) & \text{otherwise} \end{cases}$$

The three structural operations preserve models in $\mathbf{EI}_K$.

Proposition 2 *If M is a model in $\mathbf{EI}_K$, then $M_{\mathbf{Ref}(\delta)}$, $M_{\mathbf{Mon}(\delta,\varsigma)}$ and $M_{\mathbf{Cut}(\varsigma_1,\varsigma_2)}$ are also in $\mathbf{EI}_K$.*

[5] This is not to say that order or multiplicity of *inference steps* are not relevant; given our *dynamic* approach, they definitely matter, as changes in order or number of inference steps can yield different results. We just mean that order and multiplicity of *premises* are irrelevant because we represent them as a set, and therefore the two mentioned operations do not yield new rules.

Proof Equivalence and both properties of formulas are immediate. Coherence for rules follows from the definitions and coherence for rules of M. For truth for rules, see Sect. A.1. □

The language $\mathcal{EID}^*$ extends $\mathcal{EID}$ by closing it under modalities for structural operations: if φ is in $\mathcal{EID}^*$, so are $\langle Ref_\delta \rangle\, \varphi$, $\langle Mon_{\delta,\varsigma} \rangle\, \varphi$ and $\langle Cut_{\varsigma_1,\varsigma_2} \rangle\, \varphi$. The formulas are read as "*there is a way of applying the structural operation after which* φ *is the case*". By expressing the precondition of each operation with $\mathrm{Pre}_{\mathbf{Mon}(\delta,\varsigma)} \equiv L\, \varsigma$ and $\mathrm{Pre}_{\mathbf{Cut}(\varsigma_1,\varsigma_2)} \equiv L\, \varsigma_1 \wedge L\, \varsigma_2 \wedge (I\ \mathrm{prem}(\varsigma_2) \rightarrow I\ \mathrm{conc}(\varsigma_1))$ (reflexivity can be applied in any situation), the semantics of the new formulas is given as follows.

Definition 3.4 Let M be a model in **EI**, and take a world w in it.

$$(M, w) \Vdash \langle Ref_\delta \rangle\, \varphi \quad \textit{iff } (M_{\mathbf{Ref}(\delta)}, w) \Vdash \varphi$$
$$(M, w) \Vdash \langle Mon_{\delta,\varsigma} \rangle\, \varphi \quad \textit{iff } (M, w) \Vdash \mathrm{Pre}_{\mathbf{Mon}(\delta,\varsigma)} \text{ and } (M_{\mathbf{Mon}(\delta,\varsigma)}, w) \Vdash \varphi$$
$$(M, w) \Vdash \langle Cut_{\varsigma_1,\varsigma_2} \rangle\, \varphi \quad \textit{iff } (M, w) \Vdash \mathrm{Pre}_{\mathbf{Cut}(\varsigma_1,\varsigma_2)} \text{ and } (M_{\mathbf{Cut}(\varsigma_1,\varsigma_2)}, w) \Vdash \varphi$$

Just as before, the *boxed* versions of the structural operation formulas are defined as the dual of their correspondent *diamond* versions.

To provide axioms for the new formulas, Proposition 2 allows us to rely on the logic EI_K once again.

Theorem 4 (Sound and complete logic for $\mathcal{EID}^*$ w.r.t. $\mathbf{EI}_K$). *For uniformity, define the precondition of the reflexivity operation as* $\mathrm{Pre}_{\mathbf{Ref}(\delta)} \equiv \top$. *The logic* EI_{KDS}, *extending* EI_{KD} *with axioms and rule of Table 4 (where STR stands for either Ref, Mon or Cut and* ς' *is the correspondent new rule), is sound and strongly complete for the language* $\mathcal{EID}^*$ *with respect to models in* $\mathbf{EI}_K$.

The relevant axioms of Table 4 are those expressing how rule sets are affected by structural operations, and from them we can derive validities analogous to those given at the end of Sect. 3.1 for the case of information sets and deduction.

In Table 5 we provide validities expressing how structural operations affect deduction at models of $\mathbf{EI}_K$. For each structural operation, the first formula indicates that the operation does not affect deduction with a rule different from the new one, and the second formula indicates how deduction with the new rule changes. For this last case, the formula covers two possibilities: the new rule was already in the original rule set (hence just deduction is needed) or it was not (hence we ask for some requisites). As an example, the second formula for Mon indicates that a sequence of monotonicity and then deduction with the new rule ς' is equivalent to a single deduction with ς'

