

A LOGIC FOR SUSPICIOUS PLAYERS: EPISTEMIC ACTIONS AND BELIEF-UPDATES IN GAMES

Alexandru Baltag
CWI, the Netherlands

ABSTRACT

This paper introduces a notion of ‘epistemic action’ to describe changes in the information states of the players in a game. For this, ideas are developed from earlier contributions. The ideas are enriched to cover not just *purely epistemic* actions, but also *fact-changing actions* (‘real moves’, e.g. choosing a card, exchanging cards etc.) and *nondeterministic actions* and *strategies* (conditional actions having knowledge tests as conditions). The author considers natural *operations with epistemic actions* and uses them to describe significant aspects of the interaction between beliefs and actions in a game. A logic is used that combines in a specific way a multiagent epistemic logic with a dynamic logic of ‘epistemic actions’. The author presents *a complete and decidable proof system* for this logic. As an application, the author analyses a specific example of a dialogue game (a version of the Muddy Children Puzzle, in which some of the children can ‘cheat’ by engaging in secret communication moves, while others may be punished for their credulity). Also presented is a sketch of a ‘rule-based’ approach to games with imperfect information (allowing ‘sneaky’ possibilities, such as cheating, being deceived and suspecting the others to be cheating).

1. INTRODUCTION

The subject of this paper is a logic that combines in a specific way a multiagent epistemic logic with a dynamic logic of ‘epistemic actions’. This work continues and improves on the ideas and techniques presented

in Baltag *et al.* (1998, 1999) and Baltag (1999, 2000). It deals with subjects arising from the work of Fagin *et al.* (1995) and it is related to the work in Gerbrandy and Groeneveld (1997), Gerbrandy (1999), Bleeker and van Eijk (2000), Ditmarsch (2000) and van Benthem (2000).

The basic idea is to try to understand and formalize the notion of *epistemic update* and study it as an object in itself, in full generality. Intuitively, an epistemic update is a way to model *changes* which may affect the epistemic structure of the world. Primarily, these are changes in the information states of various agents, *actions that change beliefs* (although they might also change the facts of the world). But as shall argue, to better understand the belief-changing effects of such actions, we need to think of the *belief components of the action itself*: the action's own epistemic structure. In particular, *actions in a game* – such as a simple 'legal' move, a secret or illegal move, a choice of a strategy, a choice of a belief about other player's strategies etc. – not only have effects on the player's beliefs, but they seem to involve *dynamic forms of beliefs: beliefs-as-actions*, which are in the same time *beliefs about actions* and *belief-changing actions*.

A rather standard and natural way to model epistemic updates is as input–output transition relations between epistemic states or models. This is a so-called 'relational' semantics (as the one introduced in the next section). But we would like to also study epistemic updates as objects in themselves, describing *general types of epistemic changes*, which can be described independently of the input and output states. There are various natural such types, most of which seem to fall under either one (or more) of the following few categories: (1) direct-information-gathering (learning by direct experiment, by 'seeing', 'hearing' or by introspection), (2) information-exchange by *communication* (sending/receiving messages, public announcements, interception of private messages etc.), (3) information-hiding (secret communication, lying, sending encrypted messages, other forms of deceiving actions), (4) information-loss and misinformation (being lied to, non-introspective learning, starting to hold wrong beliefs, having gratuitous suspicions). In general, the interesting type of actions that our system can capture are 'half-transparent-half-hidden-actions'. For example, a move in a game can be such that some players 'see' some part (or feature) of what is happening but not the whole move; nevertheless, if the 'move' is legal they will necessarily 'suspect' it, i.e. regard it as a possibility.

As announced, we model the seeming complexity of such actions by endowing them with an *internal epistemic structure*. First, we divide actions into two categories: *simple actions* and *general actions*. Simple actions are deterministic and their effects and appearance are 'uniform', i.e. independent of the context; the general actions are nondeterministic sums of simple actions and can be modelled semantically as finite sets of simple actions.

A simple action will be given by specifying three distinct pieces of information: (1) its *presupposition* or precondition of happening; this

refers to the actual world *before* the action, and it defines the applicability of this particular action to this particular world: not every action can happen in every world; (2) the action's '*content*', describing the way the action changes the 'facts' of the world; (3) the action's *possible appearances* to the agents; i.e. the agent's views or beliefs about the very action that is taking place. The preconditions are modelled as functions assigning to each action α some sentence pre_α . The meaning of this function is that action α is possible only in a state satisfying pre_α . The 'content' of a simple action, describing the factual change induced by the action, is given by a function \cdot_0 associating to each simple action α some set α_0 of atomic sentences with the meaning that the truth values of the atomic sentences $P \in \alpha_0$ are 'flipped', i.e. changed into the opposite values by the action α . The way we model the 'appearance' of a simple action is via epistemic 'possibility' relations between actions. Usually, epistemic accessibility relations in a Kripke structure are used to represent the uncertainty of each agent concerning the current state of the system. In a similar manner, we endow our actions with accessibility relations (called 'suspicion relations') to represent each agent's uncertainty concerning the current action taking place. So we consider arrows $\alpha \rightarrow_a \beta$ between actions α, β to denote the fact that, if the current action is α then agent a thinks that β may be the current action. In other words, action α 'appears' to a as being indistinguishable from β . (This is *not* necessarily an equivalence relation, as a might be deceived into thinking that the current action is not possible, so α itself might not be among his epistemic alternatives.)

As we shall see, one way to model the *update of a state by a simple action* is as an operation of '*conditional multiplication*' of the two Kripke structures (the static and the dynamic one): the space of output-states is taken to be a subset of the Cartesian product of the two structures, in which we have deleted the 'impossible pairs', i.e. the pairs (s, α) arising from input-states s which did *not* fulfill the preconditions of the action α . We endow this set of output-states with a Kripke structure, by taking the 'product arrows': $(s, \alpha) \rightarrow_a (t, \beta)$ iff $s \rightarrow_a t$ and $\alpha \rightarrow_a \beta$; finally, we use the change-functions to update the 'facts', i.e. the truth-values of the atomic sentences in the new states. As for the general actions, which are nondeterministic sums of simple actions, they will induce non-deterministic updates: namely, the update of a state by such a general action will be the *set* of all possible output-states, obtainable through updating the initial state by every simple term of the nondeterministic sum.

This semantics reflects the idea of 'multiplicating independent uncertainties'. I introduce natural operations with actions, develop a 'calculus of epistemic actions' and state a 'normal form' representation theorem.

As an application, I use this setting and the logic to study modified versions of The Muddy Children Puzzle: some children cheat, by sending

signals to tell their friends they are dirty; the others might not suspect it, which can lead to a totally wrong line of reasoning on their part, ending in a wrong answer; or they could be more cautious and suspicious, which allows them to use other agent's wrong answers to find the truth more quickly than in the classical puzzle. Another application is to games with imperfect information (and potential misinformation), in the context of which one can use epistemic actions to formalize a notion of 'rule-based' game, given not in the usual extensional tree form, but as a set of conditional actions, providing the rules and the moves of the game. I introduce strategies in the same rule-based manner and provide a formalization of (nonprobabilistic) Nash equilibrium in modal logic.

II. A LOGIC FOR EPISTEMIC ACTIONS

I introduce here a modal language to describe the update of epistemic structures by epistemic actions. Our language L is obtained by putting together standard epistemic logic (with 'common knowledge' operators) with a dynamic logic of epistemic actions. For agents a and sets of agents A , we have the standard epistemic modalities \Box_a (the *belief*, or *knowledge*, operator) and \Box_A^* (the *common belief*, or *common knowledge*, operator). The sentence $\Box_a\varphi$ will denote the fact that *agent a believes that φ* , while $\Box_A^*\varphi$ will mean that *φ is common knowledge among all the agents of the group A* . In addition, we inductively build a set of *action-expressions* to denote *epistemic actions*, i.e. 'programs' updating epistemic situations. We build complex action-expressions from basic ones, using dynamic-logic-type program constructions, which correspond to natural operations with epistemic actions. For each such action-expression α , we have a 'dynamic-logic'-type modality $[\alpha]\varphi$; the sentence $[\alpha]\varphi$ denotes the fact that *after action α , sentence φ becomes true*, or more precisely, that if α can be executed then every possible output-state satisfies φ .

II.1. Syntax

We assume as given a set $AtProp$ of atomic propositions, denoted by P, Q, \dots , and a finite set Ag of agents, denoted by a, b, \dots . As before, we use capital letters $A, B, \dots \subseteq Ag$ to denote *finite sets of agents*.

We define, by simultaneous recursion, a set L of *propositions* over $AtProp$ (propositions denoted by φ, ψ, \dots) and a set Act_L of *action-expressions over L* (expressions denoted by α, β, \dots):

$$\begin{array}{l} \varphi, \psi ::= P \quad | \quad \neg\varphi \quad | \quad \varphi \wedge \psi \quad | \quad [\alpha]\varphi \quad | \quad \Box_a\varphi \quad | \quad \Box_A^*\varphi \\ \alpha, \beta ::= flipP \quad | \quad ?\varphi \quad | \quad \alpha + \beta \quad | \quad \alpha \cdot \beta \quad | \quad \alpha^a \quad | \quad \alpha^{*A}. \end{array}$$

Informally, the meanings of our action – constructions are: ‘test φ ’ $?\varphi$ is the action that *tests the truth* of a proposition φ , i.e. the program which accepts an epistemic state as input iff φ is true (in which case it returns the same state). The action *flip* P changes the value of the atomic sentence P at the current state, leaving everything else unchanged. The sum $\alpha + \beta$ is the nondeterministic composition (‘sum’ or ‘choice’: perform either α or β) of the two actions, while $\alpha \cdot \beta$ is their *sequential composition* (‘product’: perform first α and then β). The action α^a is the action of *one-step (nonintrospective, not-necessarily-trueful) ‘learning’* (suspicion): ‘agent a suspects α ’, i.e. a starts to believe (without introspection) that some action α might be happening (while in reality no action happens, except for a getting suspicious). We choose to call this action ‘suspicion’ instead of learning, since its default assumption is that α did *not happen* (unless we change the default by first sequentially composing α with this suspicion as in $\alpha \cdot \alpha^a$, i.e. unless we explicitly mention that α did happen). So, by itself, this is an action which ‘appears’ to a as if α is happening, while in fact nothing is happening (and everybody else sees that nothing is happening). The action α^{*A} is the action of *mutual (common, or public) learning* of an action inside a given group: the agents in the group A *commonly (and truthfully) learn* that α is happening (and indeed α is actually happening). So this action is like α , but with the proviso that it is ‘transparent’ to all the agents in the group A (while this action’s appearance to all the other agents is the same as α ’s).

I shall use the following abbreviations, for sets A of agents and sets P of atomic propositions:

$skip =: ?(true)$, where $true$ is any universally true sentence

$$\sum_{i=1}^n \alpha_i =: \alpha_1 + \alpha_2 + \dots + \alpha_n$$

$$\prod_{i=1}^n \alpha_i =: \alpha_1 \cdot \alpha_2 \dots \alpha_n$$

$$\alpha^A =: \prod_{a \in A} \alpha^a \text{ (the action of general believing)}$$

$$\alpha^{+A} =: (\alpha^{*A})^A \text{ (the action of common, mutual believing)}$$

$$\alpha^{+a} =: \alpha^{+\{a\}} \text{ (the action of fully introspective believing)}$$

$$P! = ?P + ?(\neg P) \cdot (\textit{flip } P) \text{ (the action ‘make } P \text{ true’)}$$

$$\neg P! = ?(\neg P) + ?P \cdot (\textit{flip } P) \text{ (the action ‘make } P \text{ false’)}$$

$$\text{if } \varphi \text{ do } \alpha \text{ else } \beta = ?\varphi \cdot \alpha + ?(\neg\varphi) \cdot \beta.$$

Given the above intuitive interpretations of our action-constructions, it is natural to think that the mutual-learning construct will have the following fixed-point property:

$$(*) \quad \alpha^{*A} = \alpha \cdot \alpha^{+A} = \alpha \cdot (\alpha^{*A})^A = \alpha \cdot \prod_{\alpha \in A} (\alpha^{*A})^a.$$

In other words, to say that the group A commonly learns α is equivalent to saying that, first, α is really happening, and then each of the agents in A privately learns that the group A commonly learns α . We shall use this intuitive identity to justify our semantics for α^{*A} (and later we can check that the identity really holds, up to epistemic bisimilarity).

II.2. Relational semantics for our logic

I introduce here the notion of an *epistemic state*, which is at the basis of our semantics. We shall interpret our propositions as *properties (or sets) of epistemic states* and our action-expressions as *binary relations between epistemic states*.

Given as above the sets $AtProp$ and Ag , an *epistemic state* (or a *pointed Kripke model*) is a quadruple $s = (W, \{\cdot_a\}_{a \in Ag}, \cdot_0, v)$, where W is a set of *possible worlds* or states, $v \in W$ is a distinguished world, called the *actual world*, each \cdot_a (for $a \in Ag$) is a map $\cdot_a : W \rightarrow \mathcal{P}(W)$ called *appearance map for agent a* , and $\cdot_0 : W \rightarrow \mathcal{P}(AtProp)$, called the *(factual) content map*.

