On the Dynamics of Information and Abilities of Players in Multi-Player Games

Valentin Goranko[1], Peter Hawke[2]

[1]Technical University of Denmark, [2]University of the Witwatersrand, South Africa

Abstract. The paper intends to contribute towards a more realistic treatment and formalization of the abilities of players to achieve objectives in multi-player games under incomplete, imperfect, or simply wrong information that they may have about the game and about the course of the play. In particular, we aim to develop a logical formalism for dealing with the interplay between the dynamics of information and dynamics of abilities, taking into account both the a priori information of players with respect to the game structure and the empirical information that players develop over the course of an actual play, and associate with these respective information relations and notions of 'a priori' and 'empirical' strategies and strategic abilities.

1 Introduction

The question of determining abilities of individual players and coalitions to achieve formally specified objectives in multi-player games has been a subject of formal logical analysis over the past 10 years, since the introduction of the Coalition Logic CL [9] and the Alternating-time Temporal Logic ATL [1]. In these logics no explicit assumptions about the players' knowledge or beliefs were made, but soon afterwards the first natural attempts to add the epistemic perspective to the scenario was realized by fusing multi-agent epistemic logics MAEL with ATL into the Alternating-time Temporal Epistemic Logic ATEL [7]. Gradually, it was realized that the interplay between information, knowledge, beliefs, and strategic abilities is much more subtle and involved than the former merger of the semantic structures underpinning MAEL and ATL allows to model. Since then, the search for more realistic approaches to formalize that interplay has resulted in numerous publications, of which we only mention here [11] and [8] as two of the first studies on strategic abilities under incomplete information and (im)perfect memory of players and coalitions, that have influenced our work. The present paper initiates a study that builds on these and other works by developing further ideas towards a more realistic treatment and formalization of the abilities of players to achieve objectives in multi-player games under incomplete, imperfect, or simply wrong information that they may have about the game and about the play. In particular, we aim to develop a more refined logical formalism for dealing with the interplay between the *dynamics of information* and *dynamics of abilities.* We see the main contributions of the present work as:

- Expanding the study of the effect of players' knowledge on their abilities by considering players' *information* about a game, that can also involve beliefs and confusions, as opposed to mere uncertainties.
- Highlighting and illustrating the conceptual distinction between the 'a priori' information of a player with respect to the game structure and the 'empirical' information of a player which may evolve over the course of an actual play.
- Respectively, considering and contrasting the abilities of players that rely only on a priori information when executing strategies of games, as opposed to players who are able to take advantage of empirical information, with the possibility of the implementation of *revised* strategies as the game proceeds.
- Proposing a formal logical system, extending and refining ATL and its epistemic extensions developed in [7] and [8], for expressing and formalizing the different types of abilities that take (or do not take) into account player's a priori or empirical information.

Since this paper is a preliminary report, and for lack of space, we emphasize on the conceptual discussion and motivation of the proposed logical formalism, rather than on technical results; such will be discussed only briefly later in the paper, and in more details in a follow-up full paper.

2 Preliminary remarks and some motivating examples

2.1 Incomplete vs imperfect information, and *a priori* vs *empirical* information

We would like to start with some brief terminological remarks, firstly regarding the concepts of 'incomplete' vs 'imperfect' information of a player in a game. Traditionally, in Game Theory the former refers to the knowledge or uncertainties of the player about the *structure and rules of the game*, while the latter refers to the knowledge or uncertainties of the player about *the course of the play* of the game, e.g. about the state in which the game currently is, or the history of the play, or the moves/ actions taken by other players [10, 5].

In the context of this study, we use the term 'information' to incorporate not only the knowledge or uncertainties, but also beliefs and confusions of the players, based on which they would determine their course of actions. Thus, we introduce the notions of 'a priori information' and 'empirical information'. A player's *a priori information*, roughly, is the information (beliefs, knowledge and uncertainty) the player has about the game as such, *prior* to the actual play of the game. Thus, one way to intuitively understand this notion is as referring to what information the player has about the *rules*, *protocol*, or *structure* of the game before the begin of play. On the other hand, *empirical information* refers to the information that a player builds by way of observations, recollections, and reasoning made over the course of play.

2.2 Some examples

We now present several examples of simple games in order to illustrate the distinction and relationship between the a priori information and the empirical information of a player.

In all examples in this section we consider simple turn-based games, played by two players, **1** and **2**, consisting in **1** first making a move, then **2** making a move. Player **1** will always choose between actions a and b, while **2** will choose between c, d and, sometimes, e. States are labelled with numbers, with 0 indicating the start state. Outcomes of the game are expressed in terms of the truth of a certain proposition φ.

Example 1: Consider the following game: player **1** moves first, then player **2** makes a series of moves. The structure of the game is represented in figure 1, which captures what we might think of as the *rules* of the game: the possible states in the game, and the actions that effect transitions between them. Finally, suppose that **2** is *uncertain* about the game in the following sense: **2** is not able to distinguish between states 2 and 3, as he considers the rules of the game. In particular, he considers himself as unable to judge if he is in state 2 (and not state 3), if he were to end up in state 2 in a play of the game - and vice versa. We can represent this situation very naturally by placing a broken line between states 2 and 3 in figure 1. Notice, then, that this uncertainty can be naturally captured within a representation of the *rules* of the game, such as figure 1.