Table 4 Axioms and rules for *reflexivity, monotonicity* and *cut* formulas

$\vdash \langle STR \rangle \top \leftrightarrow \mathrm{Pre}_{\mathbf{STR}}$	$\vdash \langle STR \rangle\, p \leftrightarrow (\mathrm{Pre}_{\mathbf{STR}} \wedge p)$
$\vdash \langle STR \rangle \neg\varphi \leftrightarrow (\mathrm{Pre}_{\mathbf{STR}} \wedge \neg\langle STR \rangle\, \varphi)$	$\vdash \langle STR \rangle\, (\varphi \vee \psi) \leftrightarrow (\langle STR \rangle\, \varphi \vee \langle STR \rangle\, \psi)$
$\vdash \langle STR \rangle \Diamond\varphi \leftrightarrow (\mathrm{Pre}_{\mathbf{STR}} \wedge \Diamond\langle STR \rangle\, \varphi)$	
$\vdash \langle STR \rangle\, L\, \varsigma' \leftrightarrow \mathrm{Pre}_{\mathbf{STR}}$	$\vdash \langle STR \rangle\, I\, \gamma \leftrightarrow (\mathrm{Pre}_{\mathbf{STR}} \wedge I\, \gamma)$
$\vdash \langle STR \rangle\, L\, \rho \leftrightarrow (\mathrm{Pre}_{\mathbf{STR}} \wedge L\, \rho)$ for $\rho \neq \varsigma'$	
From $\vdash \varphi$, infer $\vdash [STR]\, \varphi$	

Table 5 Formulas relating structural operations and deduction

Reflexivity with ς_δ *the rule* $(\{\delta\}, \delta)$	
• $\langle Ref_\delta\rangle\, \langle \mathrm{D}_\sigma\rangle\, \varphi \;\leftrightarrow\; \langle \mathrm{D}_\sigma\rangle\, \langle Ref_\delta\rangle\, \varphi$	for $\sigma \neq \varsigma_\delta$
• $\langle Ref_\delta\rangle\, \langle \mathrm{D}_{\varsigma_\delta}\rangle\, \varphi \;\leftrightarrow\; (\langle \mathrm{D}_{\varsigma_\delta}\rangle\, \varphi \vee (I\,\delta \wedge \langle Ref_\delta\rangle\, \varphi))$	
Monotonicity with ς' *the rule* $(\mathrm{prem}(\varsigma) \cup \{\delta\}, \mathrm{conc}(\varsigma))$	
• $\langle Mon_{\delta,\varsigma}\rangle\, \langle \mathrm{D}_\sigma\rangle\, \varphi \;\leftrightarrow\; \langle \mathrm{D}_\sigma\rangle\, \langle Mon_{\delta,\varsigma}\rangle\, \varphi$	for $\sigma \neq \varsigma'$
• $\langle Mon_{\delta,\varsigma}\rangle\, \langle \mathrm{D}_{\varsigma'}\rangle\, \varphi \;\leftrightarrow\; (\langle \mathrm{D}_{\varsigma'}\rangle\, \varphi \vee (I\,\delta \wedge L\,\varsigma \wedge \langle \mathrm{D}_\varsigma\rangle\, \langle Mon_{\delta,\varsigma}\rangle\, \varphi))$	
Cut with ς' *the rule* $(\,(\mathrm{prem}(\varsigma_2) - \{\mathrm{conc}(\varsigma_1)\}) \cup \mathrm{prem}(\varsigma_1),\; \mathrm{conc}(\varsigma_2)\,)$	
• $\langle Cut_{\varsigma_1,\varsigma_2}\rangle\, \langle \mathrm{D}_\sigma\rangle\, \varphi \;\leftrightarrow\; \langle \mathrm{D}_\sigma\rangle\, \langle Cut_{\varsigma_1,\varsigma_2}\rangle\, \varphi$	for $\sigma \neq \varsigma'$
• $\langle Cut_{\varsigma_1,\varsigma_2}\rangle\, \langle \mathrm{D}_{\varsigma'}\rangle\, \varphi \;\leftrightarrow$	
$(\langle \mathrm{D}_{\varsigma'}\rangle\, \varphi \vee (I\; \mathrm{prem}(\varsigma_1) \wedge L\,\varsigma_1 \wedge (I\; \mathrm{conc}(\varsigma_1) \rightarrow \langle \mathrm{D}_{\varsigma_2}\rangle\, \langle Cut_{\varsigma_1,\varsigma_2}\rangle\, \varphi)))$	

(if ς' was already present) or to a sequence of deduction with ς and then monotonicity with the agent having explicitly knowledge about the added premise δ and the original rule ς. See Sect. A.2 for comments about the proofs.

4 Update

So far, our language can express just *internal* dynamics. We can express how deductive steps modify explicit knowledge, and even how structural operations extends the rules the agent can apply, but we cannot express how knowledge is affected by *external* interaction. Here, we add the other fundamental source of information; in this section, we extend the language to express updates.

Updates are the result of the agent's social nature. We get new information because of interaction with our environment; information that does not necessarily follow from what we explicitly have. In *Public Announcement Logic* (PAL), an announcement is interpreted as an operation, removing worlds where the announcement does not hold. In our case, we distinguish different kinds of knowledge: implicit knowledge (given by the accessibility relation) and explicit knowledge (the information sets). We can define operations affecting explicit and implicit knowledge in different forms, and therefore expressing different ways the agent processes external information. Here, we present one of the possible definitions, what we call *explicit observations*.

4.1 True explicit observations

The previously defined operations just add formulas or rules to the corresponding sets, but do not modify the accessibility relation and therefore do not affect implicit knowledge. True explicit observations, on the other hand, do modify the accessibility relation by removing worlds where the observation does not hold. With respect to explicit knowledge, they always add the observation (a formula or a rule).