Since the atomic sentences $P \in AtProp$ are supposed to describe ‘facts of the world’, the factual content $w_0 \subseteq AtProp$ of a given world w will be interpreted as defining *the set of all ‘true’ facts of the world w* . (Usually, the same information is given by specifying a ‘valuation’ map $|\cdot| : AtProp \rightarrow \mathcal{P}(W)$, and then the factual content can be defined by putting $w_0 = \{P \in AtProp : w_0 \in |P|\}$. Clearly, the two approaches are equivalent: if we take factual content as basic, we can define the valuation by $|P| = \{w \in W : P \in w_0\}$.) For a world w , the set $w_a \subseteq W$ is called *the appearance of world w to agent a* and intuitively consists of all the worlds that are ‘indistinguishable’ from w to agent a : if the actual world is w then agent a thinks any of the worlds $w' \in w_a$ might be the actual one. The worlds $w' \in w_a$ are called the *epistemic alternatives of the world w* (for agent a). A binary relation $\rightarrow_a \subseteq W \times W$ of (*epistemic*) *indistinguishability for agent a* can be defined as:

$$w \rightarrow_a w' \quad \text{iff } w' \in w_a.$$

For the sake of generality, we don’t assume that these relations have any special properties (e.g. reflexivity, transitivity etc.): we would like to cover under our approach both false beliefs and true knowledge, and

both introspective and non-introspective beliefs.¹ Observe that, in the definition of an epistemic state, we can alternatively take the indistinguishability relations and the valuations as basic, define epistemic states as quadruples $(W, \{\rightarrow_a\}_{a \in Ag}, | \cdot |, \nu)$, and then define the appearance functions, by taking $w_a = \{w' : w \rightarrow_a w'\}$. Indeed, this corresponds to the more standard definition of Kripke structures in terms of ‘accessibility’ relations (and valuation maps).

We denote by *Mod* the class of all pointed models (i.e. epistemic states). We shall use systematic ambiguity to identify an epistemic state with its ‘top’ possible world; this is consistent, as long as we don’t reuse names of possible worlds. This allows us to ‘lift’ the functions \cdot_a and \cdot_0 (and so the relations \rightarrow_a) from *inside* a given model to functions defined *on* pointed models (epistemic states): for instance, for epistemic states s, s' we put $s \rightarrow_a s'$ iff, whenever we have $s = (W, \{\cdot_a\}_{a \in Ag}, \cdot_0, \nu)$, then we also have $s' = (W, \{\cdot_a\}_{a \in Ag}, \cdot_0, \nu')$, for some ν' s.t. $\nu \rightarrow_a \nu'$.

So we can freely talk about appearance maps \cdot_a and accessibility relations \rightarrow_a at the level of epistemic states (instead of worlds). This allows us to *abstractly specify an epistemic state* s , without giving any explicit epistemic model, but just by specifying two things: (1) the *content* s_0 , i.e. the set of atomic propositions P holding at s ; (2) for each agent a , the *appearance* s_a , i.e. the set of all states accessible from s via a -arrows.

The reason we choose to work with epistemic states and relations between them, instead of states in a given (fixed) Kripke structure (as is the more standard approach in modal logic and dynamic logic), has to do with the ‘open’ character of learning actions: they may ‘change’ the epistemic structure of the world in many (possibly infinitely many) ways; but on the other hand, we do not want to include all these possible output-states in the initial structure; on the contrary, we would like to keep our structures small for as long as possible, so in a given structure we only include the worlds that are considered as possible at a *given moment*; to model the output-states of actions that change the epistemic situation we will have to go beyond the input-structure, owing to the lack of enough states. No finite Kripke structure will suffice to model the iterated effects of our actions. Thus, we choose to model actions as *relations between Kripke models (epistemic states)*, instead of relations *inside* a given model.

As usual, we define an *knowledge model (state)*, or *S5-model (state)*, to be a model (state) in which all the accessibility arrows are equivalence

¹ This generality may in fact produce some confusion, e.g. by our free and naive use of the terms ‘knowledge’ and ‘belief’ throughout this paper. Let us clarify this here: most of the times we shall use the two terms as virtually synonymous, even talking about ‘possibly un-truthful learning’; but in fact, we shall sometimes make the difference, and stress the word ‘knowledge’ when we assume the S5-axioms. This should be clear from the context, although it might be helpful to notice that every time we do that, we immediately add the illuminating parenthesis: (S5).

relations, while an *S4-model (state)* is one in which all the relations are transitive and Euclidean.

To define common knowledge, we need to introduce *iterated accessibility relations* between epistemic states: for each group $A \subseteq Ag$ of agents, we define the relation \rightarrow_A^* between epistemic states, as the reflexive-transitive closure of the union $\bigcup_{a \in A} \rightarrow_a$; in other words: we have $s \rightarrow_A^* s'$ iff there exists an A -chain $s = s_0 \rightarrow_{a_0} s_1 \rightarrow_{a_1} \dots \rightarrow_{a_n} s'$, with $a_i \in A$ for every i . Correspondingly to the appearance map, we can now define an *iterated appearance map* $\cdot_A : Mod \rightarrow \mathcal{P}(Mod)$ for a given set of agents: $s_A = \{s' \in Mod : s \rightarrow_A^* s'\}$.

I now give the semantics, by simultaneously defining the following relations: a *truth-relation (satisfaction)* $\models \subseteq Mod \times L$ between epistemic states and formulas, and, for each action-expression $\alpha \in Act_L$, a binary *transition relation* $\Rightarrow^\alpha \subseteq Mod \times Mod$ between epistemic states. We read $\Rightarrow^\alpha s'$ as follows: *if the input-state is s then s' is one of the possible output-states of applying action α* . The definition is by double recursion, on the complexity of formulas and on the complexity of action-expressions:²

II.2 Truth

For propositional and epistemic modal operators we have the usual recursive conditions, while for the dynamic modalities we use the input-output labelled transition relations: the meaning of $[\alpha]\varphi$ is that every \Rightarrow^α -transition starting in the current state ends in a state satisfying φ . So we define $s \models \varphi$, by recursion on the complexity of $\varphi \in L$:

$s \models P$	iff	$P \in s_0$ for atomic sentences
$s \models \neg\varphi$	iff	$s \not\models \varphi$
$s \models \varphi \wedge \psi$	iff	$s \models \varphi$ and $s \models \psi$
$s \models \Box_a \varphi$	iff	$s' \models \varphi$ whenever $s \Rightarrow_a s'$
$s \models \Box_A^* \varphi$	iff	$s' \models \varphi$ whenever $s \Rightarrow_A^* s'$
$s \models [\alpha]\varphi$	iff	$s' \models \varphi$ whenever $s \Rightarrow^\alpha s'$.

II.2.2. Transition relations

For each action-expression, we define transition relations $s \Rightarrow^\alpha t$, which must reflect the above-mentioned intuitive meanings of our actions. The semantics for 'test' actions $?\varphi$, non-deterministic choice (union) $\alpha + \beta$ and sequential composition $\alpha \cdot \beta$ is essentially the standard one in dynamic logic. The action *flip* P will output a state that is completely similar to the input-state, except that the truth-value of P is reversed (from true to false or vice-versa). The action α^a of 'suspecting α ' will

²This use of double recursion for simultaneously defining truth and the transition relations is *not* a peculiarity of our logic: in fact, this applies as well to the standard semantics of dynamic logic, although this point is not usually stressed.

output a state t that is in every respect similar to the input-state s , except that the state's appearance to agent a has changed: namely, a thinks that α happened, so that a 's epistemic alternatives for the output-state t are precisely all the possible outputs of applying action α to all a 's epistemic alternatives for the input-state s . Finally, we use the above-mentioned (intuitively desirable) identity (*), saying that $\alpha^{*A} = \alpha \cdot \prod_{a \in A} (\alpha^{*A})^a$, to give the semantics for mutual learning. So the output-state t will again be a state that will be 'similar' to some output w of applying action α (to the same input s), namely similar with respect to the atomic facts and to its appearance to all the outsiders $b \notin A$. But only the appearance of t to the insider agents $a \in A$ will be different; namely, they are consciously and mutually learning α , so this mutual learning action α^{*A} is 'transparent' to all the insiders. Hence, their epistemic alternatives for the output-state t will come as the result of updating their own epistemic alternatives for every possible output of α (applied to the input s) with the very action α^{*A} of mutual learning which is taking place.

In the following, we use the notation $\Phi \Delta \Psi =: (\Phi \setminus \Psi) \cup (\Psi \setminus \Phi)$ for the *symmetrical difference of two sets* of atomic sentences (consisting of all sentences which are in one and only one of the two sets).

$s \Rightarrow^{? \varphi} t$	iff	$t = s$ and $s \models \varphi$;
$s \Rightarrow^{\Pi P} t$	iff	$t_0 = s_0 \Delta \{P\}$ (or equivalently $s_0 \Delta t_0 = \{P\}$), and $t_a = s_a$ for every a ;
$s \Rightarrow^{\alpha + \beta} t$	iff	$s \Rightarrow^{\alpha} t$ or $s \Rightarrow^{\beta} t$;
$s \Rightarrow^{\alpha \cdot \beta} t$	iff	$s \Rightarrow^{\alpha} t$ or $\Rightarrow^{\beta} t$, for some epistemic state w ;
$s \Rightarrow^{\alpha^a} t$	iff	$t_0 = s_0$, $t_b = s_b$ for every agent $b \neq a$ and $t_a = \{t' : s' \Rightarrow^{\alpha} t' \text{ for some } s' \notin s_a\}$;
$s \Rightarrow^{\alpha^{*A}} t$	iff	there exists some epistemic state w s.t.: $s \Rightarrow^{\alpha} w$, $t_0 = w_0$; $t_b = w_b$ for every agent $b \notin A$, $t_a = \{t' : w' \Rightarrow^{\alpha^{*A}} t' \text{ for some } w' \in w_a\}$, for every $a \in A$.

The last clause might appear to be circular, and in fact it is itself a *coinductive definition*, which must be understood as defining $\Rightarrow^{\alpha^{*A}}$ as the *largest relation* on epistemic states which satisfies the above given (fixed-point) property. This fully defines the semantics of our logic.

II.3. Preconditions, appearance, change, choice

I now define some useful auxiliary functions on action-expressions $\alpha \in Act_L$: the *precondition* pre_{α} , the *appearance* α_a of action-expression α to a given agent a , the *change* α_0 induced by α in the factual content of

the world (also called the *content* of α) and the *choice* set $|\alpha|$ (of all possible choices of simple deterministic ‘resolutions’ of α).

These notations are technically useful (for instance, in stating our axioms), but they also have some independent intuitive justification. In the next section, the intuitions associated with these functions will be used to provide an interesting alternative (but equivalent) semantics for our logic. But for now, they should be understood as simple *syntactic notations*.³

II.3.1. Precondition

The precondition function $pre : Act_L \rightarrow L$ associates with each action-expression a sentence, its *precondition* pre_α , which intuitively defines its *domain of application*:

$$\begin{aligned} pre_{? \varphi} &= \varphi \\ pre_{flipP} &= true \\ pre_{\alpha^a} &= true \\ pre_{\alpha + \beta} &= pre_\alpha \vee pre_\beta \\ pre_{\alpha \cdot \beta} &= pre_\alpha \wedge [\alpha]pre_\beta \\ pre_{\alpha^{*A}} &= pre_\alpha \end{aligned}$$

The intuitions underlying this notation are the following: the ‘test’ action $? \varphi$ can only happen in a state in which φ is true; a ‘pure change of facts’ $flipP$ or an action of ‘pure suspicion’ α^a can *always* happen (hence we assign them the universally true precondition); the sequential composition $\alpha \cdot \beta$ can happen only if, first, α can happen, and then, after α is executed, β can happen; in other words, the precondition of $\alpha \cdot \beta$ is the conjunction of the precondition of α and the sentence asserting that the execution of β makes true the precondition of β . Finally, the action α^{*A} (the truthful common-learning of α by the members of the group A) can happen if and only if α itself can happen.

³To stress the analogy with epistemic states, we shall use the same notations here for the content and appearance of an action-expression as for content \cdot_c and appearance \cdot_a functions for epistemic states. No confusion is possible: as mentioned, these are now just *syntactic notations*, while the corresponding functions for states were semantic objects describing their Kripke structures. As announced, in the next section, we shall convert these syntactic notions into semantical ones, and we shall *use again the same notations* to denote the corresponding notions of ‘content’ and ‘appearance’ of real, semantic actions. There as here, these apparently ambiguous notations will be consciously used to reinforce the analogy between epistemic actions and epistemic states, but there won’t be any possibility of real confusion: the functions in the next section will be defined on *epistemic actions*, semantic objects which are formally distinct both from epistemic states (subjects of our first definitions above for content and appearance) and from action-expressions (for which the announced syntactic notations are introduced here).

To relate this syntactic definition to our semantics, denote the *domain* of a relation $R \subseteq Mod \times Mod$ by $dom(R) = \{s \in Mod : sRt \text{ for some } t \in Mod\}$, the *R-image* of a state s : $R(s) = \{t \in Mod : sRt\}$, and the *interpretation* of a sentence φ by $\|\varphi\| = \{s \in Mod : s \models \varphi\}$ (the class all epistemic states satisfying the sentence). Then we can easily observe that: $dom(\Rightarrow^? \varphi) = \|\varphi\|$, $dom(\Rightarrow^{flipP}) = dom(\Rightarrow^{\alpha^a}) = Mod$, $dom(\Rightarrow^{\alpha + \beta}) = dom(\Rightarrow^\alpha) \cup dom(\Rightarrow^\beta)$, $dom(\Rightarrow^{\alpha^{*A}}) = dom(\Rightarrow^\alpha)$, and $dom(\Rightarrow^{\alpha \cdot b}) = dom(\Rightarrow^\alpha) \cap \{s \in Mod : \Rightarrow^\alpha(s) \subseteq dom(\Rightarrow^\beta)\}$. These identities justify the above definition, and indeed one can easily check by induction that for every action-expression α we have: $\|pre_\alpha\| = dom(\Rightarrow^\alpha)$.