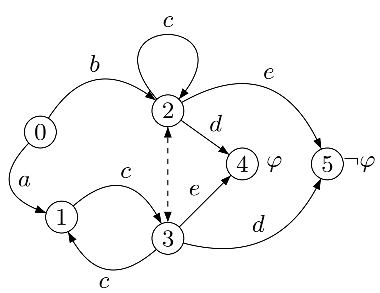

Fig. 1. Example 1. The rules of the game, along with the a priori uncertainty of player 2

Let us consider questions of strategic ability. In particular, does **2** have the ability to guarantee the outcome φ, whatever **1** chooses? Notice, first, that if **2** had *perfect information* about the game, then he certainly could guarantee φ. However, by simply taking into account figure 1 (ie. the rules) and the a priori uncertainty of **2** represented there, it seems clear that **2** does *not* have this ability: if his a priori uncertainty persists through the play of the game, and he ends up in either state 2 or 3, then he will not be certain whether to perform action d or action e in order to achieve φ.

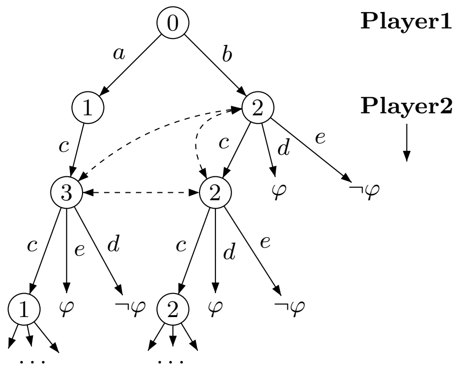

Fig. 2. Example 1. The tree of possible plays of the game

Consider now figure 2, in which the structure in figure 1 has been unravelled into a tree, explicitly representing the possible plays of the game. In this figure, we have again represented **2**'s uncertainty between states 2 and 3. Again, this uncertainty is clearly enough to ensure that **2** will not be able to guarantee φ - if the uncertainty persists. However, when considering the possible plays explicitly, it becomes evident that there may be several reasons why **2**'s information may *evolve* as the game is played. Perhaps **2** will discover, if he actually ends up in state 2, that his memory is good enough to remember the path he took to the current state - and since there is no sequence of states that leads to both state 2 and state 3, his uncertainty will disappear. On the other hand, suppose that the player's memory of the history of the game is not very good (and again that he is in state 2) - in this case, he might *test* which state he is in, by performing action c, which has very different results in state 2 and state 3 (this is, in a sense, a more *active* use of memory). Finally, suppose that, as it happens, **2** finds himself in state 1, after **1** selects action a. In this case, **2** is successfully able to *revise* his strategy such that it will be one that will guarantee the achievement of φ, even if his memory of the history of play is not very good: by committing to play action c immediately, and then c in any state that **2** cannot discern between states 2 and 3, he will guarantee the achievement of φ, so long as he sticks to this strategy (this last point is adapted from [4]).

In summary, there are various ways in which **2** may be able to use *experience of actual play* to revise his initial uncertainty about the game, with a marked effect on the *ability* of the player.

Example 2: Let us consider a second, more nuanced example of where the information of a player may evolve over the course of the game in such a way that the player's informational situation is improved, and his ability along with it. Consider the game in figure 3. Here, after **1** moves, player **2** has the option of two actions to perform. Whichever action he chooses, whatever the outcome, he can then initiate the game again by performing action c. This may be thought of as a 'repeated game', then, in this sense. Further, in this game, **2** again does not have perfect a priori information about the game. Unlike example 1, however, in this case he suffers from a *confusion* about the game, as opposed to 'mere' uncertainty': he is under the mistaken impression that actions a and b, as performed by **1**, have the same effect: they both lead to state 2. We represent this situation in the figure with a broken arrow from state 1 to state 2, indicating that **1** considers state 2 the (only) state he considers possible after the action a is performed.

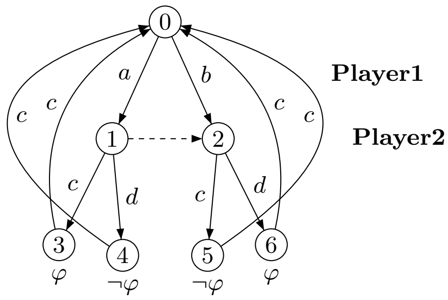

Fig. 3. Example 2. The rules of the game.

Does **2** have the ability to guarantee φ? Again, if his confusion persists throughout the game, then surely not: he will always perform the wrong action in state 1 (namely, action d). Again, his a priori information is not enough to guarantee success. However, we may again consider some possible effects of the experience of actual play on the player's ability. Suppose that the player is able to *learn from experience* in the following manner: after playing the game a few times, the player may eventually come to the realization that **1**'s action a does not have the effect that **2** has been expecting - and thereby revise his confusion about the structure of the game, and his ability to (indefinitely) guarantee φ along with it. Again, then, experience of actual play could conceivably alter the player's abilities.