Definition 4.1 (*Explicit observation operation*). Take a model $M = \langle W, R, V, Y, Z\rangle$ in **EI**, and let χ be a formula of (a rule based on) $\mathcal{I}$. The model $M_{\chi!}{=}\langle W', R', V', Y', Z'\rangle$ is given by

- $W' := \{ w \in W | (M, w) \Vdash \chi \}$ $(W' := \{ w \in W | (M, w) \Vdash \mathrm{TR}(\chi) \})$,
- $R' := R \cap (W' \times W')$ • $V'(w) := V(w)$ for every $w \in W'$
- $Y'(w) := Y(w) \cup \{\chi\}$ $(Y'(w) := Y(w))$ for every $w \in W'$,
- $Z'(w) := Z(w)$ $(Z'(w) := Z(w) \cup \{\chi\})$ for every $w \in W'$.

Proposition 3 *Let M be a model in $\mathbf{EI}_K$ and let χ be a formula (a rule). If M is in $\mathbf{EI}_K$, so is $M_{\chi!}$.*

Proof Equivalence is immediate as well as the properties for rules (formulas), since $Z(Y)$ is not affected in the remaining worlds. Coherence for formulas (rules) holds because χ is added uniformly, and truth for formulas (rules) holds because of the definition of W'. □

The language $\mathcal{EID}^{*!}$ closes $\mathcal{EID}^{*}$ under modalities $\langle \chi ! \rangle$ for explicit observations. The new formulas are read as "*there is a way of explicitly observing χ after which φ is the case*". In case χ is a formula, define its precondition as $\mathrm{Pre}_{\chi!} \equiv \chi$; in case χ is a rule, define its precondition as $\mathrm{Pre}_{\chi!} \equiv \mathrm{TR}(\chi)$. The semantics for the new formulas is given as follows.

Definition 4.2 Let M be a model in **EI**, and take a world w in it.

$$(M, w) \Vdash \langle \chi ! \rangle \varphi \quad \textit{iff} \quad (M, w) \Vdash \mathrm{Pre}_{\chi!} \text{ and } (M_{\chi!}, w) \Vdash \varphi$$

The formula $[\chi !] \varphi$ is defined as the dual of $\langle \chi ! \rangle \varphi$, as usual.

Theorem 5 (Sound and complete logic for $\mathcal{EID}^{*!}$ w.r.t. $\mathbf{EI}_K$). *The logic EI_{KDSO}, extending EI_{KDS} with axioms and rule of Table 6, is sound and strongly complete for $\mathcal{EID}^{*!}$ with respect to models in $\mathbf{EI}_K$.*

The relevant axioms are those indicating how explicit knowledge about formulas and rules is affected: the agent always knows the observation explicitly after observing it, and any other explicit knowledge was already present before the observation. The axioms look similar to those for deduction and structural operations, but there is an important difference: the precondition. While in the case of deduction and structural operations the agent needs to have enough explicit knowledge to extract the new piece,

Table 6 Axioms and rules for *explicit observation* formulas

$\vdash \langle \chi ! \rangle \top \leftrightarrow \mathrm{Pre}_{\chi!}$	$\vdash \langle \chi ! \rangle p \leftrightarrow (\mathrm{Pre}_{\chi!} \wedge p)$
$\vdash \langle \chi ! \rangle \neg\varphi \leftrightarrow (\mathrm{Pre}_{\chi!} \wedge \neg\langle \chi ! \rangle \varphi)$	$\vdash \langle \chi ! \rangle (\varphi \vee \psi) \leftrightarrow (\langle \chi ! \rangle \varphi \vee \langle \chi ! \rangle \psi)$
$\vdash \langle \chi ! \rangle \Diamond\varphi \leftrightarrow (\mathrm{Pre}_{\chi!} \wedge \Diamond\langle \chi ! \rangle \varphi)$	
If χ is a formula:	
$\vdash \langle \chi ! \rangle I\,\chi \leftrightarrow \mathrm{Pre}_{\chi!}$	$\vdash \langle \chi ! \rangle L\,\rho \leftrightarrow (\mathrm{Pre}_{\chi!} \wedge L\,\rho)$
$\vdash \langle \chi ! \rangle I\,\gamma \leftrightarrow (\mathrm{Pre}_{\chi!} \wedge I\,\gamma)$ for $\gamma \neq \chi$	
If χ is a rule:	
$\vdash \langle \chi ! \rangle L\,\chi \leftrightarrow \mathrm{Pre}_{\chi!}$	$\vdash \langle \chi ! \rangle I\,\gamma \leftrightarrow (\mathrm{Pre}_{\chi!} \wedge I\,\gamma)$
$\vdash \langle \chi ! \rangle L\,\rho \leftrightarrow (\mathrm{Pre}_{\chi!} \wedge L\,\rho)$ for $\rho \neq \chi$	
From $\vdash \varphi$, infer $\vdash [\chi !]\varphi$	