11.3.2. Appearance of an action

Given an agent a , we define a function $\cdot_a : Act_L \rightarrow Act_L$, giving the *appearance of action-expression α to agent a* . The intended intuitive interpretation of the action-expression α_a is the ‘apparent action’ from a ’s point of view: the way the action denoted by α appears to agent a :

$$\begin{aligned}
(? \varphi)_a &= skip \\
(flipP)_a &= skip \\
(\alpha^a)_a &= \alpha \\
(\alpha^a)_b &= skip \text{ (for } b \neq a) \\
(\alpha + \beta)_a &= \alpha_a + \beta_a \\
(\alpha \cdot \beta)_a &= \alpha_a \cdot \beta_a \\
(\alpha^{*A})_a &= \alpha_a \cdot \alpha^{*A} \text{ (for } a \in A) \\
(\alpha^{*A})_b &= \alpha_b \text{ (for } b \notin A).
\end{aligned}$$

The intuition behind this definition is the following: a ‘pure test’ $? \varphi$ or a pure ‘change of facts’ $flipP$ have no intrinsic epistemic effect, since they are unobservable by the agents (who will thus think that nothing, i.e. *skip*, happens). The action-expression α^a is supposed to represent the action in which agent a thinks that α is happening; so the appearance of this action to agent a is precisely α ; on the other hand, α^a is a ‘private’ epistemic action of agent a , an action which cannot be observed by any outsider; so its appearance to any other agent $b \neq a$ is *skip*: outsiders think nothing happens. The appearance of a non-deterministic sum (‘either α or β ’) to an agent a is the nondeterministic sum of the two appearances; similarly, the sequential composition of two actions (‘ α followed by β ’) appears to an agent as the sequential composition of the appearances of the two actions. Finally, we can use

the intuitive identity (*) to obtain the appearance of the mutual-learning action $\alpha^{*A} = \alpha \cdot \prod_{a \in Ag} (\alpha^{*A})^a$: its appearance to the ‘insiders’ is the same as the appearance of α followed by this very action of mutual-learning α^{*A} ; while to the ‘outsiders’, this appearance is exactly the same as the appearance of α itself: they learn nothing more.

In order to define a notion of ‘content’ (factual change) of an action-expression and a notion of non-deterministic choice, we first need to introduce a notion of *simplicity*. As announced in the introduction, ‘simple’ actions are deterministic actions which have ‘uniform’ appearance and ‘uniform’ effects on the facts of the world, in the sense that the appearance and effect are independent of the current state. This will be made precise in the next section, but for now it is enough to syntactically define simple expressions as the ones which do not contain any ‘real’ non-determinism (any *non-epistemic occurrences* of $+$, i.e. occurrences outside the scope of a pure suspicion operator), although they may contain ‘epistemic non-determinism’ (i.e. $+$ is allowed inside the scope of such epistemic operators).

II.3.3. Simple action-expressions

The set Act_L^0 of *simple* action-expressions (expressions denoted by $\sigma, \rho, \sigma', \dots$) is a subset of Act_L , inductively defined by:

$$\sigma, \rho ::= flipP \mid ?\varphi \mid \sigma \cdot \rho \mid \alpha^a \mid \sigma^{*A}$$

where α is any arbitrary (not necessarily simple) action-expression.

It can be easily checked that the transition relations \Rightarrow^α corresponding to simple action-expressions are always *deterministic* (i.e. they are partial functions). Moreover, one can also check that these relations change the truth-values of the atomic facts in a uniform manner, independent of the input-state. Using the same notation as above $\Phi \Delta \Psi =: (\Phi \setminus \Psi) \cup (\Psi \setminus \Phi)$ to denote the symmetrical difference of two sets of atomic sentences, we can easily check that: if we have both $s \Rightarrow^\alpha t$ and $s' \Rightarrow^\alpha t'$, then we also have that $s_0 \Delta t_0 = s'_0 \Delta t'_0$. This shows the ‘uniformity’ of simple action’s effects, and allows us to define the following notion.

II.3.4. Content (fact-change effect) of a simple action

We introduce a function $\cdot_0 : Act_L^0 \rightarrow \mathcal{P}(AtProp)$, called the *content*, or the *change*, function. For a simple expression σ , its content σ_0 will consist precisely of the atomic facts whose truth-values are changed by the transition relation \Rightarrow^α : whenever we have $s \Rightarrow^\alpha t$, we will also have that $\sigma_0 = s_0 \Delta t_0$ (where Δ is again the symmetrical difference). But we can define this function in a purely syntactical manner, by induction on

simple action-expressions:

$$\begin{aligned} (? \varphi)_0 &= \emptyset \\ (flip P)_0 &= \{P\} \\ (\alpha^a)_0 &= \emptyset \\ (\sigma \cdot \rho)_0 &= \sigma_0 \Delta \rho_0 \\ (\sigma^{*A})_0 &= \sigma_0. \end{aligned}$$

The intuitive meaning of this definition is the following: a ‘test’ action $? \varphi$ (if possible at all) or a ‘pure suspicion’ action α^a do not change in any way the ‘facts’ (the objective state of the world); the action denoted by $flip P$ changes only one fact, namely the truth-value of P (from true to false and vice-versa). A sequential composition $\sigma \cdot \rho$ changes first the (truth-values of all the) ‘facts’ that (the action denoted by) σ would change; then it changes (the truth-values of) all the facts that ρ would change. As a result, the facts that both σ and ρ would change remain unchanged (since their truth-values are twice flipped); similarly, the facts that neither of the two actions would change remain unchanged; while the truth-values of the facts changed by one and only one of the two actions are flipped. Finally, the ‘objective content’ of σ^{*A} is the same as that of σ : the effect of this action on the facts of the world is the same as the effect of action σ .

11.3.5. Choice (resolution of non-determinism)

I define a *choice function* $|\cdot| : Act_L \rightarrow \mathcal{P}(Act_L^0)$, taking general action-expressions into *sets of simple action-expressions*. Since simple action-expressions always denote deterministic actions, the choice set $|\alpha|$ can be understood as the set of *all possible simple deterministic ‘resolutions’* of our nondeterministic action.

$$\begin{aligned} |\sigma| &= \{\sigma\} \text{ (for } \sigma \in Act_L^0\text{)} \\ |\alpha + \alpha'| &= |\alpha| \cup |\alpha'| \\ |\alpha \cdot \alpha'| &= \{\sigma \cdot \sigma' : \sigma \in |\alpha|, \sigma' \in |\alpha'|\} \\ |\alpha^A| &= \{\sigma \cdot \alpha^{*A} : \sigma \in |\alpha|\} \end{aligned}$$

In other words, simple actions are their own (unique) ‘resolution’: there is no real choice to be made; the nondeterministic sum $\alpha + \alpha'$ ‘sums’ up all the choices that are possible in either α or α' ; so it can be resolved in any of the ways α or α' are resolved. The sequential composition of nondeterministic actions $\alpha \cdot \alpha'$ can be resolved by resolving first action α

and then action α' . Finally, by (*), $\alpha^{*A} = \alpha \cdot (\alpha^{+A})$, where α^{+A} is a simple action-expression (being a product of simple expressions); so indeed, we can use the resolution of this product to get a resolution of the mutual-learning action α^{*A} .

I can also introduce an associated *choice relation* $\rightarrow \subseteq Act_L \times Act_L^0$, describing the 'choice' of some simple action, i.e. the transition from a possibly nondeterministic action-expression α to any of its simple components. This is defined by putting:

$$\alpha \rightarrow \sigma \text{ iff } \sigma \in |\alpha|.$$

Alternatively, one can inductively define it by:

$$\begin{aligned} &\sigma \rightarrow \sigma \text{ (for every } \sigma \in Act_L^0) \\ &\text{if } \alpha \rightarrow \sigma \text{ then } \alpha + \alpha' \rightarrow \sigma \\ &\text{if } \alpha' \rightarrow \sigma' \text{ then } \alpha + \alpha' \rightarrow \sigma' \\ &\text{if } \alpha \rightarrow \sigma \text{ and } \alpha' \rightarrow \sigma' \text{ then } \alpha \cdot \alpha' \rightarrow \sigma \cdot \sigma' \\ &\text{if } \alpha \rightarrow \sigma \text{ then } \alpha^{*A} \rightarrow \sigma \cdot \alpha^{+A}. \end{aligned}$$

II.3.6. Epistemic alternatives of simple action-expressions

By analogy with epistemic states (remembering that, for a state s , its appearance s_a is the set of all its epistemic alternatives), we can formally define now the *agent a 's epistemic alternatives for a simple action-expression σ* to be all the elements of the set $|\sigma_a|$ (i.e. all the simple resolutions of σ 's appearance to a). Correspondingly (as in the case of epistemic states), we can define (*epistemic*) *indistinguishability arrows between simple action-expressions* $\rightarrow_a \subseteq Act_L^0 \times Act_L^0$ for each agent a :

$$\sigma \rightarrow_a \sigma' \text{ iff } \sigma' \in |\sigma_a|.$$

One can easily see that these relations have the following properties (which can alternatively be taken as providing an inductive definition of the epistemic arrows between simple expressions):

$$\begin{aligned} &? \varphi \rightarrow_a \text{skip} \\ &\text{flip} P \rightarrow_a \text{skip} \\ &\text{if } \sigma \rightarrow_a \sigma' \text{ and } \rho \rightarrow_a \rho' \text{ then } \sigma \cdot \rho \rightarrow_a \sigma' \cdot \rho' \\ &\text{if } \alpha \rightarrow \sigma \text{ then } \alpha^a \rightarrow_a \sigma \\ &\text{if } b \neq a \text{ then } \alpha^a \rightarrow_b \text{skip} \\ &\text{if } \sigma \rightarrow_a \sigma' \text{ and } a \in A \text{ then } \sigma^{*A} \rightarrow_a \sigma' \cdot \sigma^{*A} \\ &\text{if } \sigma \rightarrow_b \sigma' \text{ and } b \notin A \text{ then } \sigma^{*A} \rightarrow_b \sigma'. \end{aligned}$$

Iterated epistemic alternatives of simple expressions

Again by analogy with epistemic states, we can introduce iterated epistemic relations between simple action-expressions: for each group

$A \subseteq Ag$ of agents, we define the relation \rightarrow_A^* as the reflexive-transitive closure of the union $\cup_{a \in A} \rightarrow_a$. (In other words, $\sigma \rightarrow_A^* \sigma'$ iff there exists a finite chain of A -arrows linking σ and σ' .) Also, put:

$$|\sigma|_A = \{\sigma' : \sigma \rightarrow_A^* \sigma'\}.$$

This is called the *set of all A -iterated epistemic alternatives* of the simple action-expression σ . It is important to observe that *both* $|\sigma_a|$ and $|\sigma|_A$ *are always finite* (if, as we have already have assumed, the set Ag of all agents is finite).

I stress once again that the concepts of precondition, appearance, (iterated) epistemic alternatives, content (change) and choice, as defined above, are all just convenient *syntactic notations* for finite sets of expressions, or finite-image relations between expressions. Nevertheless, these notations institute a formal analogy between simple action-expressions (as syntactical objects) and epistemic states (which are semantical objects): they have both a ‘factual content’ (which is a set of atomic facts) and an ‘appearance’ (a set of epistemic alternatives) for each agent. As we have seen, epistemic states are completely determined by these two pieces of information (*content, family of all appearances*). One can easily see that to completely determine the effect of a simple action we need a third piece of information: its precondition. This suggests that we could think of the semantic counterpart of a simple action-expression σ (i.e. its underlying transition relation \Rightarrow^σ) as being something very much like an epistemic state: a triplet (*content, appearances, precondition*). Indeed, these intuitions will be used in the next section to provide an *alternative semantics* for our logic, one that is *closer in spirit and in structure to our syntax* than the relational semantics.

III. THE PRODUCT SEMANTICS

I shall give now an alternative, but equivalent, semantics for this logic. We consider this semantics as having its own independent motivation, as well as heuristical, philosophical and technical importance. It is an improvement of the semantics first introduced in Baltag *et al.* (1998) and developed in Baltag (1999, 2000). It is interesting that the origins of this semantics are related to the highly technical work in Baltag *et al.* (1998) on the completeness and decidability of epistemic action logics. The main ideas for this semantics occurred as a side-effect of attempts to axiomatize the interplay of knowledge, common knowledge and action.

The basic concept is that of a *simple epistemic action*, which will be the semantic counterpart of our simple action-expressions. Roughly

speaking, this concept is a dynamic analogue of the notion of *epistemic state*. The intuition is that an action can have *different appearances* to various agents, which we can model in a similar manner to the one used for epistemic states: namely, as sets of *epistemic alternatives* for each agent. (But of course the alternatives of an action are themselves possible actions, not states.) Each epistemic alternative for the output-state will come as the output of an epistemic alternative of the current action applied to an epistemic alternative of the input-state. This is the idea of a 'product-semantics': the uncertainties regarding the state and the ones regarding the action are to be multiplied. The resulting 'static' Kripke structure (of the output-state) is a product of the initial 'static' Kripke structure (of the input-state) with the given 'dynamic' Kripke structure (of the current action).