Example 3: In both of the previous examples, we considered how experience may *refine* the information of a player, thereby granting the player a broader range of abilities. Indeed, if it were always the case that experience refined information (confusions, uncertainties), then it would be the case that if a player had a winning strategy that were compatible with his a priori information about the game, then this strategy would remain a winning strategy for the game even after the player experiences actual play. Put another way, experience would always add to, and never subtract from, the store of winning strategies that a player is able to execute based purely on his a priori information. However, realistically, it is not the case that experience only reduces confusion and uncertainty - indeed, experience of actual play may *create* confusions or uncertainties for a player that were not identifiable a priori.

Consider the following example: **2** moves after **1**, as depicted on the left in figure 4. In this case, suppose that **2** has perfect a priori information about the game: he has no initial uncertainty or confusion about the nature of the game he is to play. Notice that **2** in fact has several winning strategies, if this happy state of perfect certainty persists: he can play d at 1 and d at 2, or d at 1 and e at 2.

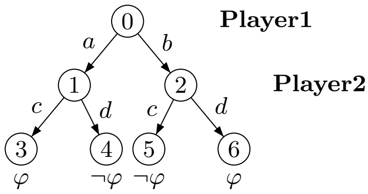

Fig. 4. Example 3. A game in which player 2 has no a priori uncertainty or confusions.

However, suppose that over the course of an actual play of the game, player **2** turns out unable to *observationally* distinguish states 1 and 2, after player player **1** made her move, and before player **2** was to play his move. This might happen, for instance, if **2** failed to observe what action was taken by player **1**, or if **2** were an artificial agent, such as a robot, and one of **2**'s sensors malfunctions, leaving **2** with reduced powers of perception. If such an event were to occur, it becomes clear that some of the strategies that guaranteed **2** success when only considering his a priori information will suddenly become useless: in practice, **2** will now only be able to follow strategies that do not require **2** to perform different actions at states 1 and 2. Thus, **2** is no longer able to win the game.

Example 4: Consider the game depicted in figure 5. Here, player **2** has a priori uncertainty about states 1 and 2: **2** is not certain what the results of **1**'s possible actions will be, and (let us suppose) **2** does not have the ability to discriminate 1 from 2 based on observation alone. If relying only on this a priori information, it is clear that this player does not have the ability to achieve φ.

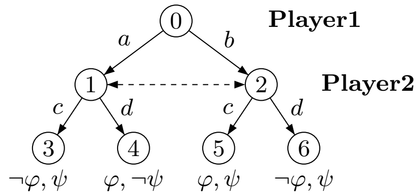

Fig. 5. Example 4. A game with a priori uncertainty of player 2 that can be cleared by reasoning.

However, let us suppose further that **2** knows something about the preferences of **1** -namely, that **1** has ψ as her goal - and that **2** is equipped with the ability to carry out basic game theoretic reasoning. In this case, **2** may use *reasoning* to distinguish 1 and 2, since **2** can conclude that **1** would not allow the game to enter state 1, because the achievement of her goal from that state is not guaranteed, whereas from state 2 it is. Thus, player 2 can predict that from state 0 player 1 will choose action b, thus steering the game into state 2, from which player 2 knows that he should choose action c to achieve his goal.

In this case, we are inclined to say that **2** *has* the ability to achieve φ. However, that ability crucially depends on the information about other players and the *reasoning abilities* of the player, added to his a priori information about the game.

Remark 1. Note that the reasoning described in this example can be naturally regarded as 'a priori reasoning', like most game theoretic reasoning, since it may take place before any actual play of the game. However, it is important not to confuse the results of 'a priori reasoning' in this intuitive sense with the notion of 'a priori information' we use in this paper. Since the results of the 'a priori reasoning' in this example applies to possible *plays* of the game, this information is 'empirical information' in the sense in which we use this term in the paper.

Some observations may be derived from the previous examples. First, the a priori information for a player, about a game, may involve both uncertainties and confusions, and can be understood as applying to the 'structure' or 'rules' of the game. Second, the information of a player may evolve from this initial, a priori information, given experience of play, and there are many mechanisms as to how this evolution may occur - some of which reduce uncertainty and confusion, some of which add to uncertainty or confusion. Finally, there are obvious relationships between a priori information and empirical information in the examples: while it is not always the case that the a priori information *by itself* determines the empirical information of the player, certainly it plays *a significant role* in determining the latter.