Table 7 Formulas relating explicit observations and deduction

If χ is a formula	
• $\langle \chi! \rangle \langle \mathrm{D}_\sigma \rangle \varphi \leftrightarrow \langle \mathrm{D}_\sigma \rangle \langle \chi! \rangle \varphi$	if $\chi \notin \mathrm{prem}(\sigma)$
• $\langle \chi! \rangle \langle \mathrm{D}_\sigma \rangle \varphi \leftrightarrow (\langle \mathrm{D}_\sigma \rangle \langle \chi! \rangle \varphi \vee (I\,\chi \rightarrow \langle \mathrm{D}_\sigma \rangle \langle \chi! \rangle \varphi))$	if $\chi \in \mathrm{prem}(\sigma)$
If χ is a rule	
• $\langle \chi! \rangle \langle \mathrm{D}_\sigma \rangle \varphi \leftrightarrow \langle \mathrm{D}_\sigma \rangle \langle \chi! \rangle \varphi$	if $\chi \neq \sigma$
• $\langle \chi! \rangle \langle \mathrm{D}_\chi \rangle \varphi \leftrightarrow (\langle \mathrm{D}_\chi \rangle \langle \chi! \rangle \varphi \vee (L\,\chi \rightarrow \langle \mathrm{D}_\chi \rangle \langle \chi! \rangle \varphi))$	if $\chi = \sigma$

observation is a more radical informational process: it just need for the observation to be *true* (*truth-preserving*).

To finish this section, Table 7 presents formulas indicating how the two informational processes considered in this paper, inference and update, interact with each other.[6] Together with Table 5, it provides principles about how external and internal dynamics intertwine when we process information, as it will be shown when revising the restaurant example in Sect. 6.

5 Comparison with other works

The present work develops a representation of explicit/implicit information in order to describe the way different processes affect them. Other works have proposed similar frameworks; this section provides a brief comparison between some of them and our proposal.

5.1 Fagin–Halpern's logics of awareness

Fagin and Halpern presented in 1988 a *logic of general awareness* ($\mathcal{L}_A$). The language is a set of atomic propositions P closed under negation, conjunction and the operators A_i ($A_i\varphi$ is read as "*the agent i is aware of* φ") and L_i ($L_i\varphi$ is read as "*the agent i implicitly believes that* φ").

The semantic model is a tuple $M = (W, \mathfrak{A}_i, \mathfrak{L}_i, V)$ with $(W, \mathfrak{L}_i, V)$ a Kripke model ($\mathfrak{L}_i$ a serial, Transitive and Euclidean relation) and $\mathfrak{A}_i$ a function assigning a set of formulas of $\mathcal{L}_A$ to each agent i in each world (her awareness set). Semantics for atomic propositions, negations, conjunctions and L_i (a box modal operator) are standard; for formulas of the form $A_i\,\varphi$ we look into the awareness set.

The main difference between $\mathcal{L}_A$ and our approach is the dynamics. First, our semantic model has a rule set function, indicating the processes the agent can use to increase her explicit information. It is not that she knows that after a rule application her information set will change; it is that she knows the *process* that leads the change. Second, $\mathcal{L}_A$ does not express changes in awareness sets (though later the authors add a relation over W to represent steps in time). Our approach uses inference as the process that extends explicit information, and represents it as a model operation modifying information sets. Third, the language of our information sets is less expressive than the

[6] The formulas cover two cases: deduction not using the new knowledge and deduction using it. See Sect. A.3 for comments about the proofs.

one of the awareness sets, but it allows us to define *updates* for representing external dynamics, a process not considered in $\mathcal{L}_A$.

5.2 Duc's dynamic epistemic logic

Duc proposed in 1997, 2001 a dynamic epistemic logic to reason about agents that are neither logically omniscient nor logically ignorant. He defined the language $\mathcal{L}_{BDE}$, based on formulas of the form $K\gamma$ ("γ *is known*") for γ a propositional formula, and closed under negation, conjunction and the diamond modal operator $\langle F\rangle$ ("φ *is true after some course of thought*"). Note how we cannot talk about the real world.

The models are tuples (W, R, Y) with R a relation over W and Y a function assigning a set of propositional formulas to each world. The definition asks for properties guaranteeing that a world where the agent is *logically omniscient* will be eventually reached via R-transitions. Semantics for negation, conjunctions and are standard; for $K\gamma$ formulas we look into the sets of formulas of the corresponding world and the operator $\langle F\rangle$ is interpreted as a diamond with R.

The framework does not represents implicit information and, while it does express changes in explicit information, it does it with a relation between worlds, different from our model operation approach. Also, the framework does not represent external dynamics, like our explicit observations.

5.3 Jago's logic for resource-bounded agents

In 2006a, 2006b, Jago presented a logic for resource-bounded agents. His semantic model extends Kripke models with a set of formulas and a set of rules in each possible world. He also considered rule-based inference as the mechanism for increasing explicit information.

There are two main differences in the approaches. First, Jago represents inference as a relation between worlds. Extending what we said before, our model-operation representation gives us a functional treatment of inference, while the relational representation forces us to ask for properties of the relation to get this behaviour. Some properties may need a more powerful language to be expressed (e.g., the uniqueness of the result of a rule application) and some others may be not preserved after updates (e.g., the existence of a world resulting from an available rule application). The second one is our external dynamics, not considered in Jago's work.