But this idea cannot be generally applied to every action: it assumes that the two uncertainties (about the current action and about the current state) are *independent*. One way this can fail is due to different action's limited *domains of application*: some actions may not be applicable to some states. In the worst case, even if the real action is applicable to the real state, some given epistemic alternative of the action might be incompatible with some of, or even all, the epistemic alternatives of the input state! In fact, *real learning is based on this phenomenon*: increase of knowledge can only come by dropping some of the prior epistemic alternative-states, i.e. by narrowing the range of possibilities. As shown in Baltag *et al.* (1999), this phenomenon can be easily taken care of by endowing our (simple) epistemic actions with *preconditions*, i.e. propositions which define their domain of application. Consequently, we have to 'prune' the above product of the two structures, by deleting all the impossible outputs (of possible actions applied to states outside their domains). The result is a restricted product operation.

But there is another way this principle can fail, owing to *nondeterminism*, or more generally, to what we will call *nonuniformity*. Take for instance a conditional, if-then-else action: φ do β . In any given context, this is not in fact nondeterministic, but both its simple effects (e.g. whether or not it 'flips' the truth-value of some atom P) and its 'appearance' (i.e. what are agent a 's epistemic alternatives for this action) may depend on the current state (or, more precisely, they will depend on whether or not the current state satisfies the condition φ). This is the 'nonuniformity' (of effects or of appearance) of our action.

Consequently, a general action-expression (involving choice $+$) cannot be interpreted as a simple action (i.e. one having the above internal epistemic Kripke structure). But by defining (general) *epistemic actions* as being *sets of simple actions*, we can interpret all our expressions α as epistemic actions $\|\alpha\|$. This is completely similar to the way we can interpret our *formulas* φ as *propositions*, i.e. sets, or classes, of epistemic states: $\|\varphi\| = \{s \in Mod : s \models \varphi\}$.

III.1. Epistemic action models

An *action-model* $(K, \{\cdot_a\}_{a \in Ag}, \cdot_0, pre)$ consists of a *finite* multiagent Kripke model $(K, \{\cdot_a\}_{a \in Ag}, \cdot_0)$ and a precondition map $pre: K \rightarrow L$, mapping each element of K to a sentence of our language. We also require the Kripke model to be *serial*: $k_a \neq \emptyset$ for any $a \in Ag$. (We do this just because the interpretations of all our action-expressions happen to be serial: no agent ‘dies’ in our epistemic actions.) To distinguish it from state-models, we call the ‘possible worlds’ $k \in K$ of an action-model ‘*possible action-tokens*’, while the analogue of the ‘factual content’ map $\cdot_0: K \rightarrow \mathcal{P}(AtProp)$ will be called the *change-function* (or just ‘content’ map) of the model. (We interpret $P \in k_0$ as encoding the fact that action-token k always ‘flips’ the truth-value of P , from true to false and vice-versa.) As before, the set k_a is called the *appearance of (action-token) k to agent a* , while its elements are called the *agent a ’s epistemic alternatives for k* . As for possible worlds, we can introduce an epistemic accessibility relation between action-tokens $k \rightarrow_a k'$ (called the ‘suspicion relation for agent a ’), by defining it as $k' \in k_a$. Finally, $pre_k = pre(k) \in L$ will be called the *precondition* (or presupposition) of action-token k .

A *simple action* is just a ‘pointed action-model’, i.e. a tuple $\sigma = (K, \{\cdot_a\}_{a \in Ag}, \cdot_0, pre, k)$, composed of (the components of) an action-model and a designated action-token $k \in K$, called the ‘the real action’. I denote by Act^0 the class of all epistemic actions. (Observe that, *except for the precondition function, a simple epistemic action is the same kind of formal object as an epistemic state!*) As for epistemic states, we use systematic ambiguity to ‘lift’ the functions \cdot_0, \cdot_a, pre and the relations \rightarrow_a from inside action-models to the level of simple epistemic actions. As before, this allows us to specify a simple epistemic action σ by just giving *three pieces of information*: the action’s *precondition* $pre_\sigma \in L$ (defining its domain), the action’s *content* (change-set) $\sigma_0 \subseteq AtProp$ (specifying which atomic sentences have their truth-values ‘flipped’) and the action’s *appearance* σ_a to each agent a .

A (*general*) *epistemic action* is just a *finite set* $\alpha \subseteq Act^0$ of simple actions.⁴ We put $Act = \mathcal{P}_{fin}(Act^0) = \{\alpha \subseteq Act^0 : \alpha \text{ is finite}\}$ to be the set of all epistemic actions. The *choice relation* $\rightarrow \subseteq Act \times Act^0$ is defined as the converse of the membership relation: $\alpha \rightarrow \sigma$ iff $\sigma \in \alpha$.

As anticipated in the previous section, the intuition is that: $\alpha \rightarrow \sigma$ means that the nondeterminism of α can be resolved by choosing the

⁴Yes, it might be confusing, but it’s formally true: a simple action is formally *not* a (general) epistemic action (but it can be an element of an epistemic action). But in practice, I won’t stress the difference between the simple action σ and the epistemic action $\{\sigma\}$. This also explains why I will be reusing the σ ’s later to denote strategies, which (far from being simple) are in fact rather complex epistemic actions. But I do hope that by then there won’t be left any possibility of confusion.

simple action σ ; that a simple action σ is possible only if pre_σ is true; that the action σ changes the facts of the world by ‘flipping’ all the truth-values of the atomic sentences in σ_0 (while leaving the others unchanged); finally, σ_a gives the ‘appearance’ of action σ to agent a , i.e. it is the set of all σ ’s epistemic alternatives from a ’s point of view: if σ were the real action happening, then agent a would believe that the nondeterministic action σ_a is happening. In the case that σ_a is really a non-deterministic action (i.e. a set of at least two simple actions), then we interpret this as *epistemic uncertainty*: a suspects that any of the simple actions in σ_a might in fact be happening. So, *epistemically*, a deterministic action may ‘look’ like a non-deterministic one.

III.2. Interpretation of an action-expression

I give now the semantics of our action-expressions in terms of action-models. We associate with each simple expression σ a simple action $\bar{\sigma} \in Act^0$, called the *correspondent* of σ ; simultaneously, we define, for each action-expression α , an *interpretation* $\|\alpha\| \in Act$. (The interpretation of simple actions σ will be just the singleton $\{\bar{\sigma}\}$.) First, to define $\bar{\sigma}$ for simple actions $\sigma \in Act_L^0$, we put:

$$\begin{aligned} pre_{\bar{\sigma}} &= pre_\sigma \\ \bar{\sigma}_0 &= \sigma_0 \\ \bar{\sigma}_a &= \|\sigma_a\|. \end{aligned}$$

For general epistemic actions $\alpha \in Act_L$, we define:

$$\|\alpha\| = \{\bar{\sigma} : \sigma \in \alpha\}.$$

This completely specifies the correspondent simple action $\bar{\sigma}$ and the interpretation $\|\alpha\|$. Roughly speaking, the interpretation map is simply taking the syntactic notations introduced in the previous section and making them into a semantics.

III.3. Truth and update

The promised alternative semantics for our logic can be given by simultaneously defining three functions: *update of a state by a simple action* (a partial function $.: Mod \times Act^0 \rightarrow Mod$), *update of a state by a general epistemic action* $.: Mod \times Act \rightarrow \mathcal{P}(Mod)$ and the *interpretation (or truth-set) of a formula* $\|\cdot\| : L \rightarrow \mathcal{P}(Mod)$. The update $s.\alpha$ of a given state s with a simple action σ gives the (unique, if at all existing) output-state resulting from applying the action to the input-state. The update $s.\alpha$ of a state with a general epistemic action gives the set of all possible

output-states that can result from the execution of α on s . Finally, the interpretation of a formula gives the class of all epistemic states satisfying the formula.

First, *the simple update*: for states $s \in Mod$ and simple actions $\sigma \in Act^0$, we put

$$s.\sigma \text{ is defined iff } s \in \|\|pre_\sigma\|\|$$

$$(s.\sigma)_0 = s_0 \Delta \sigma_0$$

$$(s.\sigma)_a = \{s', \sigma' : s' \in s_a, \sigma' \in \sigma_a\}.$$

This is indeed a formalization of the above-mentioned idea of ‘multiplying the uncertainties’: after using the precondition function to eliminate the ‘impossible outputs’ (of simple actions applied to inputs which do not satisfy their preconditions) and using the content function to appropriately change the facts of the input-state, we describe the appearance of the output-state to each agent as the ‘product’ of the two appearances (of the initial state and of the action) to the same agent. In other words, the epistemic alternatives of the output are all the consistent outputs of applying the epistemic alternatives of the action to the epistemic alternatives of the input. This is indeed a sort of restricted product of the two Kripke structures.

Next, *the general update*: for general epistemic actions $\alpha \in Act$, we define

$$s.\alpha = \{s.\sigma : \sigma \in \alpha\}.$$

This formalizes the idea that the output of a non-deterministic action is just the set of the possible outputs of all its simple deterministic resolutions.

Finally, *the interpretation* $\|\| \varphi \|\|$ of a formula $\varphi \in L$ is defined by:

$$\|\| P \|\| = \{s \in Mod : P \in s_0\}$$

$$\|\| \neg \varphi \|\| = \{s \in Mod : s \notin \|\| \varphi \|\|\}$$

$$\|\| \varphi \wedge \psi \|\| = \|\| \varphi \|\| \cap \|\| \psi \|\|$$

$$\|\| \square_a \varphi \|\| = \{s \in Mod : s_a \subseteq \|\| \varphi \|\|\}$$

$$\|\| \square_A \varphi \|\| = \{s \in Mod : s_A \subseteq \|\| \varphi \|\|\}$$

$$\|\| [\alpha] \varphi \|\| = \{s \in Mod : s.\alpha \subseteq \|\| \varphi \|\|\}$$

This is just the extensional version of our previous definition of truth for the logic of epistemic actions.

We can easily check that *our two semantics are equivalent*, in the following sense.

Proposition 3: For every $s, s' \in Mod$, $\alpha \in Act_L$, $\sigma \in Act_L^0$ and $\varphi \in L$, we have:

$$\begin{aligned} s \Rightarrow^\alpha s' &\text{ iff } s' \in s. \|\alpha\| \\ s \Rightarrow^\alpha s' &\text{ iff } s.\bar{\sigma} = s' \\ s \models \varphi &\text{ iff } s \in \|\varphi\|. \end{aligned}$$

In conclusion: the *purpose* of introducing this alternative semantics is to have a semantical notion of action which can capture in a compact way general types of epistemic change. Unlike transition relations between epistemic states, our epistemic actions are *finite objects*, which nevertheless describe changes that can affect infinitely many epistemic states: usually, they can be applied again and again, and their domain of action is usually a proper class of states. Given any finite input-state and any epistemic action, we can easily compute the output via the above-described update operation (a ‘product’ of the two Kripke structures). As we shall see later, a rule-based *game* can be specified semantically by giving a finite set of epistemic actions (together with some winning conditions), while a *game-playing situation* is specified by giving a game and an initial epistemic state.

IV. EXAMPLES AND PROPERTIES OF EPISTEMIC ACTIONS

To give some examples of action-models corresponding to natural actions, let us fix our set of agents $Ag = \{a, b, c\}$.

IV.1. (Private, truthful, conscious, introspective) learning

Agent a learns (discovers) that some proposition φ is true. The act of learning is done *in private*: while it is happening, nobody else knows, or even suspects, that it is happening. (Accordingly, after this action, agents b and c remain in the same information-state as before.) The act of learning is indeed learning and not just a belief-revision, in the sense that it is *truthful*: φ is actually true. The act of learning is *conscious* and *introspective*, in the sense that agent a knows what she is doing and knows that nothing else happens in the meantime.

This action α can be represented in our language as conscious-introspective-truthful-and-secret learning action: $\alpha = (? \varphi)^{*a}$. In terms of action models, it can be described by a structure with two action-tokens, $K = \{k, l\}$. Here k represents the ‘real’ action that is taking place (learning of φ by agent a), action which has as presupposition the truth of φ , $pre_k = \varphi$: one cannot truthfully learn something false. (If we wanted to model a notion of ‘truthful and informative (non-redundant) learning,

we would have to add as extra-presupposition the fact that agent a doesn't know φ before the action, i.e. we would put $pre_k = \varphi \wedge \neg \Box_a \varphi$.) On the other hand, l represents the action that agents b and c think that is taking place, namely nothing: $l = skip$ will just be the 'trivial' action in which nothing changes. This trivial action can 'happen' anywhere, $pre_l = true$. Also, the trivial action is completely 'transparent', in the sense that, if it happens, then everybody knows it is happening; so it is its own only successor: $l \rightarrow_a l, l \rightarrow_b l, l \rightarrow_c l$ (and no others). On the contrary, action-token k 'looks like' the trivial one l from the point of view of b and c , i.e. $k \rightarrow_b l, k \rightarrow_c l$, while the same action-token k is 'transparent' to a , who knows that k is happening, so she considers k as its own only alternative: $k \rightarrow_a k$.

In Figure 1 the action-tokens are represented by boxes that surround their own presuppositions and the star is used to mark the designated 'top' action-token (the 'actual action'). We do not explicitly draw the choice relation and the change-functions, as they are trivial: the actions are deterministic and 'purely epistemic' (no change of facts). So $k \rightarrow k$, $l \rightarrow l$ and $k_0 = l_0 = \emptyset$.



Fig. 1

IV.2. Secure group announcements with no suspicion

Suppose a and b get together, without c suspecting this (or, alternatively, suppose a and b have common access to a secret, reliable and secure communication channel). Agent a makes a sincere announcement φ at this gathering (or sends a sincere message over this channel). Here, 'sincere' means that a actually believes φ to be true, and we actually assume more, namely that a and b trust each other. As mentioned, c does not suspect that this is happening: he trusts a and b and does not even consider the possibility of such a secret communication. (Or, alternatively, one can say that the act of communication is done in such a misleading way, that it appears to c as if nothing happened, and that nothing could happen.) This action can be described as $\alpha = (\Box_a \varphi)^{*[a, b]}$ and can be represented by the action model shown in Figure 2.