3 LODIA: Logic Of Dynamics of Information and Abilities

We now present a logical system for dealing in a more refined and realistic way with issues of imperfect and incomplete information, uncertainties, and beliefs when assessing the players' abilities. Roughly, the key ideas behind that logic are as follows. First, for the structures upon which the logic will be interpreted, we use a variation of the concurrent (epistemic) game structures, as presented in [1] and [7], which we call (for the lack of better name) *concurrent game information structures*. The key addition to these structures is two 'information relations' per player. The first, called 'a priori information relation' relates one state to another if the player considers the second state a possible *structural alternative* for the first, i.e. alternative in the structure of the game. This relation can be used to represent structural uncertainties and beliefs. The second, called 'empirical information relation', relates a *run*, i.e., initial segment of a play of the game, to (possibly) another one which the player considers a possible 'observational alternative' to the first one. This relation can be used to represent empirical uncertainties and beliefs.

Remark 2. Note that we do not impose any general conditions, such as reflexivity, symmetry, transitivity or seriality on any of these relations, though in many specific cases such conditions would naturally arise, e.g., uncertainties are represented by equivalence relations, etc. Indeed, we consider the notion of 'belief' suggested by the information relations to be more general than that usually employed in doxastic logic. For instance: in general, if a player thinks it possible that the (play of the) game is in state y while, in fact, it is in state x, then he need not have a symmetric confusion when the game is indeed in state y; likewise, if the player thinks that the game is in state y when it is in state x, and thinks that the game is in state z when it is in state y, then he need not thinks that the game is in state z when it is in state x.

Using these relations, we will refine the notion of strategic abilities of players encompassed by the logic ATL, in order to distinguish between 'objective' abilities (what the player can achieve *if* they had perfect information), 'a priori' abilities (what the player can achieve based on her a priori information about the game), and 'empirical' abilities (what the player can achieve given that the player can take advantage of, or suffer disadvantage from, experience of actual play).

3.1 Concurrent Information Game Structures

Definition 1 (CGS, [1]) *A* concurrent game structure *(CGS) is a tuple* $\langle k, Q, d, \delta, \Pi, \pi \rangle$*, where:*

- k *is the number of players; we then take the set of players to be* $\mathcal{A} = \{\mathbf{1}, \mathbf{2}, \ldots, \mathbf{k}\}$.
- Q *is a non-empty set of* states.
- $d : Q \times \mathcal{A} \longrightarrow \mathbb{N}$ *is a function where* $d(q, a) > 0$ *(hereafter written as* $d_a(q)$*) represents the number of actions available to player* a *at state* q.
  The set of actions available to a *at* q *will be denoted* $D_a(q) = \{1, 2, \ldots, d_a(q)\}$.
  A joint action at a given state q*, denoted by* σ_q *(or simply by* σ *when* q *is fixed by the context), is a tuple* $\langle j_1, j_2, \ldots, j_k \rangle$*, where* $j_i \leq d_i(q)$ *for every* $i \leq k$*, consisting of a collection of actions, one for each player, that may be performed at state* q*. Given a joint action* $\sigma = \langle j_1, j_2, \ldots, j_k \rangle$*, we write* σ^i *to indicate* j_i*. We write* $D(q)$ *for the set of joint actions at* q*, and* $D(Q)$ *for the set of all joint actions from states in* Q.
- $\delta : Q \times D(Q) \rightarrow Q$ *is the* transition function*, that maps a state* q *and a joint action at* q *to the* successor state *in* Q.
  The set of all successor states of q *will be denoted by* $succ(q)$.
- Π *is a (countable) set of atomic propositions.*
- $\pi : \Pi \longrightarrow \mathcal{P}(Q)$ *is a* valuation *assigning a set of states in* Q *to each atomic proposition.*

One can think of a CGS as capturing the *rules* of a game.

Definition 2 (Game) *A* game *is a tuple* $\langle \mathcal{S}, q \rangle$ *consisting of a CGS with a set of states* Q *and a initial state* $q \in Q$.

Definition 3 *Given a concurrent game structure $\langle k, Q, d, \delta, \Pi, \pi\rangle$, we say that state q' is a* successor *of state q if there is a move vector σ such that $\delta(q, \sigma) = q'$. We denote the set of successors to state q by $succ(q)$.*

Definition 4 (Runs and plays of a game) *Given a CGS $\mathcal{S}$, a* play *of $\mathcal{S}$ is an infinite sequence $\lambda = q_0q_1q_2\ldots$ of states such that $q_{i+1} \in succ(q_i)$ for every $i \in \mathbb{N}$. A* (finite) run *in $\mathcal{S}$ is a finite initial segment of a play: $q_0q_1q_2\ldots q_n$. The last state of a run ρ will be denoted by $l(\rho)$. One-state runs will be identified with the respective states.*

For a play λ, we use $\lambda[i]$, $\lambda[0, i]$ and $\lambda[i, \infty]$ to denote, respectively the ith state of λ, the initial prefix (run) $q_0q_1\ldots q_i$ of λ and the infinite suffix $q_iq_{i+1}\ldots$ of λ.