5.4 van Benthem's acts of realization

In van Benthem (2008b), the author considers a language similar to our $\mathcal{EI}$ but without formulas about rules. The semantic model is of the form $(W, W^{acc}, \sim, V)$ where $(W, \sim, V)$ is a Kripke model and W^{acc} is a set of *access worlds*: pairs (w, X) with $w \in W$ and X a set of factual formulas (the access set). Formulas are interpreted at access worlds in the usual way, with $I\gamma$ true at (w, X) *iff* $\gamma \in X$.

The author defines two model operations: *implicit observation* (removing worlds as an announcement in PAL) and *explicit observation* (removing worlds and also adding the formula to the access sets of the remaining ones, like our explicit information operation). Then, he notices that the two operations overlap in their effects on the model, and proposes two more "orthogonal" operations: one simply removing worlds ("*bare observation*") and another one simply adding true formulas to access sets (an "*act of realization*"). This *act of realization* is more general than our deduction: *any* formula belonging to the implicit information can be added to the access set; in particular, validities, can be added at any point. Our framework allows us to add a formula only if it is the conclusion of an applicable rule.

6 Final remarks and further work

Let us represent the restaurant example with our framework. The initial setting can be given by a model M with six possible worlds, each one of them indicating a possible distribution of the dishes, and all of them indistinguishable from each other.

For explicit knowledge, consider atomic propositions of the form $\mathsf{p_d}$ where p stands for a person (*f*ather, *m*other or *y*ou) and d stands for some dish (*m*eat, *f*ish or *v*egetarian). The waiter explicitly knows each person will get only one dish, so we can put the rules $\rho_1 : \{\mathsf{y_f}\} \Rightarrow \neg\mathsf{y_v}$, $\rho_2 : \{\mathsf{f_m}\} \Rightarrow \neg\mathsf{f_v}$ and similar ones in each world. Moreover, he explicitly knows that each dish corresponds to one person, so the rule $\sigma : \{\neg\mathsf{y_v}, \neg\mathsf{f_v}\} \Rightarrow \mathsf{m_v}$ can be also added, among many others. Let w be the real world, where $\mathsf{y_f}$, $\mathsf{f_m}$ and $\mathsf{m_v}$ are true. The formula $\neg I\ \mathsf{m_v} \wedge \neg\Box\mathsf{m_v}$, indicating that the waiter does not know (neither explicitly nor implicitly) that your mother has the vegetarian, is true at w.

While approaching the table, the waiter can increase the rules he knows. This does not give him new explicit facts, but it will allow him to infer faster later. From *Cut* over ρ_1 and σ, he gets $\varsigma_1 : \{\mathsf{y_f}, \neg\mathsf{f_v}\} \Rightarrow \mathsf{m_v}$. Then, $\langle Cut_{\rho_1,\sigma}\rangle \neg I\ \mathsf{m_v}$ and $\langle Cut_{\rho_1,\sigma}\rangle\ L\ \varsigma_1$ are also true at w. Moreover, he can apply *Cut* again, this time with ρ_2 and ς_1, obtaining the rule $\varsigma_2 : \{\mathsf{y_f}, \mathsf{f_m}\} \Rightarrow \mathsf{m_v}$ and making $\langle Cut_{\rho_1,\sigma}\rangle\ \langle Cut_{\rho_2,\varsigma_1}\rangle \neg I\ \mathsf{m_v}$ and $\langle Cut_{\rho_1,\sigma}\rangle\ \langle Cut_{\rho_2,\varsigma_1}\rangle\ L\ \varsigma_2$ true at w.

After the answer to the question "*Who has the fish?*", the waiter explicitly knows that you have the fish. Four possible worlds are removed, but he still does not know that your mother has the vegetarian. (We have $\langle Cut_{\rho_1,\sigma}\rangle\ \langle Cut_{\rho_2,\varsigma_1}\rangle\ \langle\mathsf{y_f}!\rangle\ (\neg I\ \mathsf{m_v} \wedge \neg\Box\mathsf{m_v})$ true at w).

Then he asks "*Who has the meat?*", and the answer removes one of the remaining worlds. Now he knows implicitly that your mother has the vegetarian dish and, moreover, he is able to infer it and add it to his explicit knowledge:

$$(M, w) \Vdash \langle Cut_{\rho_1,\sigma}\rangle\ \langle Cut_{\rho_2,\varsigma_1}\rangle\ \langle\mathsf{y_f}!\rangle\ \langle\mathsf{f_m}!\rangle\ (\Box\mathsf{m_v} \wedge \langle \mathrm{D}_{\varsigma_2}\rangle\ I\ \mathsf{m_v})$$

Two structural operations, two explicit observations and one inference are all that the waiter needs.

The proposal can be extended in several ways. The first one is by extending the internal language beyond the propositional one. As we mentioned, we chose it because it makes the definition of updates with *true* information (Sect. 4) possible. In general, a true observation in the full explicit/implicit information language cannot be simply added to an information set, since it may become false after being observed (witness Moore sentences, like $p \wedge \neg\Box p$). A first attempt would be to keep in the new information set those formulas that are true in the new model, but the definition would face circularity: the new information set should contains just the formulas that are still true, but in order to decide whether an explicit information formula $I\,\gamma$ is true or not, we need this new information set. A further analysis providing a solution to this limitation will greatly increase the expressivity of the framework.