Here, $K = \{k, l\}$ as before, $k \rightarrow_a k, k \rightarrow_b k, k \rightarrow_c l, l \rightarrow_a l, l \rightarrow_b l, l \rightarrow_c l$, $pre_k = \Box_a \varphi$ (since the announcement is 'sincere', so the presupposition is




Fig 2

that a believes φ), $pre_l = true$ (the universally true condition). As before, $k_0 = l_0 = \emptyset$.

IV.3. *Message-passing over unreliable channels (but still no suspicion)*

As before, agent a sends a message to agent b , without c suspecting that this is happening. The message is again 'sincere' and a and b trust each other. The communication channel is secure, but not completely reliable: messages can be lost before reaching b . But in fact, the message *is* received by b . The picture is more complex this time (Figure 3).

The reason is that a cannot distinguish between the real action-token k and the alternative action t in which b does not receive the message. If t were the 'real' action, then b 's view of the action would be the same as c 's: i.e. they would be both misled into thinking that 'nothing happened' (i.e. they will believe the 'trivial' action l is the one that is happening).




Fig. 3

IV.4. Reliable, secure group announcements with a suspicious outsider

As in example 2, but now c is suspicious: he doesn't trust a and b so much, so he suspects this group announcement might be happening. He does not necessarily believe it is happening, but he doesn't exclude such a possibility. On the other hand, a and b know this, and moreover they have common knowledge of this suspicious character of c . This is shown in Figure 4.

Another example (from van Ditmarsch, 2000) is the action 'show your card': agent a shows her red card to b , in the presence of c ; c witnesses the act of showing the card, but does not actually see the card. However, c knows that a has either a red or black card (either because of prior information, or because he learns it during the act of showing: maybe in the same time a publicly announces that she has either a red or a black card). This is shown in Figure 5.

IV.5. Group announcements with a (secure) wiretap

As in the last example, but now c is not only suspicious, but extremely curious: he actually wiretaps the conversation between a and b (or violates their mail etc.). So c knows about the announcement, while a and b don't suspect this: they just do not consider wiretapping as a real possibility; but they still know that c is suspicious, so they do suspect that c suspects something. But (owing to his wiretapping) c knows all this (including their suspicion about his suspicion). This is shown in Figure 6.

IV.6. Fact-changing: exchange of cards

All the previous examples were 'purely epistemic', with no changing of facts. Suppose that it is common knowledge that there are only two cards




Fig. 4




Fig. 5




Fig. 6

left in a game, a red one and a black one, and that a has one of them and b has the other. Suppose that, in fact, a has the red card before this action but that a and b publicly exchange their cards, in the presence of c (who sees the exchange, but not the cards). The picture is as shown in Figure 7, where the sets inside represent the content of the actions (the atomic facts whose values are flipped).

IV.7. *S5-actions and S4-actions*

How about the case of 'fully introspective actions' and of 'knowledge actions' (in which nobody is deceived, in addition to full introspection)? An *knowledge-action* (or *S5-action*) is one in which all the accessibility relations are equivalence relations. Similarly, a *belief-action* (or *S4-*

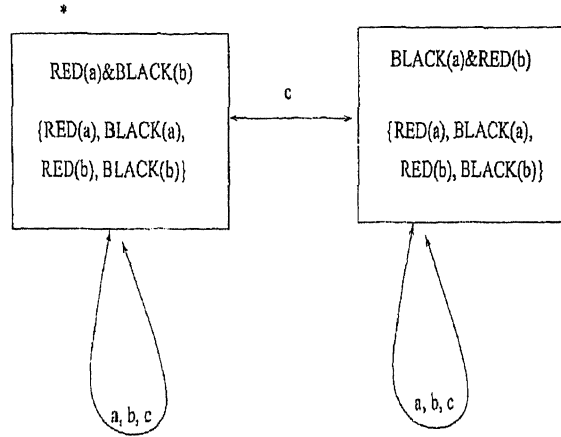


Fig. 7

action) is one in which all the accessibility relations are transitive and Euclidean.

IV.8. Non-deterministic actions

Let us change the previous example such that the same type of card-exchange takes place, but that we are not given any information concerning who has the red card. This is a non-deterministic action, which can be represented as the set (or the sum) of the two alternative actions present in the previous example. We don't know anymore which one is the 'top', the actual action. We can also represent this by explicitly drawing the choice transition (Figure 8).

IV.9. Example of 'product' update

Suppose we have three agents, *a*, *b* and *c*, and that there is some relevant fact *P*, known only to *a* (for instance *P* might be the fact that agent *a* has an Ace in a poker game); moreover, suppose it is public knowledge that *a* knows whether *P* or not (say, because the 'rules of the game' are such that everybody knows his/her own card). The initial epistemic state can be represented by the Kripke structure shown in Figure 9.

The possible states or worlds are represented by circles, the accessibility relations by arrows, and the 'actual' world is the one in which *P* holds (while *b* and *c* cannot distinguish between it and the other 'possible world', in which *P* fails). Suppose now that, without *c* knowing or suspecting anything, *a* tells *b* that *P* holds; moreover, *a* and *b* are mutually trusting each other, so that it is common knowledge among

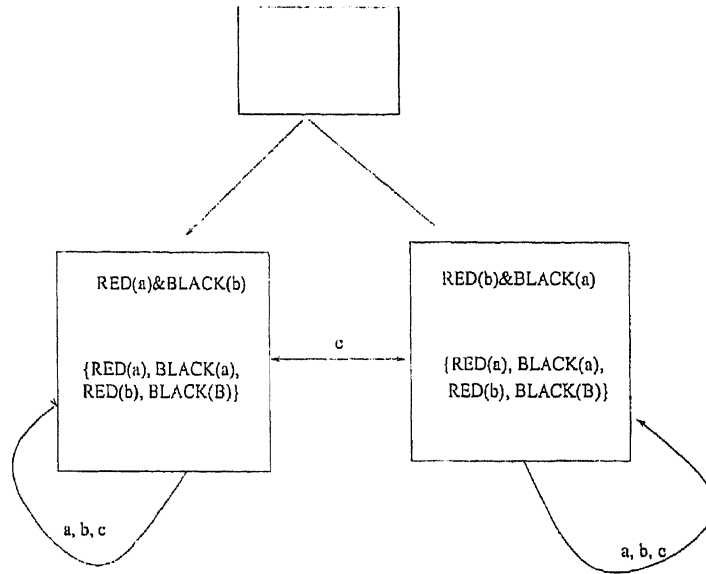


Fig. 8




Fig. 9

them that what a says is actually true. This is a ‘secure group announcement with no suspicion’, of the kind described in example 2 above: the picture of this action is precisely the one in Example 2, if we take the announcement φ to be the atomic sentence P itself. Using the above definition for update, one can easily compute the output epistemic state (Figure 10), where the ‘actual’ world is the one on top, in which both a and b know in common that P , while c considers as possible only the ‘old worlds’ (the states on the bottom, identical copies of the old states, in which b didn’t know whether P holds or not). As expected, this




Fig. 10

action is misleading c , inducing him to have false beliefs about the world (as shown by the non-reflexive arrows).

IV.10. Bisimulation of epistemic actions

The standard notion of observational equivalence for epistemic states is *bisimilarity*. We remind the definition of this important concept, by introducing it in a slightly nonstandard manner, via the notion of powerset-lifting of a relation: this is a way to naturally ‘lift’ any binary relation on objects to a relation on *sets of objects*.⁵

Definition: Given a binary relation $R \subseteq C \times C$ on a class C , the *powerset-lifting of R* is a binary relation $\check{R} \subseteq \mathcal{P}(C) \times \mathcal{P}(C)$ between *subsets of C* , defined by:

$$A \check{R} B \text{ iff } \forall a \in A \exists b \in B a R b \quad \text{and} \quad \forall b \in B \exists a \in A a R b.$$

An *epistemic bisimulation* is a binary relation $R \subseteq Mod \times Mod$ between epistemic states s.t.:

$$\text{if } s R t \text{ then } s_0 = t_0 \quad \text{and} \quad s_0 \check{R} t_a \text{ (for all } a \in Ag).$$

It is easy to see that this definition is equivalent to the standard ‘back-and-forth’ conditions defining bisimulation between pointed Kripke

⁵ More precisely, the importance of this notion is related to the Extensionality Axiom in set theory: if we take R as our notion of *identity for objects* (i.e. we identify objects modulo R) then the Extensionality Axiom implies that the resulting notion of *identity between sets of objects* is given by \check{R} .

models. The powerset – lifting \check{R} of any epistemic bisimulation R is called a *set-bisimulation*.⁶ Two states (or sets of states) are said to be *bisimilar*, written as $s \sim t$, if there exists some bisimulation (or set-bisimulation) relating them. We write $s \sim t$ for bisimilar states, and $S \sim T$ for sets of states.

One can easily define an analogue relation of observational equivalence for epistemic actions, by adding identity of preconditions as an extra-requirement for bisimulation.

Definition: A *simple-action bisimulation* is a binary relation $R \subseteq Act^0 \times Act^0$ between simple epistemic actions s.t.:

$$\text{if } \alpha R \rho \text{ then } pre_\alpha = pre_\rho, \sigma_\alpha = \sigma_\rho \quad \text{and} \quad \sigma_a \check{R} \rho_a \text{ (for all } a \in Ag).$$

The powerset-lifting $\check{R} \subseteq \mathcal{P}(Act^0) \times \mathcal{P}(Act^0)$ of any simple-action bisimulation R is called an *epistemic action bisimulation*. (Observe that \check{R} is indeed a binary relation between general epistemic actions.) Two (simple) actions are said to be *bisimilar*, if they are related by some (simple) action bisimulation. We write $\sigma \sim \rho$ for bisimilar simple actions, and $\alpha \sim \beta$ for bisimilar (general) epistemic actions.

I mention here, without proof, the following results. Proofs of older versions of these results (for logical systems that are similar to the present one, but lacking nondeterminism and fact-changing actions) can be found in Baltag (1999) and Baltag *et al.* (1998, 1999), and the proofs for the present version are contained in Baltag (2000). The proof of completeness uses a *terminating rewriting system* for sentences and action-expressions and a *filtration argument* similar to the one by Kozen and Parikh to prove the completeness of PDL.

Proposition 2 (Bisimilar actions applied to bisimilar states yield bisimilar outputs): If $s \sim t$ and $\alpha \sim \beta$ then $s.\alpha \sim t.\beta$. In words: given two bisimilar actions acting on two bisimilar input-states, every possible output of the first action applied to the first input-state is bisimilar to some output of the second action applied to the second input-state (and vice-versa).

Proposition 3: S5-actions applied to S5-states yield S5-outputs. Similar, S4-actions applied to S4-states yield S4-outputs.

⁶ Observe that \check{R} is a binary relation between *sets of states*; as mentioned in the previous footnote, if we take a bisimulation R to be our notion of equivalence for epistemic states, then we should take the corresponding set-bisimulation \check{R} as our notion of equivalence between sets of states. We can thus read the above definition of bisimulation as imposing a minimal requirement for the relation R to be acceptable as a good notion of observational equivalence between epistemic states: if we identify two states (via R) then we should identify their contents (via $=$) and their appearance-sets (via \check{R}).

Proposition 4 (complete equational system): There exists a complete equational calculus of epistemic actions; i.e. there exists an equational system, containing equations between terms involving the action operations of our syntax (test, flip, sum, sequential composition, suspicion α^a and common learning α^{*A}), system which is sound and complete with respect to action bisimilarity.

Proposition 5 (completeness and decidability): The axiomatic proof system presented below provides a sound and complete axiomatization for this logic. The proof method implies also that the logic is decidable.

TABLE I
The proof system for the logic of epistemic actions

<i>Basic axioms</i>	
All propositional validities	
($[\alpha]$ -normality)	$\vdash [\alpha](\varphi \rightarrow \psi) \rightarrow ([\alpha]\varphi \rightarrow [\alpha]\psi)$
(\Box_a -normality)	$\vdash \Box_a(\varphi \rightarrow \psi) \rightarrow (\Box_a\varphi \rightarrow \Box_a\psi)$
(\Box_A^* -normality)	$\vdash \Box_A^*(\varphi \rightarrow \psi) \rightarrow (\Box_A^*\varphi \rightarrow \Box_A^*\psi)$
<i>Composition axiom</i>	$\vdash [\alpha \cdot \beta]\varphi \leftrightarrow [\alpha][\beta]\varphi$
<i>Nondeterministic choice</i>	$\vdash [\alpha + \beta]\varphi \leftrightarrow [\alpha]\varphi \wedge [\beta]\varphi$
<i>Mix axiom</i>	$\vdash \Box_A^*\varphi \rightarrow \varphi \wedge \bigwedge_{a \in A} \Box_a \Box_A^*\varphi$
<i>Simple-action axioms</i>	
Let σ be a simple action-expression.	
(Change of facts)	
If $P \notin \sigma_0$ then	$\vdash [\sigma]P \leftrightarrow (pre_\sigma \rightarrow P)$
If $P \in \sigma_0$ then	$\vdash [\sigma]P \leftrightarrow (pre_\sigma \rightarrow \neg P)$
(Partial functionality)	$\vdash [\sigma]\neg\chi \leftrightarrow (pre_\sigma \rightarrow \neg[\sigma]\chi)$
(Action-knowledge)	$\vdash [\sigma]\Box_a\varphi \leftrightarrow (pre_\sigma \rightarrow \Box_a[\sigma_a]\varphi)$
<i>Modal rules</i>	
(<i>Modus ponens</i>)	From $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$, infer $\vdash \psi$
($[\alpha]$ -necessitation)	From $\vdash \varphi$, infer $\vdash [\alpha]\varphi$
(\Box_a -necessitation)	From $\vdash \varphi$, infer $\vdash \Box_a\varphi$
(\Box_A^* -necessitation)	From $\vdash \varphi$, infer $\vdash \Box_A^*\varphi$
<i>Action-common-knowledge rule</i>	
Let σ be a simple action-expression, φ be a sentence and A be a set of agents. Consider some sentences χ_ρ for all $\rho \in \sigma _A$ (i.e. all ρ such that $\sigma \rightarrow_A^* \rho$, including σ itself). Assume that:	
1. $\vdash \chi_\rho \rightarrow [\rho]\varphi$.	
2. If $a \in A$ and $\rho' \in \rho_a$, then $\vdash (\chi_\rho \wedge pre_\rho) \rightarrow \Box_a \chi_{\rho'}$.	
∴ From these assumptions, infer $\vdash \chi_\sigma \rightarrow [\sigma]\Box_A^*\varphi$.	