A q-play/ q-run *is a play/run where $q_0 = q$. We denote the set of q-plays by $Play(q)$ and the set of all q-runs by $Run(q)$. We also denote the set of all plays in $\mathcal{S}$ by $Play(\mathcal{S})$ and the set of all runs in $\mathcal{S}$ by $Run(\mathcal{S})$.*

Definition 5 (CIGS) *A* concurrent informational game structure *(CIGS) is a tuple $(\mathcal{S}; \{\overset{a}{\rightsquigarrow}_{\mathbf{i}}\}_{\mathbf{i}\in\mathcal{A}}, \{\overset{e}{\rightsquigarrow}_{\mathbf{i}}\}_{\mathbf{i}\in\mathcal{A}})$, where $\mathcal{S}$ is a CGS with a set of players $\mathcal{A}$, and for every $\mathbf{i} \in \mathcal{A}$:*

- $\overset{a}{\rightsquigarrow}_{\mathbf{i}} \subseteq Q \times Q$ *is a* a priori information relation *for the player* $\mathbf{i}$;
- $\overset{e}{\rightsquigarrow}_{\mathbf{i}} \subseteq Run(\mathcal{S}) \times Run(\mathcal{S})$ *is an* empirical information relation *for the player* $\mathbf{i}$*, which coincides with $\overset{a}{\rightsquigarrow}_{\mathbf{i}}$ when restricted to one-state runs.*

The intuition: $q_1 \overset{a}{\rightsquigarrow}_{\mathbf{i}} q_2$ if the player $\mathbf{i}$ considers the state q_2 as a possible *a priori* alternative of the state q_1 in the game structure $\mathcal{M}$. Likewise, $\rho_1 \overset{e}{\rightsquigarrow}_{\mathbf{i}} \rho_2$ if the player $\mathbf{i}$ considers the run ρ_2 a possible alternative of the run ρ_1 in $\mathcal{M}$. Initially, the observational information of the player about a play is simply her structural information about the game. As the play progresses, the player may on one hand gain some additional information about the structure of the game, and on the other hand develop some uncertainties or wrong beliefs about the history and the current state of the play.

Some particular cases:

- If $\overset{a}{\rightsquigarrow}_{\mathbf{i}}$ is the equality, then the player $\mathbf{i}$ has a complete (a priori) information about the game $\mathcal{M}$. Likewise, if $\overset{e}{\rightsquigarrow}_{\mathbf{i}}$ restricted to the runs of a given play in $\mathcal{M}$ is the equality (i.e. does not associate any run of the play with any different run), then the player $\mathbf{i}$ has a perfect (empirical) information about that play; and if the whole $\overset{e}{\rightsquigarrow}_{\mathbf{i}}$ is the equality, then the player $\mathbf{i}$ has a perfect (empirical) information about every play of the game $\mathcal{M}$.
- If the player $\mathbf{i}$ has no wrong beliefs, but only uncertainties about the game, then $\overset{a}{\rightsquigarrow}_{\mathbf{i}}$ is an equivalence relation of *a priori indistinguishability.* Likewise, if the player $\mathbf{i}$ has no wrong beliefs, but only uncertainties about the states of a play, then $\overset{e}{\rightsquigarrow}_{\mathbf{i}}$ is an equivalence relation of *empirical indistinguishability.*
- Furthermore, if the player can keep a count of the number of moves made in the play, or has a 'clock' showing how many time units have passed since the beginning of the play (where every time unit corresponds to one transition from a state to a successor state), then $\overset{e}{\rightsquigarrow}_{\mathbf{i}}$ can only relate runs of the same length (this property is known in the literature as *synchronicity* [3]). In general, however, it is conceivable that a player may 'forget' the length of the current run.

Definition 6 (Strategies, [1]) *Given a CIGS $\mathcal{S}$ with a set of states Q and a player $\mathbf{i} \in \mathcal{A}$, a* memoryless strategy for $\mathbf{i}$ in $\mathcal{S}$ *is a function $f_{\mathbf{i}} : Q \longrightarrow \mathbb{N}$; a* (memory-based) strategy for $\mathbf{i}$ in $\mathcal{S}$ *is a function $F_{\mathbf{i}} : Run(Q) \longrightarrow \mathbb{N}$. In both cases, the strategy selects an action that the player should perform from the given state, resp. run.*

A memoryless strategy $f_{\mathbf{i}}$ is a priori uniform *if for any $q_1, q_2 \in Q$, if $q_1 \overset{a}{\rightsquigarrow}_{\mathbf{i}} q_2$ then $f_{\mathbf{i}}(q_1) = f_{\mathbf{i}}(q_2)$. A memory-based strategy $F_{\mathbf{i}}$ is* empirically uniform *if for any $\rho_1, \rho_2 \in Run(Q)$, if $\rho_1 \overset{e}{\rightsquigarrow}_{\mathbf{i}} \rho_2$ then $F_{\mathbf{i}}(\rho_1) = F_{\mathbf{i}}(\rho_2)$.*

The set of plays in a given CIGS that are consistent with a given strategy $f_{\mathbf{i}}$ for a player $\mathbf{i}$ will be denoted by $out(q, f_{\mathbf{i}})$; see e.g., [1] for a precise definition. In the case of a memory-based strategy $F_{\mathbf{i}}$, by $out(q, F_{\mathbf{i}})$ we denote the set of outcome plays in the game $(\mathcal{S}, q)$ starting from the state q.