We have analyzed the case in which the information is true, but this is not the general situation. By removing such restriction we can talk about *beliefs* (information no necessarily true). This would allow us to explore dynamics of these different notions (van Benthem 2007 provides an account for *belief revision* and Velázquez-Quesada 2009 provides dynamics for different notions of information). Moreover, besides truth-preserving inference, there are other inferences, like *default reasoning* or *abduction*. We can also represent them in order to study how all of them work together.

For the external dynamics, our finer representation of knowledge allows us to define different kinds of observations. Besides our explicit observations, we can also define implicit ones. A more expressive internal language would allow us to represent more kinds of observations, all of them differing in how introspective is the agent about the observation.

In the context of agent diversity (Liu 2008), our framework allows us to represent agents having different rules and therefore different reasoning abilities. The idea works also for external dynamics: agents may have different observational powers. It will be interesting to explore how agents with different reasoning and observational abilities interact.

Acknowledgements The author thanks Johan van Benthem for his invaluable ideas and suggestions that led to the present work. He also thanks the organizers and participants of the ILLC's *Seminar on Logics for Dynamics of Information*, as well as DGL'08 and LaII'08 workshops; their comments have helped to improve earlier versions.

Appendix

A Technical appendix

A.1 Closure of structural operations

We will prove the *truth for rules* property for the three operations. Note that in the three cases it is enough to show that the rules are truth-preserving in M because the truth-value of the translation depends just on the valuation, which is preserved by the operations.

Reflexivity Recall that $\varsigma_\delta = (\{\delta\}, \delta)$ and pick any $\rho \in Z'(w)$. If ρ is already in $Z(w)$, we have $(M, w) \Vdash \mathrm{TR}(\rho)$ since M is in $\mathbf{EI}_K$. Otherwise, ρ is ς_δ, but we obviously have $(M, w) \Vdash \delta \to \delta$.
Monotonicity Recall that $\varsigma' = (\mathrm{prem}(\varsigma) \cup \{\delta\}, \mathrm{conc}(\varsigma))$ and pick any $\rho \in Z'(w)$. If ρ is already in $Z(w)$, we have $(M, w) \Vdash \mathrm{TR}(\rho)$. Otherwise, ρ is ς' and we should have $\varsigma \in Z(w)$; therefore, we have $(M, w) \Vdash \bigwedge \mathrm{prem}(\varsigma) \to \mathrm{conc}(\varsigma)$ and hence $(M, w) \Vdash (\bigwedge \mathrm{prem}(\varsigma) \wedge \delta) \to \mathrm{conc}(\varsigma)$.
Cut Recall that $\varsigma' = ((\mathrm{prem}(\varsigma_2) - \{\mathrm{conc}(\varsigma_1)\}) \cup \mathrm{prem}(\varsigma_1)\, , \mathrm{conc}(\varsigma_2))$ and pick any $\rho \in Z'(w)$. If $\rho \in Z(w)$, we have $(M, w) \Vdash \mathrm{TR}(\rho)$ since M is in $\mathbf{EI}_K$. Otherwise, ρ is ς' and we have $\{\varsigma_1, \varsigma_2\} \subseteq Z(w)$.
Let $\bigwedge \mathrm{prem}(\varsigma')$ be true at w in M; then, every premise of ς' is true at w in M. This includes every premise of ς_1 and every premise of ς_2 except $\mathrm{conc}(\varsigma_1)$. But since every premise of ς_1 is true at w in M and ς_1 is in $Z(w)$, truth for rules of M tells us that $\mathrm{conc}(\varsigma_1)$ is true at w in M and hence every premise of ς_2 is true at w in M. Now, since ς_2 is in $Z(w)$, truth for rules of M tell us that $\mathrm{conc}(\varsigma_2)$, that is, $\mathrm{conc}(\varsigma')$, is true at w in M. Then we have $(M, w) \Vdash \mathrm{TR}(\varsigma')$.

A.2 Structural operations and deduction

The validity of the formulas follows from the bisimilarities between models stated below. In our case, the bisimulation concept extends the standard one by asking for related worlds to have the same information and rule set: given two models $M_1 = \langle W_1, R_1, V_1, Y_1, Z_1\rangle$ and $M_2 = \langle W_2, R_2, V_2, Y_2, Z_2\rangle$, a non empty relation $B \subseteq (W_1 \times W_2)$ is a bisimulation if and only if it is a standard bisimulation between $\langle W_1, R_1, V_1\rangle$ and $\langle W_2, R_2, V_2\rangle$ and, if Bw_1w_2, then $Y_1(w_1) = Y_2(w_2)$ and $Z_1(w_1) = Z_2(w_2)$.

Let $M = \langle W, R, V, Y, Z\rangle$ be a model in $\mathbf{EI}_K$, and take $w \in W$. Models of the form $M_{\mathbf{STR}\sigma}$ are the result of applying first the structural operation STR and then the deduction operation with rule σ, and analogously for models of the form $M_{\sigma\mathbf{STR}}$. In all cases, the bisimulation is the identity relation over worlds reachable from w.