Proposition 6 (completeness for S5 (and S4) actions and models): If we restrict our class of models to knowledge models (i.e. S5-models) and the class of epistemic actions to knowledge (S5) actions, then we can obtain a sound and complete proof system for this class by adding to the system above the standard multi-agent S5 axioms. Similar remarks apply to S4-models and S4-actions.

Proposition 7 (normal form representation theorem): Every (finite) action is bisimilar to a (finite) sum of products of tests γ , change-actions (of the form $\text{change} = \prod_{P \in \mathcal{P}} \text{flip } P$) and suspicion-actions β^a . Moreover, this representation is unique (up to reordering and bisimulation). More precisely, if we put $\text{change}(\sigma) = \prod_{P \in \sigma_n} \text{flip } P$ then we have:

$$\alpha = \sum_{\sigma \in \alpha} \gamma(\text{pre}_\sigma) \cdot \text{change}(\sigma) \cdot \prod_{a \in Ag} (\sigma_a)^a.$$

V. EPISTEMIC ACTIONS AND INFORMATION FLOW IN GAMES

V.1. Dialogue games: an analysis of a muddy children game

As an application of the method, I give an analysis of a ‘modified muddy children’ puzzle, similar to the one given by Gerbrandy to the classical version of this puzzle. There are four children a, b, c, d , the first three are the muddy ones. Each can see the others but not himself. The father comes and says publicly: ‘At least one of you is muddy’. Then they play a game, in rounds. In each round they all simultaneously announce publicly one of the following: ‘I know I am muddy’, ‘I know I am not muddy’, ‘I don’t know (whether I am muddy or not)’. After many rounds (say four for convenience), the game stops. The ones who gave a correct ‘definite’ answer (‘muddy’ or ‘not’) win (say 10 points), the ones who gave a wrong answer lose (–10 points), the ones who still don’t know finish with 0 points.

In the classical puzzle, it can be proved that in certain assumptions (namely, that it’s common knowledge that all children are sincere in their answers, that they are ‘good logicians’ and that they do not ‘take guesses’, but they answer only they know it) then all the dirty children win in three rounds and the others win in the fourth run. But one of the not so easily observable assumption is the absence of secret communications. Even if the children are sincere and do not ‘cheat’ by lying or guessing, there are some more subtle forms of cheating. Let’s suppose for instance that, after the first round (but before the third), children a and b (very good friends, trusting and helping in each other even at the price of

cheating, because ... a friend in need is a friend indeed ...) decide to 'cheat' by sending each other secret signals to communicate the message: 'You are dirty'. Naturally, in the second round, they both answer 'Yes, I know I am dirty' and win. Child *c* is also a very trustful person, so trustful that she cannot imagine that such a dirty and secret communication between her dirty colleagues could have taken place. So, in the third round, she is confused: thinking that *a* and *b* used only their reasoning abilities to answer, she concludes that (the only way for this to have happened is if) *a* and *b* were the only dirty ones. So she hurries to answer 'I know I am not muddy' and she loses! The fourth child *d* is the only 'clean' one, and he has two possibilities: he either 'gets suspicious', i.e. starts entertaining the possibility (which soon becomes a certainty, after *c*'s wrong answer) that *a* and *b* cheated; or he could still go on and think this is impossible. In the first case, his action of suspicion will help him to win in the end: after the third round, he gets convinced that *a* and *b* cheat, that *c* is deluded and that himself (*d*) is clean, which he actually will say in the fourth run, winning! But in the second, he will 'go crazy': he will never understand what happened: after the third run, his set of beliefs is not only false, but is actually inconsistent!

Let, for each agent $i \in Ag = \{a, b, c, d\}$, $D_i, Win_i, Lose_i$ be some atomic sentences, meaning 'i is dirty', 'i wins 10 points', 'i loses 10 points'. We make the following abbreviations (some of which are inspired from Gerbrandy's analysis):

$$O_i = \neg Win_i \wedge \neg Lose_i \text{ ('i didn't lose or win yet')}$$

$$Win_I = \bigwedge_{i \in I} Win_i \text{ ('all agents in } I \text{ are winning')}$$

$$Lose_I = \bigwedge_{i \in I} Lose_i \text{ ('all agents in } I \text{ are losing')}$$

$$D_I = \bigwedge_{i \in I} D_i \wedge \bigwedge_{i \notin I} \neg D_i \text{ ('I is the set of all dirty kids')}$$

$$\Box_i^+ \varphi = \Box_i \Box_i^* \varphi \text{ ('i reflexively and introspectively believes } \varphi \text{')}$$

$$Yes_i = \Box_i^+ D_i \text{ ('i believes he is dirty')}$$

$$No_i = \Box_i^+ \neg D_i \text{ ('i believes he is not dirty')}$$

$$Yes_I = \bigwedge_{i \in I} Yes_i$$

$$No_I = \bigwedge_{i \in I} No_i$$

$$?_i = \neg Yes_i \wedge \neg No_i \text{ ('}i \text{ doesn't know if he's dirty or not)}$$

$$Right_i = (D_i \wedge Yes_i) \vee (\neg D_i \wedge No_i) \text{ ('}i \text{ is right in his/her belief')}$$

$$Wrong_i = (D_i \wedge No_i) \vee (\neg D_i \wedge Yes_i) \text{ ('}i \text{ is wrong in his/her belief')}$$

$$father = D_a \vee D_b \vee D_c \vee D_d \text{ ('one of you is dirty')}$$

$$vision = \bigwedge_{i \neq j \in Ag} \Box_{Ag}^*(D_i \rightarrow \Box_j^* D_i)$$

$$Ans_i = \{Yes_i, No_i, ?_i\}$$

$$Ans = \left\{ \bigwedge_{i \in Ag} \chi_i : \chi_i \in Ans_i \text{ for all } i \in Ag \right\}.$$

'Vision' says it is common knowledge that everybody sees the others (and so knows, with full introspection, whether or not the others are dirty). Ans_i is the set of i 's possible answers (in one round of questioning); Ans is the set of possible 'global answers' to father's question (tuples of answers of each agent in one round).

For actions we introduce the following abbreviations:

$$cheat_{i,j} = (D_i? \cdot D_j?)^{*(i,j)} \text{ ('agents } i \text{ and } j \text{ cheat')}$$

$$suspicion_d = \left(\sum_{i,j \neq d} cheat_{i,j} + skip \right)^{*d}$$

('d suspects cheating is happening')

$$WinRule = \prod_{i \in Ag} [\text{if } (0_i \wedge Right_i) \text{ do } (Win_i)! \text{ else (if } (0_i \wedge Wrong_i) \text{ do}$$

$(Lose_i)! \text{ else skip}]$

$$PA = \prod_{\chi \in Ans} (? \chi \cdot WinRule)^{*Ag} \text{ ('public answering')}.$$

$Cheat_{i,j}$ is the cheating communication action, by which i and j secretly exchange messages concerning their own dirtiness; $suspicion_d$ is the action by which d starts to suspect that cheating (by any of the other

pairs of players) might be happening (but he is not sure); *WinRule* is the winning rule of the game: the player whose answer is right wins, the player who is wrong loses, and the undecided ones can continue to 'play'. *PA* is the main move of this game: the public action of all the players publicly answering father's question in the same round; we made the winning rule part of it, since we think of each round as being immediately followed by the exclusion from the game of all the 'decided' players (who win or lose according to the public winning rule); only the undecided can continue to play.

Assume $Ag = \{a, b, c, d\}$. Then the classical muddy children puzzle is explained by the following theorem of our logic:

$$\vdash D_{a,b,c} \wedge \Box_{Ag}^* vision \rightarrow [father^{*Ag}][PA][PA][PA](Win_{a,b,c} \wedge [PA]Win_d).$$

This theorem that, if the initial state is such that a, b and c are only dirty players and they all can see each other, than after father's public announcement, followed by two rounds of public answering, all the dirty players 'win' (by correctly answering in the third round), and then the clean player also wins in the fourth round. This shows that, if no cheating occurs, then indeed after four rounds everybody 'wins' (i.e. finds out if he/she's dirty or not). In the case that cheating by a and b does take place, the following theorem explains the (false, but rational) conclusion of c (that she is clean), conclusion which leads to her 'defeat':

$$\vdash D_{a,b,c} \wedge \Box_{Ag}^* vision \rightarrow [father^{*Ag}][PA][cheat_{a,b}][PA](Win_{a,b} \wedge [PA]Lose_c).$$

If d starts to suspect what is happening, then he will truthfully conclude that he is clean, winning in the fourth round:

$$\vdash D_{a,b,c} \wedge \Box_{Ag}^* vision \rightarrow [father^{*Ag}][PA][cheat_{a,b}][suspicion_d][PA][PA][PA]Win_d.$$

In fact, it is even more natural to model d 's 'suspicious' behaviour by assuming that he suspects cheating at every stage of the game (as he cannot know for sure when it'll happen). If, instead of *PA*, we take the main move of the game to be the action of 'public answering with d -suspicion', defined by $PA_{-d} = suspicion_d \cdot PA$, then we can similarly prove that:

$$\vdash D_{a,b,c} \wedge \Box_{Ag}^* vision \rightarrow [father^{*Ag}][PA_{-d}][cheat_{a,b}][PA_{-d}][PA_{-d}][PA_{-d}]Win_d.$$

So our logic can faithfully represent the player's reasoning in this example.

V.2. Rule-based games

As usually defined in game theory, a game is just a tree, with nodes corresponding to the possible successive states of the game and labelled arrows between them, corresponding to the possible moves. Each nonterminal node is labelled with the name of some player $a \in Ag$, who is supposed to move at that node. For each player a , it is given an 'information partition' of the set of all states, which is essentially the same as having equivalence relations \sim_a on this set for each player.

There are some in-built problems with this model. First, it assumes that no player ever 'cheats' and no player is ever 'deceived': the epistemic relations being equivalence relations, there is no way to model false beliefs inside this model. There are known proposals for ways to deal with this issue, and with the more general issue of the interplay between beliefs, actions, strategies, beliefs about strategies etc. in a game; but all these proposals go way beyond this simple tree-model of a game.

Secondly, there is a computational problem, related to the enormous size of the set of all possible states in most natural games. Since the model has to contain at once *all possible future states of the game*, the size of such a model will typically be huge, and by default the associated logics will become computationally intractable, if not plainly undecidable. This 'brute force' approach to modelling games can still be useful for many purposes, but it provides no way to express softer, more subtle, ways to play a game, based on reasoning about the rules of the game, on local (and temporally circumscribed) reasoning about strategies and mutual beliefs of the player *as they appear at the moment of playing*. There is no principled way to calculate in advance the next possible states and the next beliefs of the players, except for just looking at the future nodes of the tree; this is like trying to anticipate the future of a game, but not by using reasoning based on rules, but by just ... playing (and in fact, playing all the possible moves).

I sketch here another proposal for modelling a notion of 'rule-based games'. Owing to the lack of a Kleene star (iteration operator) from our logic (see the last section for the reasons we chose not to have one), we have to restrict ourselves to *games of bounded length*. Our syntax will essentially be the one of our epistemic action logic, for some finite set Ag of 'players'; we only have to add some special atomic sentences $\{P_a^k\}_{a \in Ag, 1 < k < m_a}$ (for some natural numbers $\{m_a\}_{a \in Ag}$), which express winning conditions: P_a^k means 'player a finishes with a payoff of at least k ' (or, say, 'player a wins at least k dollars').

In this section I will restrict our class of epistemic states to *S4*-states. (Correspondingly, we will only consider *S4*-actions.) Hence, we assume as additional axioms the standard multimodal *S4*-axioms.

Definition: A *game* will be a tuple $G = (\{M_a\}_{a \in Ag}, \{P_a^k\}_{1 \leq k \leq m_a \leq N})$ consisting of the following: for each player a , a finite set $M_a \subseteq Act$ of epistemic *S4*-actions, called *the possible moves of player a*, and satisfying $\alpha_a \sim \alpha$ for every $\alpha \in M_a$ (i.e. 'players know their moves'⁷); for each player a and number $1 \leq k \leq m_a$, some set of epistemic states P_a^k , understood as the set of *all states in which player a wins a payoff $\geq k$* ; and an upper bound N for the number of rounds of the game. Usually, these components will be assumed to satisfy a list of extra-conditions, enumerated below.