Remark 3. Note that uniformity treats the information relations just like equivalence relations. Nevertheless, we resist simply *defining* the relations as equivalence relations, for at least the reason that this would in many cases be a misleading representation of the scenarios we are interested in modeling. Furthermore, avoiding the assumption of equivalence relations is likely to make a difference to, first, any procedure for computing empirical information from a priori information and, second, doxastic and/or epistemic operators if these were to be introduced into LODIA as a natural extension.

3.2 Syntax and Semantics of LODIA

For simplicity of the exposition, here we only present the 'one-player' version of the logic. The extension to coalitional abilities is technically straightforward, but conceptually these operators are rather more involved and will be discussed in a follow-up work.

Definition 7 (Syntax of LODIA) *The language $\mathcal{L}_{\boldsymbol{LODIA}}$ of **LODIA** is defined by the following mutual induction on state formulae and path formulae:*

$$\textit{State formulae}: \ \varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \langle\!\langle \mathbf{i} \rangle\!\rangle \theta \mid \langle\!\langle \mathbf{i} \rangle\!\rangle^a \theta \mid \langle\!\langle \mathbf{i} \rangle\!\rangle^e \theta$$

$$\textit{Path formulae}: \ \theta ::= \mathcal{X}\varphi \mid \mathcal{G}\varphi \mid \varphi\,\mathcal{U}\varphi$$

where $p \in \Pi$ and $\mathbf{i} \in \mathcal{A}$.

*Informal semantics of **LODIA**.* The intended interpretation of the formulas involving strategic operators is as follows:

- $\langle\!\langle \mathbf{i} \rangle\!\rangle \theta$: "player **i** has the *objective ability* to achieve outcome θ, without taking into account the epistemic constraints on the player".
  This is the standard ATL interpretation of strategic abilities.
- $\langle\!\langle \mathbf{i} \rangle\!\rangle^a \theta$: "player **i** has the *a priori ability* to achieve outcome θ in the play, only taking into account the a priori information of the player about the game".
- $\langle\!\langle \mathbf{i} \rangle\!\rangle^e \theta$: "player **i** has the *empirical ability* to achieve outcome θ in the play, taking into account the empirical information of the player accumulated during the play".

Given that the player may clear uncertainties and confusions, but also develop new uncertainties or false beliefs during the play, none of these abilities is generally stronger than the other. Actually, the use of the term 'ability' here may be questioned; a more acceptable, though longer, alternative could be 'objective / a priori / empirical assessment of the player's ability'.

Definition 8 (Formal semantics of LODIA) *Let $\mathcal{S}$ be a CIGS and q any state in Q. Then the truth-interpretation of $\mathcal{L}_{\boldsymbol{LODIA}}$-formulae extends the truth-interpretation of ATL-formulae [1] as follows:*

- *$\mathcal{S}, q \models \langle\!\langle \mathbf{i} \rangle\!\rangle \theta$ iff there exists a memoryless strategy $f_{\mathbf{i}}$ for **i** such that $\mathcal{S}, \lambda \models \theta$ for every play $\lambda \in out(q, f_{\mathbf{i}})$.*
- *$\mathcal{S}, q \models \langle\!\langle \mathbf{i} \rangle\!\rangle^a \theta$ iff there exists an a priori uniform strategy $f_{\mathbf{i}}$ for **i** such that $\mathcal{S}, \lambda \models \theta$ for every play $\lambda \in out(q, f_{\mathbf{i}})$.*
- *$\mathcal{S}, q \models \langle\!\langle \mathbf{i} \rangle\!\rangle^e \theta$ iff there exists an empirically uniform strategy $F_{\mathbf{i}}$ for **i** such that $\mathcal{S}, \lambda \models \theta$ for every play $\lambda \in out(q, F_{\mathbf{i}})$.*

The definition of $\langle\!\langle \mathbf{i} \rangle\!\rangle^e$ is meant to take into account that the abilities of a player to achieve certain outcomes in a play of the game would naturally evolve during that play, and that evolution is captured by the notion of empirically uniform strategy.

4 LODIA at work: some remarks

4.1 On validities and other important formulae in LODIA

We wish to emphasize that the logic **LODIA** is *not* of particular interest and importance for us with respect to its validities. We are certainly not inclined to concentrate our studies on capturing such validities through axiomatic system. So, we will restrict our discussion on these to some brief remarks.

To begin with, the validities involving each of the three types of ability operators in **LODIA** alone are obviously included in the validities of the standard coalitional ability operator $\langle\langle A\rangle\rangle$ in the standard CGS-based semantics of the Alternating time temporal logic ATL* (reduced to single-player coalitions) given in [1]; besides, it is unlikely that any of the new versions of that operator will take away any essential property from it. Seeing this should be quite routine for the ATL-like fragment of **LODIA**, by just checking the validity of the ATL axioms from [6] for each of $\langle\langle\mathbf{i}\rangle\rangle$ and $\langle\langle\mathbf{i}\rangle\rangle^a$; indeed, when restricted to ATL each of these has the semantics based on CGSs with incomplete information, and in is known from [6] that memoryless and memory-based semantics for ATL (with complete information) are equivalent. For the full ATL* the validities have not yet been captured axiomatically, so the same approach does not work, but we conjecture that a suitable semantic construction can produce a standard CGS-based model from any model satisfying a given formula involving any of $\langle\langle\mathbf{i}\rangle\rangle$ and $\langle\langle\mathbf{i}\rangle\rangle^a$. We are not going to pursue that problem here any further.