Reflexivity. Let ς_δ be the rule $(\{\delta\}, \delta)$:
- If $\sigma \neq \varsigma_\delta$, then $(M_{\mathbf{Ref}(\delta)_\sigma}, w) \leftrightarroweq (M_{\sigma\,\mathbf{Ref}(\delta)}, w)$.
- If $\varsigma_\delta \in Z(w)$, then $(M_{\mathbf{Ref}(\delta)_{\varsigma_\delta}}, w) \leftrightarroweq (M_{\varsigma_\delta}, w)$.
- If $\delta \in Y(w)$, then $(M_{\mathbf{Ref}(\delta)_{\varsigma_\delta}}, w) \leftrightarroweq (M_{\varsigma_\delta\,\mathbf{Ref}(\delta)}, w)$.

Monotonicity. Let ς' be the rule $(\mathrm{prem}(\varsigma) \cup \{\delta\}, \mathrm{conc}(\varsigma))$:
- If $\sigma \neq \varsigma'$, then $(M_{\mathbf{Mon}(\delta,\varsigma)_\sigma}, w) \leftrightarroweq (M_{\sigma\,\mathbf{Mon}(\delta,\varsigma)}, w)$.
- If $\varsigma' \in Z(w)$, then $(M_{\mathbf{Mon}(\delta,\varsigma)_{\varsigma'}}, w) \leftrightarroweq (M_{\varsigma'}, w)$.
- If $\delta \in Y(w)$ and $\varsigma \in Z(w)$, then $(M_{\mathbf{Mon}(\delta,\varsigma)_{\varsigma'}}, w) \leftrightarroweq (M_{\varsigma\,\mathbf{Mon}(\delta,\varsigma)}, w)$.

Cut. Let ς' be the rule $(\,(\mathrm{prem}(\varsigma_2) - \{\mathrm{conc}(\varsigma_1)\}) \cup \mathrm{prem}(\varsigma_1),\ \mathrm{conc}(\varsigma_2)\,)$:
- If $\sigma \neq \varsigma'$, then $(M_{\mathbf{Cut}(\varsigma_1,\varsigma_2)_\sigma}, w) \leftrightarroweq (M_{\sigma\,\mathbf{Cut}(\varsigma_1,\varsigma_2)}, w)$.
- If $\varsigma' \in Z(w)$, then $(M_{\mathbf{Cut}(\varsigma_1,\varsigma_2)_{\varsigma'}}, w) \leftrightarroweq (M_{\varsigma'}, w)$.
- If $(\mathrm{prem}(\varsigma_1) \cup \{\mathrm{conc}(\varsigma_1)\}) \in Y(w)$ and $\varsigma_1 \in Z(w)$, then $(M_{\mathbf{Cut}(\varsigma_1,\varsigma_2)_{\varsigma'}}, w) \leftrightarroweq (M_{\varsigma_2\,\mathbf{Cut}(\varsigma_1,\varsigma_2)}, w)$.

The involved operations (structural ones and deduction) preserve worlds, accessibility relations and valuations. To show that the identity relation over worlds reachable from w is a bisimulation, we just need to show that they have the same information and rule set in both models.

Consider as an example the third bisimilarity for monotonicity. For information sets, take any γ in the information set of w at $M_{\mathbf{Mon}(\delta,\varsigma)_{\varsigma'}}$; by definition, either it was already in that of w at $M_{\mathbf{Mon}(\delta,\varsigma)}$ or else it was added by the deduction operation. In the first case, it is in w at M (structural operations do not modify information sets); then it is also in w at M_ς and hence it is in w at $M_{\varsigma\mathbf{Mon}(\delta,\varsigma)}$. In the second case, γ should be conc(ς'), but then we have the premises of ς' (and hence those of ς) in w at $M_{\mathbf{Mon}(\delta,\varsigma)}$. Then, they are already in w at M and, by hypothesis, we have ς in w at M, so conc$(\varsigma) =$ conc(ς') is in w at M_ς and hence it is in w at $M_{\varsigma\mathbf{Mon}(\delta,\varsigma)}$.

For the other direction, take γ in w at $M_{\varsigma\mathbf{Mon}(\delta,\varsigma)}$. Then it is in w at M_ς and therefore either it was already in w at M or else it was added by the deduction operation. In the first case, γ is preserved through the monotonicity and the deduction operations, and therefore it is in w at $M_{\mathbf{Mon}(\delta,\varsigma)_{\varsigma'}}$. In the second case, γ should be conc(ς), and then we should have prem(ς) and ς in the correspondent sets of w at M. By hypothesis we have δ in w at M, so we have all the premises of ς' in w at M and therefore they are also in w at $M_{\mathbf{Mon}(\delta,\varsigma)}$. Since we have ς in w at M, we have ς' in w at $M_{\mathbf{Mon}(\delta,\varsigma)}$ too. Hence, we have conc$(\varsigma') =$ conc(ς) in w at $M_{\mathbf{Mon}(\delta,\varsigma)_{\varsigma'}}$. The case for rules is similar.