The conjunction of the requirements that all moves $\alpha \in M_a$ are *S4*-actions and that $\alpha_a \sim \alpha$ implies the following: every move $\alpha \in M_a$ is a set of a -reflexive simple actions (i.e. s.t. $\sigma \rightarrow_a \sigma$) and moreover this set is closed under a -accessibility (i.e. if $\sigma \in \alpha$ and $\sigma \rightarrow_a \sigma'$ then $\sigma' \in \alpha$); in other words, each a -move is a union of \rightarrow_a -equivalence classes. This, in its turn, implies the validity of the following two schemas:

$$\Box_a[\alpha]\varphi \rightarrow [\alpha]\Box_a\varphi$$

$$\langle \sigma \rangle \Box_a\varphi \rightarrow \Box_a[\sigma]\varphi$$

for every a -move $\alpha \in M_a$ and every simple component $\sigma \in \alpha$. The first schema expresses a 'perfect-recollection'-type postulate: if a player knows (or believes) that after he'll play some move α sentence φ will become true, then after he really plays move α he will know (believe) φ to be true. In other words, a player cannot be 'totally surprised' by his own move: performing it will not contradict any of his prior beliefs. The second schema says that, if after a player's chosen move is actually executed (and realized) as a concrete simple action σ the player comes to know (believe) sentence φ , then he must have known (believed) already before the move that such a realization of his move as σ would make φ to be true. In other words, a player cannot be 'deceived' by any (simple-action-)realization of his own move; all the simple actions σ subsumed by his move are nondeceiving: $\sigma \rightarrow_a \sigma$. As a result of these two schemas, the only way a player can 'learn' from his own move is by learning which of the simple actions subsumed by his move is actually executed; but no

⁷This means the appearance of their own moves to themselves is correct. But note that this does *not* imply that they know the concrete realization, or resolution, of their own actions in each context: the moves might be nondeterministic, without the player being able to distinguish between its different simple resolutions; he only knows the *type* of his own action.

simple component-action in itself is adding any new knowledge to its author.⁸

We put $\text{Play}_\alpha = \|\text{pre}_\alpha\|$ (the class of all states in which move α is 'playable'), $\text{Play}_a = \cup_{\alpha \in M_a} \text{Play}_\alpha$ (the class of all states in which player a is to play first), $\text{Play} = \cup_{a \in A} \text{Play}_a$ (the class of all 'playable' states), $M = M_G = \cup_{a \in A} M_a$ (the set of all possible moves), $P_a = \cup_k P_a^k$ (the set of all winning states for player a), $T = \cap_{a \in A} P_a$ (the set of all possible 'outcomes', i.e. terminal states, in which everybody wins something). For a set of outcomes $\Theta \subset T$, we define the *minimal payoff* Θ_a of player a in Θ by $\Theta_a =$ the largest k s.t. $\Theta \subset P_a^k$.

We also define the following game-related action-terms: the *game-action* $\gamma = \gamma_G = \sum_{\alpha \in M} \alpha$, the *player-action* $\gamma(a) = \sum_{\alpha \in M_a} \alpha$ and the *anti-player-action* $\gamma(-a) = \sum_{b \neq a} \gamma(b)$. Notice that, for each player a , we clearly have $\gamma = \gamma(a) + \gamma(-a)$. We can also introduce a sentence Play_a , saying that 'player a is the one to play', and a sentence P_a , saying that 'player a wins some payoff':

$$\text{Play}_a = \text{pre}_{\gamma(a)} = \bigvee_{\alpha \in M} \text{pre}_\alpha$$

$$P_a = \bigvee_{k=1}^{m_a} P_a^k.$$

(As a consequence, we have $\|\text{Play}_a\| = \text{Play}_a$ and $\|P_a\| = P_a$.)

As announced, we list some conditions for games:

1. $\text{Play}_a \cap \text{Play}_b = \emptyset$ for $a \neq b$ ('no two players can move in the same time');
2. $s \in P_a^k$ iff $s_b \subseteq P_a^k$ ('every player knows every other player's winning payoff');
3. if $s \in P_a^k$ then $s_\alpha \subseteq P_a^k$, for every move $\alpha \in M$ ('once you win, you win: game's over for you');
4. $s \in P_a$ iff $s_\gamma = \emptyset$ ('you win a payoff iff you cannot make a new move');

⁸ Example: the move 'pick a card' (at random, from a pack, without looking at the cards in the pack) in a card-game. Let's assume that this move includes turning face up the card you have picked and looking at it. The simple actions subsumed by this move are all the actions of type 'pick card X ', where X is any particular card (e.g. a Queen of Swords). The player *cannot choose* any of these simple actions (he cannot choose a particular card of his liking), but only their type (he can choose whether or not to perform the nondeterministic move of picking a card at random). This justifies our terminology: what we call 'moves' in a game can only be such 'types' of actions, which are susceptible to be 'chosen' by the agents. A player can indeed learn from his own such move (e.g. can learn there was a Queen of Swords in the pack), but cannot learn from any of the deterministic simple actions subsumed by the move: he already knows that, if he'll choose a Queen of Swords, then there must have been such a queen in the pack; what he learns, in fact, is to distinguish which of these simple-action-components of his move is actually realized.

5. every $\alpha \in M_a$ has as precondition a sentence (equivalent to one) of the form $\Box_a^* \phi_a$, for some sentence ϕ_a ('a player knows when he can make a given move');
6. $P_a^j \subseteq P_a^k$ for $j \geq k$ ('if you win an amount bigger than k , then you win at least k ').

Of course, in the formal semantics we will interpret the special atoms using the payoff sets: $\|P_a^k\| =: P_a^k$. It is easy to see that the conditions listed above in the definition of a game correspond to special *modal sentences* in our language.

Proposition 8: A game satisfies any of the above six conditions iff the corresponding condition in the list below is valid (i.e. true in all epistemic states):

- 1'. $Play_a \rightarrow \neg Play_b$, for $a \neq b$
- 2'. $P_a^k \leftrightarrow \Box_b P_a^k$
- 3'. $P_a^k \rightarrow [\alpha] P_a^k$ for every move $a \in M$
- 4'. $P_a \leftrightarrow \neg Play_a$
- 5'. $pre_\alpha \leftrightarrow \Box_a pre_\alpha$
- 6'. $P_a^j \rightarrow P_a^k$, for $j \geq k$.

One of the natural requirements which are usually imposed on game-trees for games with imperfect information is that *at states that are indistinguishable for player a , the sets of a 's available moves are the same*. One can easily see that this is a consequence of our conditions, more precisely of condition 5. The requirement is usually stated in the S5 context of standard game theory, but in our more general S4 context we can define 'indistinguishability for player a ' as 'bisimilarity of appearances for player a ', by putting: $s \sim_a t$ iff $s_a \sim t_a$ (where \sim is set-bisimilarity). Then we can prove that condition 5 implies that:

for every a -move $\alpha \in M_a$, if $s \sim_a t$ and $s \in Play_\alpha$ then $t \in Play_\alpha$.

V.3. Game-playing situations

The set of all possible initial states of the game G is $Start_G = \{s \in Mod : s.\gamma_G^N \subseteq P\}$. Again, it is to see that this set is definable by the modal logic sentence $\bigwedge_{a \in Ag} [\gamma_G^N](\bigvee_k P_a^k)$, which says that after at most N moves one has to reach a terminal state. A *game-playing situation* is a pair (G, s) of a game and an initial epistemic state $s \in Start_G$. It is clear that, given a game-playing situation (G, s) , we can use our product-update operation to compute the set $(G, s)^+$ of all possible 'next' game-playing situations:

$$(G, s)^+ =: \{(G_{-1}, s') : s' \in \gamma_G.s\}$$

where G_{-1} is the game that looks just like G , except that the upper bound for the number of rounds is $N - 1$. Applying repeatedly this operation, we can define the set $(G, s)^{(n)}$ of all possible n -step game-playing situations by: $(G, s)^{(0)} = (G, s)$ and $(G, s)^{(n+1)} = ((G, s)^{(n)})^+$. Finally, by applying this N times, we can compute the set of all possible outcomes of the game G played on the initial state s :

$$\mathcal{O}(G, s) = (G, s)^{(N)}.$$

V.4. Strategies and strategy profiles

The set of all *partial (rule-based) strategies of player a* (strategies denoted by $\sigma(a), \sigma'(a), \dots$), is defined inductively by:

$$\sigma(a), \sigma'(a) ::= \alpha \in M_a \quad | \quad \text{if } \Box_a \varphi \text{ do } \sigma(a) \text{ else } \sigma'(a).$$

So a -strategies are *conditional actions*, with conditions being given by *knowledge tests* for agent a and actions being given by a 's *possible moves*.⁹ This concept is formally different from the standard semantic concept of strategy in game theory, but it's related. In game theory, strategies are functions from equivalence classes (modulo the indistinguishability relation \rightarrow_a , assumed to be $S5$) into possible moves. Our definition refers to strategies as action-terms of a specific kind. We think our definition is more natural in a rule-based approach to games. Clearly, one can recover the associate game-theoretic strategy from these strategic terms: just take the interpretations (truth-sets) of the conditions (the knowledge tests) in our conditional expressions for strategies; in the $S5$ -case these will split up into \rightarrow_a -equivalence classes (since the interpretation of a knowledge test for a is closed under \rightarrow_a -equivalence). So we can just define the strategy, by mapping an equivalence class into the move which occurs in our term-strategy as being conditioned by a knowledge test which includes the given equivalence class. For finite games, we can also go the other way around: we can recover one definition from the other (since finite epistemic structures are characterizable up to bisimilarity in epistemic logic with common-knowledge operators).

Total strategies are those strategies σ which exhaust all the possibilities. Logically, they can be characterized by the validity of the sentence $\text{pre}_{\sigma(a)} \leftrightarrow \text{Play}_a$.

For a set $A \subseteq Ag$ of players, an A -*profile of strategies* is an indexed tuple $\vec{\sigma} = \{\sigma(a)\}_{a \in A}$, s.t. each $\sigma(a)$ is a total strategy of player a . Then we write $\vec{\sigma}_a = \sigma(a)$. A *full profile* is an Ag -profile of strategies. Given an A -profile and some player $a \in A$, we can define an $A \setminus \{a\}$ -profile $\vec{\sigma}_{-a} = \{\sigma_b : b \neq a\}$. Similarly to the game-action and the player-action, for a

⁹ This is very close to the notion of *knowledge-based protocols*, introduced by Fagin *et al.* (1995).

given profile $\vec{\sigma} = \{\sigma_a\}_{a \in A}$ we can define a *profile-action*:

$$\gamma(\vec{\sigma}) =: \sum_{a \in A} ?(Play_a) \cdot \sigma_a.$$

A *profile-playing situation* is a pair $(\vec{\sigma}, s)$ of a full profile of strategies for some game G and an initial state $s \in Start_G$. Similarly to the operations $(G, s)^+$, $(G, s)^{(n)}$ and $\mathcal{O}(G, s)$ defined above for game-playing situations, we can define the following: the set $(\vec{\sigma}, s)^+ =: \{(\vec{\sigma}, s') : s' \in \gamma(\vec{\sigma})\}$ of all the possible next profile-playing situations; the set $(\vec{\sigma}, s)^{(n)}$ of all possible n -step profile-playing situations, set defined by $(\vec{\sigma}, s)^0 =: (\vec{\sigma}, s)$ and $(\vec{\sigma}, s)^{(n+1)} =: ((\vec{\sigma}, s)^{(n)})^+$; and finally the set of all possible outcomes of playing the strategy profile $\vec{\sigma}$ on the initial state s :

$$\mathcal{O}(\vec{\sigma}, s) =: (\vec{\sigma}, s)^{(N)}$$

where N is the length of the game.

V.5. Best possible answers and Nash equilibria in games with incomplete information

It is rather easy to write a dynamic logic formula (for instance, using our game-action and some similar concept for strategy-actions) which captures the concept of a strategy profile being a Nash equilibrium in games with perfect information. Moreover, one can do this while avoiding quantification over strategies. (This is important: in general, their number is huge. In fact, remember that we chose to not even introduce the semantical concept of a strategy, as it is usually done, so we cannot do this quantification: we would have to quantify over all the infinitely many equivalent action-expressions that describe the same strategy.)

But it is harder to do it for games with imperfect information, as the ones we are dealing with here. In fact, to do it, we have to restrict ourselves back to *S5*-games and states. So, for the rest of this section, we restrict our states (and actions) to *S5*-states (and *S5*-actions) (i.e. states s such that $s \rightarrow_a s$) and we postulate as an additional axiom the specific *S5* axiom $(\Box_a \varphi \rightarrow_a \varphi)$. An *S5* game-playing situation is a pair (G, s) s.t. G has only *S5*-moves and the starting state s is *S5*. To any such *S5* game situation we can associate in the obvious way a game-tree in the traditional sense, of a game with incomplete information.

Let us first recall, using our notations, the definitions for best possible answer to a strategy profile', 'Nash equilibrium' and 'subgame perfect equilibrium'. (In the following, we are using the notation \mathcal{O}_a introduced above for the minimal payoff of player a in the set of outcomes \mathcal{O} .)