Furthermore, the interaction between the three versions of the ability operator is rather superficial in general, essentially captured by the following validities:

- $\langle\langle\mathbf{i}\rangle\rangle^a\theta \to \langle\langle\mathbf{i}\rangle\rangle\theta$
- $\langle\langle\mathbf{i}\rangle\rangle^e\theta \to \langle\langle\mathbf{i}\rangle\rangle\theta$
- $\langle\langle\mathbf{i}\rangle\rangle^a\mathcal{X}\varphi \leftrightarrow \langle\langle\mathbf{i}\rangle\rangle^e\mathcal{X}\varphi$

The first two state that the a priori or empirical ability to achieve something implies the objective ability to achieve that same thing. The third captures the condition that a priori information and empirical information overlap in the single-move case: something can be achieved in one step using empirical information just in case it can be achieved in one step using a priori information, because if there is no history to the game, there is no difference between a priori information and empirical information. Note, however, that the rule for equivalent replacement from classical logic does not apply here, and the third validity cannot be used for such replacement in the scope of other ability operators.

Here are a few other natural (although, generally invalid) interaction formulae in the language of **LODIA**:

- $\langle\langle\mathbf{i}\rangle\rangle\theta \to \langle\langle\mathbf{i}\rangle\rangle^a\theta$
- $\langle\langle\mathbf{i}\rangle\rangle\theta \to \langle\langle\mathbf{i}\rangle\rangle^e\theta$
- $\langle\langle\mathbf{i}\rangle\rangle^e\theta \to \langle\langle\mathbf{i}\rangle\rangle^a\theta$
- $\langle\langle\mathbf{i}\rangle\rangle^a\theta \to \langle\langle\mathbf{i}\rangle\rangle^e\theta$

The first formula is valid on the class of structures in which player **i** has *complete a priori information.* The second formula is valid on the class of structures in which **i** has *perfect empirical information.* The third formula is valid on the class of structures where empirical information can only ever *subtract* from a priori information, that is the player can only forget or fail to observe or use what he knew about the game, but cannot gain additional insight or knowledge from the play. Conversely, the fourth formula is valid on the class of structures where empirical information can only ever *add* to a priori information, that is the player can only gain additional insight or knowledge during the play, but can never forget or lose in any other way his a priori information about the game.

4.2 On model-checking in LODIA

Our interest in a logic such as **LODIA** is ultimately arising from its use for specifying and verifying abilities of agents (and coalitions) in concrete game structures; that is, in solving the *model-checking problem* relative to such a logic. Thus, our ultimate aim with using **LODIA** is to compute abilities of players in a given model.

For lack of space, here we only briefly discuss the model-checking problem for **LODIA**; it will be treated in more detail in a follow-up work. First, we note again that the fragment of **LODIA** involving only the operator $\langle\langle\rangle\rangle$ for objective ability is essentially ATL* with complete information; the fragment involving only the operator $\langle\langle\rangle\rangle^a$ for a priori ability is essentially ATL* with incomplete information and no memory (recall); and the fragment involving only the operator $\langle\langle\rangle\rangle^e$ for empirical ability is essentially ATL* with incomplete information and memory (recall). It is claimed in [1] that the model-checking problem for the latter is generally undecidable, even for the ATL-fragment.

Thus, one major challenge here is to restrict the semantics by making further assumptions, specific to the problem at hand, that would regain decidability of the model checking, and then to develop practically useful methods for it. We emphasize again that the empirical information of a given player is *not explicitly given* as part of the game model but has to be *computed*, in accordance with the a priori information of a player along plus, possibly, further assumptions about the player's abilities to build up that empirical information during the play. Thus, another major challenge in the practical model-checking of formulae of **LODIA** is the *computing of the empirical information dynamically in the course of the play* and then applying it to determine the truth of formulae involving the operator for empirical ability.

There are at least two important types of assumptions that must be taken into account for the model-checking procedure:

1. The abilities of players with respect to observations, memory, and reasoning during the play of the game. A natural hierarchy of such *comprehension abilities* is:
   (a) *No abilities.* Basically, this is the case of a pre-programmed device performing a sequence of actions no matter how the environment or the other players react to these, and how the play progresses.
   (b) *Observational abilities only.* The player is typically a device with sensors that can observe the current state, and possibly the actions of the other players, but cannot remember the history of the play, neither can perform reasoning on these observations.
   (c) *Observational and memory abilities.* The player has sensors and memory that enable him to observe and remember his observations, thus building up a 'local history' of the play, but cannot perform reasoning based on these.
   (d) *Observational, memory, and reasoning abilities.* A truly intelligent agent, that can observe, remember, and reason on the play of the game.
2. Additional assumptions on the scenario being modeled, regarding e.g.: players' knowledge or beliefs about each other's preferences, knowledge, beliefs, etc. before and during the play of the game; or, the possible occurrence of *contingent events* over the course of play, such as the unexpected failure of players' observational ability or memory, or unexpected gain of knowledge (e.g., a card game player spotting a card of another player), etc.