Now suppose a world u is reachable from w through the accessibility relation at $M_{\mathbf{Mon}(\delta,\varsigma)_{\varsigma'}}$. Since neither the relations nor the worlds are modified by the operations, u is reachable from w at M and therefore u is reachable from w at $M_{\varsigma\mathbf{Mon}(\delta,\varsigma)}$, too. Now we use the coherence properties: since $\delta \in Y(w)$ and $\varsigma \in Z(w)$, we have δ and ς in the corresponding sets of u, and then we can apply the argument used for w to show that u has the same information and rule set on both models.

A.3 Explicit observation and deduction

Just as the case of structural operations and deduction, the validity of the formulas follows from the bisimilarities stated below.

If χ is a formula:	• If $\chi \notin$ prem(σ), then $(M_{\chi!_\sigma}, w) \leftrightarrows (M_{\sigma\,\chi!}, w)$. • If $\chi \in$ prem(σ) and $\chi \in Y(w)$, then $(M_{\chi!_\sigma}, w) \leftrightarrows (M_{\sigma\,\chi!}, w)$
If χ is a rule:	• If $\chi \neq \sigma$, then $(M_{\chi!_\sigma}, w) \leftrightarrows (M_{\sigma\,\chi!}, w)$. • If $\chi = \sigma$ and $\chi \in Z(w)$, then $(M_{\chi!_\sigma}, w) \leftrightarrows (M_{\sigma\,\chi!}, w)$.

The proof is similar to the case of structural operations and deduction, keeping in mind that observations remove worlds.

References

Baltag, A., Moss, L. S., & Solecki, S. (1999). *The logic of public announcements, common knowledge and private suspicious*. SEN-R9922. Amsterdam: CWI.

Baltag, A., & Smets, S. (2008). A qualitative theory of dynamic interactive belief revision. In G. Bonanno et al. (Eds.), *LOFT7* (pp. 13–60). Amsterdam: AUP.

Dretske, F. I. (1981). *Knowledge and the flow of information*. Oxford: Basil Blackwell.

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, Institut für Informatik, Universität Leipzig, Leipzig, Germany.

Fagin, R., & Halpern, J. Y. (1988). Belief, awareness, and limited reasoning. *Artificial Intelligence, 34*(1), 39–76.

Floridi, L. (2005). Consciousness, agents and the knowledge game. *Minds and Machines, 15*(3–4), 415–444.

Gerbrandy, J. (1999). Dynamic epistemic logic. In L. S. Moss, J. Ginzburg, & M. de Rijke (Eds.), *Logic, language and computation* (pp. 67–84). Stanford: CSLI.

Hintikka, J. (1962). *Knowledge and belief*. Ithaca, N.Y.: Cornell University Press.

Jago, M. (2006a). Logics for resource-bounded agents. Ph.D. thesis, Department of Philosophy, University of Nottingham, Nottingham, UK.

Jago, M. (2006b). Rule-based and resource-bounded: A new look at epistemic logic. In T. Ågotnes & N. Alechina (Eds.), *Proceedings of the workshop on logics for resource-bounded agents, ESSLLI 2006* (pp. 63–77), Malaga, Spain.

Liu, F. (2008). Changing for the better. Preference dynamics and agent diversity. Ph.D. thesis, ILLC, Amsterdam. ILLC Dissertation series DS-2008-02.

Moortgat, M. (1997). Categorical type logics. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language* (pp. 93–177). Amsterdam: Elsevier.

Plaza, J. A. (1989). Logics of public communications. In M. L. Emrich et al. (Eds.), *International symposium on methodologies for intelligent systems* (pp. 201–216). Tennessee: Oak Ridge National Laboratory.

Sim, K. M. (1997). Epistemic logic and logical omniscience. A survey. *International Journal of Intelligent Systems, 12*(1), 57–81.

van Benthem, J. (2006). Epistemic logic and epistemology: The state of their affairs. *Philosophical Studies, 128*, 49–76.

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

van Benthem, J. (2008a). Logic and reasoning: Do the facts matter? *Studia Logica, 88*(1), 67–84.

van Benthem, J. (2008b). Merging observation and access in dynamic logic. *Journal of Logic Studies, 1*(1), 1–17.

van Benthem, J. (2008c). Tell it like it is: Information flow in logic. *Journal of Peking University (Humanities and Social Science Edition), 1*, 80–90.

van Benthem, J., & Kooi B. (2004). Reduction axioms for epistemic actions. In R. Schmidt et al. (Eds.), *Advances in modal logic* (*Number UMCS-04-09-01 in Technical Report Series*) (pp. 197–211). Manchester: University of Manchester.

van Ditmarsch, H., van der Hoek, W., & Kooi B. (2007). *Dynamic epistemic logic. Synthese library series* (Vol. 337). Berlin: Springer.

van Ditmarsch, H. P. (2000). Knowledge games. Ph.D. thesis, University of Groningen. ILLC Dissertation Series DS-2000-06.

Velázquez-Quesada, F. R. (2009). *Dynamics in explicit and implicit information* (Submitted).