Definition: Given a profile-playing situation $(\vec{\sigma}, s)$ (for a game G), a 's strategy $\vec{\sigma}_a$ is said to be a *best possible answer in the initial state s* to the adversary's $Ag \setminus \{a\}$ -profile of strategies $\vec{\sigma}_{-a}$ iff for any other a -strategy σ'_a we have:

$$O(\vec{\sigma}, s)_a \geq O(\vec{\sigma}_{-a}, \sigma'_a, s)_a.$$

The profile-playing situation $(\vec{\sigma}, s)$ is a *Nash equilibrium* if for every agent a , $\vec{\sigma}_a$ is a best possible answer in s to $\vec{\sigma}_{-a}$. Finally, the strategy profile $\vec{\sigma}$ is a *subgame-perfect equilibrium* if $(\vec{\sigma}, s)$ is a Nash equilibrium for every playable state s .

By unfolding this definition and translating it into our language, we can immediately give a first logical translation of the notion of 'best possible answer'; this translation will involve quantification over strategies, so it will in fact be a schema, not a formula in our language: in the initial state s , $\vec{\sigma}_a$ is a best possible answer to $\vec{\sigma}_{-a}$ iff for every other a -strategy σ'_a and every $1 \leq k \leq m_a$, we have that:

$$[\gamma(\vec{\sigma}_{-a}) + \gamma(\sigma'_a)]^N P_a^k \rightarrow [\gamma(\vec{\sigma})]^N P_a^k.$$

But we can improve on this translation, by using the epistemic modalities in the SS -context, to obtain an equivalent formula, in which we *do not have to quantify over strategies*. Indeed, let us put $\gamma(\vec{\sigma}, \alpha) = (\gamma(\vec{\sigma}_{-a} + \alpha))^a$, for $\alpha \in M_a$, and $\otimes_a^g \varphi = \bigvee_{\alpha \in M_a} \Box_a \langle \gamma(\vec{\sigma}, \alpha) \rangle \varphi$. Then the promised equivalent for being the best answer is the following:

$$\bigwedge_{1 \leq k \leq m_a} (\otimes_a^g)^N P_a^k \rightarrow \Box_a \langle \gamma(\vec{\sigma}) \rangle^N P_a^k$$

where $\langle \alpha \rangle$ is the dual of $[\alpha]$ (i.e. $\langle \alpha \rangle \varphi = \neg[\alpha] \neg \varphi$).

In words, this says that a strategy $\vec{\sigma}_a$ for player a is the best answer to a profile of adversary's strategies iff: whenever there exists a strategy σ' s.t. that a knows that σ' will ensure some minimal payoff then he/she also knows that already the given strategy σ' will ensure at least the same payoff. Unfolding this even more, to get closer to the actual formula, we obtain that: whenever there exists some a -move s.t. a knows that after using this move against the given profile of adversary's strategies, there will still exist some a -move s.t. a will know that after using against the same profile a will know etc. ... that after N such moves he can get some minimal payoff, then a can also be sure that the given strategy $\vec{\sigma}_a$ will ensure at least the same payoff.

This is a way of using iterated knowledge to express properties of (knowledge-based)-strategies: the quantifier over a -strategies is replaced by N iterations of the box \Box_a . Now, by using this expression for 'being a best possible answer', we can get the desired modal expression for Nash equilibrium.

Proposition 9: In the game-tree with incomplete information associated to an S5 game-playing situation, a profile playing situation $(\vec{\sigma}, s)$ is a Nash equilibrium iff the following modal sentence is true at the initial state s :

$$\bigwedge_{a \in Ag} \bigwedge_{k=1}^{m_a} (\Box_a \langle \gamma(\vec{\sigma}, a) \rangle)^N P_a^k \rightarrow \Box_a \langle \gamma(\vec{\sigma}) \rangle^N P_a^k$$

Similarly, a strategy profile $\vec{\sigma}$ is a subgame-perfect equilibrium iff the above modal sentence is valid (i.e. true at all playable states).

Remark: that in this proposition, the epistemic operator \Box_a (to be read now as ‘knowledge’, since it is truthful: we are back in S5) plays an important role: we cannot drop it and talk simply in terms of strategies and dynamic logic. There might indeed exist a way to play a ’s moves in the game (i.e. $\gamma(a)$) s.t., when the others still play their equilibrium strategies, a could achieve a higher payoff than playing his equilibrium strategy. But this can only happen ‘by accident’, as agent a cannot plan on it: there is no way to systematically ensure this will happen, i.e. there is no way for a to *know in advance* that it will happen.

The more interesting cases are of course the more nonclassical, non-S5 ones. We are currently working towards a more general approach of this kind to ‘games with hidden moves’ (cheating, suspicion of cheating), based on our logic of epistemic actions.

VI. COMPARISON WITH OTHER WORK

The origins of the subject are in Fagin *et al.* (1995), where the authors analyse knowledge in distributed systems, using a mixture of epistemic logic $S5_m$ and temporal logic. The fundamental issues, examples and insights that gave rise to the logic discussed here come from the work of Fagin and colleagues. But their approach is rather different, being based mainly on temporal logic, instead of dynamic logic. Moreover, their approach runs into several problems. First, the resulting logic is *too strong*: in general, it is not decidable.¹⁰ Secondly, from a different conceptual perspective, their logic seems to be *not expressive enough*: there is no notion of *updating knowledge (information)*; one cannot talk about the change of information induced by specific actions, but only about what happens ‘next’ (or ‘always’ or ‘sometimes’ in the future), and this is *only* determined by the model. (In their setting, the semantics is given by ‘runs’, i.e. temporal sequences inside a huge Kripke structure, describing all possible evolutions of the system.) When they actually

¹⁰Only 32 of the 96 logics of knowledge and time analysed by Halpern and Vardi [XXXX] contain common-knowledge operators; out of these, all but 12 are undecidable.

analyse concrete examples (e.g. the Muddy Children Puzzle), they do *not* use their formal logic only, but also 'external' (semantic) reasoning about models; in effect, they simply 'update' their structures from the outside (in the meta-language) according to informal intuitions, but without any attempt to give a systematic treatment of this operation. Both technically and philosophically, the present approach is essentially different: our models are simpler and easier to handle, as we are trying, not just to keep them finite, but to keep them as small as possible. In effect, we do *not* incorporate all the possible runs into the system as they do (as well as the game-theoreticians). Instead, our epistemic state models contain only information about the present moment (and the agent's uncertainties about it). The rest of the information is stored somewhere else: part of it, in the action-models, which describe only the properties of actions and the uncertainties concerning them. Other pieces of information (e.g. about the available actions, or 'moves', the rules of the game, constraints about the future etc.) are contained in our concept of a rule-based game. Strategies, strategy profiles, beliefs about strategies etc. are separate actions; in on-going work we are trying to formalize what is called in game theory the 'epistemic type' of a player.

One of the seminal ideas of the present work comes from a paper of Gerbrandy and Groeneveld (1997). The idea was to combine Fagin-style epistemic logic with the work of Veltman (1996) on update semantics. The authors introduce special kinds of epistemic actions, namely *public announcements* ('group updates'). Their logic is strong enough to capture all the reasoning involved in the Muddy Children Puzzle. In his PhD dissertation, Gerbrandy improves and extends these ideas with a 'program-update' logic.

Our own work started from observing some odd (or at least not always desirable) features of Gerbrandy's and Groeneveld's public announcements. Namely, they have 'group-learning' actions of the form $\mathcal{L}_A\varphi$, with the intended meaning 'the agents in the group A learn in common that φ is true'. The problem is that, with their definition, agents that are *outside* the group A (the 'outsiders') do not in any way *suspect* that the group-announcement is happening. Of course, they wouldn't *know* it is happening (since they are not part of the 'inside' group), but (by the definition in Gerbrandy and Groeneveld, 1997) these outsiders are not even allowed to *consider the possibility* that such an announcement might be happening. As a result, they are totally 'mislead': after this action, in the resulting Kripke structure, the outsiders 'live in an ideal world' (i.e. they do not 'access' the actual world anymore). To put it differently, even if the initial model was a 'knowledge structure' (i.e. $S5_m$ -model), updating it with any announcement with at least two insiders and one outsider will result in a non-knowledge (non- $S5_m$, more specifically nonreflexive) structure: the outsider acquires false beliefs about the world. Such an interesting 'deceiving' situation is indeed possible in real

life and we would like to still have such an action in our logic; but we wouldn't want to impose that *every* group-announcement-with-outsiders be necessarily deceiving! Moreover, we would like to give the outsiders a better chance: not only that they could *suspect* something is going on, but on the basis of this suspicion they might *act*, attempting to confirm their suspicions. They could, for instance, wiretap or intercept the communications of the 'insiders'.

The work of H. P. van Ditmarsch, although related in content, did not influence the work reported here, as we discovered it later, and enjoyed some comments and communications with him on these issues. All the actions (for the game of Cluedo) introduced in his PhD dissertation are very special cases of our actions. Although he did not study the meta-theoretical properties of his logic, completeness and decidability of his logic follows trivially as a particular case of our work.

VII. CONCLUSIONS AND FUTURE PROJECTS

In our work, developed in Baltag *et al.* (1998, 1999) and Baltag (1999), the present paper and our on-going work in Baltag (2000), we further generalized in several directions the ideas arising from Gerbrandy and Groeneveld (1997). The main conceptual and technical novelty consisted in our product-semantics for update, in which we have endowed actions with their own, internal epistemic structure. In addition to its philosophical importance, this idea has clear technical advantages: it offers a simple, compact way to represent epistemic changes and to compute their effect; it has greatly simplified our prior work on completeness and decidability for various logics, some proposed by J. Gerbrandy and H. van Ditmarsch, some arising from our own work; in its 'syntactical' version, the idea of endowing actions with an epistemic 'appearance' was useful in formulating simple, intuitive axioms to describe the interplay between knowledge (belief) and change. The space does not permit us here to go into a discussion of the axioms, but we would like to stress the importance we attribute to our Action-Knowledge Axiom. We think it captures a new, important insight about the relation between prior knowledge (or belief), posterior knowledge and knowledge (or belief) of the action itself. The Action-Common-Knowledge Rule is an 'iterated' version of the above-mentioned axiom, a rule with a certain inductive (or rather co-inductive) nature; it ensures a way for checking, before some action is taken, what are the conditions in which this action might lead to new facts becoming common knowledge.

In this paper and the related on-going paper (Baltag, 2000) (to which I have relegated the proofs and the technical details), I am generalizing further the approach developed in our previous work, by adding *non-determinism* and *change of facts*, by extending our previous proofs of

completeness to the present setting; and I apply it to sketch the beginning of an *analysis of information-flow in games*, analysis based on a formalization of rule-based games in terms of epistemic actions.

In ongoing work, I am using variables for actions, which allows me to introduce a fixed-point operator to describe actions that involve epistemic circularity and self-reference. In that more general context, the mutual-learning-operator (originating in some form from Gerbrandy) can simply be defined, via a fixed point expression. I have chosen to keep it simpler here, for reasons of readability. Secondly, I am working on enriching the present approach in order to deal with softer game-theoretic issues. In particular, I am looking at the logic obtained by adding Kleene star (iteration) on top of ours: this would indeed be very useful in a game theoretic context, but the problems of completeness and decidability for this logic are open. More generally, at present I have an extension of the epistemic action logic into an 'epistemic process algebra' and a modal logic going with it: it is obtained by adding a parallel composition and a (process-algebra-style) communication operator between actions. In fact, this development is even more interesting from a game-theoretic perspective: not only we are able to capture in simple, concise formulas the outcome of a game, the projected outcome, counter-factual reasoning, equilibria concepts, rationalizability, epistemic types of players, belief-revision in games, but the setting seems to open up new possibilities, new 'kinds' of games, in which the strategies, the mutual beliefs, the rules of the game, the winning conventions, the number and identity of players etc, are all revisable in the midst of playing the game.

In other ongoing work, I propose a generalization of epistemic actions to capture *probabilistic epistemic actions* (in which the epistemic arrows are replaced by probability distributions and in which our update operation is combined with Bayesian belief-revision). The most promising approach (but potentially very hard) involves combining the probabilistic belief-update with the process-algebra approach and trying to use some of the recent work on process algebras for continuous probabilistic systems, for getting logics that are closer to the hard-core classical game theory.

As a last remark, I would like to express my pious hope that this paper could be the beginning of a more general study of the 'logic of cheating at games': a logic for the suspicious player.

REFERENCES

- Baltag, A. (1999). 'A logic of epistemic actions'. in: *Proceedings of the Workshop on Foundations and Applications of Collective Agent Based Systems*, 11th European Summer School on Logic, Language and Information (ESSLLI'99), Utrecht University.

- Baltag, A. (2000). 'The dynamic logic of epistemic actions', unpublished manuscript.
- Baltag, A., Moss, L. S., Solecki, S. (1998). 'The logic of public announcements, common knowledge and private suspicions (extended abstract)', in: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'98)*, pp. 43–56. Morgan Kaufmann Publishers.
- Baltag, A., Moss, L. S., Solecki, S. (1999). 'The logic of public announcements, common knowledge and private suspicions', CWI technical report SEN-R9922.
- Bleeker, A. and van Eijk, J. (2000). 'Epistemic action and change', presented at LOFT4.
- Fagin, R., Halpern, J., Moses, Y. and Vardi, M. (1995). *Reasoning about Knowledge*. MIT Press.
- Gerbrandy, J. (1999). *Bisimulations on Planet Kripke*, PhD dissertation, University of Amsterdam.
- Gerbrandy, J. and Groeneveld, W. (1997). 'Reasoning about information change', *JLLI*, vol. 6, pp. 147–69.
- van Benthem, J. (2000). 'Information update as relativisation', unpublished manuscript.
- van Ditmarsch, H. P. (2000). *Knowledge games*, PhD dissertation, Groningen University.
- Veltman, F. (1996). 'Defaults in update semantics', *JPL*, vol. 25, pp. 221–61.