Clearly, each type of assumptions is crucial for computing the empirical information relation of a given player, and therefore for his empirical abilities to achieve goals. The higher a player is in the comprehension abilities hierarchy, the richer and more difficult to compute is the dynamics of his empirical information and abilities.

In order to illustrate this point, let us revisit some of the examples discussed in section 2.

– Example 1: this example suggests several ways in which an empirical relation may be generated on the assumption that a priori uncertainty amounts to *observational indistinguishability.* If the player is assumed to record all passed stages of a play in his memory and to use that

memory actively, then this will generate an empirical relation that is essentially the identity relation on runs. If, moreover, the player is assumed to be able to *revise* his strategy depending on the course of play, this will generate an empirical information relation that can distinguish the run of states 013 from the runs 02 and 022 - but quite possibly not the other way round.
- Example 2: this example suggests that if the assumption is made that the player has some memory (so that he can remember past mistakes) and is able to revise his strategy after mistakes are executed (after some number n of mistakes), then the empirical relation generated will match the a priori confusions for runs up to a certain length, and then become the identity relation thereafter.
- Example 3: this example suggests that if the assumption is made that *contingent events* such as a failure of a player's observational powers occurs at certain stages of play, then the empirical information relation will be less refined than the a priori information relation in a corresponding manner. It is an interesting technical challenge how such contingencies ought to be modeled.
- Example 4: this example suggests that if the assumption is made that the player **1** has certain preferences and that **2** knows about these preferences and can perform simply reasoning tasks, then the empirical information relation that is generated is again the identity relation.

5 Conclusion

In this paper we have only marked the first steps towards a more realistic formalization and logical treatment of the dynamic interplay between knowledge and abilities of players to achieve objectives in multi-player games.

There are various natural extensions and variations of the logic **LODIA**, including:

- The *coalitional extension* of **LODIA**, involving operators for strategic abilities, in all 3 senses presented above, for any groups of players, just like in ATL*.
- The *epistemic and doxastic extensions* of **LODIA** involve operators for a priori structural knowledge and beliefs, as well as empirical knowledge and beliefs of players and coalitions. These operators come in all familiar flavors: individual, group, common, and distributed knowledge, thus extending and refining the logics ATEL [7], ATOL and ATEL-R* [8], etc.
  A particularly interesting case for study here is to analyze the dynamics of empirical information in the setting of extensions of ATL with *dynamic epistemic logics* [2].

Finally, we concede that the semantics of **LODIA** remains too abstract for any practical purposes. A more concrete and refined semantics would allow for a more realistic interpretation of the concepts introduced in this paper; this is a subject of further work.

References

1. Alur, R., T. A. Henzinger, O. Kupferman, "Alternating-time Temporal Logic", in *Journal of the ACM* **49**, pp. 672-713 (2002). Updated, improved and extended text. Available at www.cis.upenn.edu/alur/Jacm02.pdf.
2. van Ditmarsch, H., W. van der Hoek, B. Kooi *Dynamic Epistemic Logic*, Springer, Synthese Library, 2007.
3. Fagin R., J. Halpern, Y. Moses, M. Y. Vardi, *Reasoning About Knowledge*, MIT Press (1995).
4. Halpern, J. Y., "On Ambiguities in the Interpretation of Game Trees", in *Games and Economic Behaviour* **20**, pp.66-96 (1997).
5. Harsanyi, J., "Games With Incomplete Information Played by Bayseian Players", in *Management Science*, Vol 14, No 3 (1957).
6. Goranko, V., G. van Drimmelen, "Complete Axiomatization and Decidability of the Alternating-time Temporal Logic", in *Theoretical Computer Science*, **353**, pp 93–117, 2006.
7. van der Hoek, W., M. Wooldridge, "Cooperation, Knowledge and Time: Alternating-time Temporal Epistemic Logic and its Applications", in *Studia Logica* **75**, pp. 125-127, Kluwer Academic Publishers (2003).

8. van der Hoek, W., Jamroga, W., “Agents Who Know How To Play”, in *Fundamenta Informaticae* **62**, pp. 1-35, IOS Press (2004).
9. Pauly, M., *Logic for Social Software*, PhD dissertation, ILLC dissertation series (2001).
10. Rasmusen, E., *Games and Information: An Introduction to Game Theory*, Chapter 2, Blackwell (1989)
11. Schobbens, P. Y., “Alternating-time logic with imperfect recall”, in *Electronic notes in Theoretical Computer Science*, **85(2)**, pp.82-93, 2004. Proceedings of LCMAS (2003).