Comparing Variants of Strategic Ability

Wojciech Jamroga and Nils Bulling

[1] Computer Science and Communication, University of Luxembourg
wojtek.jamroga@uni.lu
[2] Department of Informatics, Clausthal University of Technology, Germany
bulling@in.tu-clausthal.de

Abstract. In this paper, we show that different semantics of ability give rise to different validity sets. The issue is important for several reasons. First, many logicians identify a logic with its set of true sentences. As a consequence, we prove that different notions of ability induce different *logics* in the traditional sense. Secondly, the study can be seen as the first systematic step towards satisfiability-checking algorithms for variants of **ATL** other than the basic variant.

1 Introduction

Alternating-time temporal logic (**ATL**) [2] is a temporal logic that incorporates some basic game theoretical notions. In **ATL** we can for instance express that a group of agents is able to *bring about* φ, i.e., they are able to enforce that property φ holds whatever the other agents might do. **ATL** has been studied extensively in previous years; however, most of the research was focused on the way such logics can be used for verification of multi-agent systems. In particular, the complexity of model checking was investigated and compared for different settings and different variants of the logic [10, 11, 7]. Studies of other formal meta-properties have been relatively scarce. Axiomatization and satisfiability were investigated in [4, 12, 9], some expressivity issues were raised in [8], and invariance of the semantics with respect to a couple of classes of models was proven in [3]. Moreover, most of these studies were limited to the basic variant of **ATL** where agents possess perfect information about the current state of the system and perfect memory about the evolution of the system so far. Formal properties of other variants of **ATL** are largely left untouched, and so is the comparison between different semantic variants of the logic.

Semantic variants of **ATL** are usually derived from different assumptions about agents' capabilities. Can the agents "see" the current state of the system, or only a part of it? Can they memorize the whole history of observations in the game? Is it enough that they can enforce the required temporal property "objectively", or should they be able to come up with the right strategy on their own? Different answers to these questions induce different semantics of strategic ability, and they clearly give rise to different analyses of a given problem domain. However, it is not entirely clear to what extent they give rise to different *logics*.

One natural question that arises in this respect is whether these semantic variants generate different sets of valid (and, dually, satisfiable) sentences.

The question is important for several reasons. First, many logicians identify a logic with the set of sentences that are true in the logic; semantics is just a possible way of defining the set, alternative to an axiomatic inference system. Thus, by comparing validity sets we compare the respective logics in the traditional sense. Moreover, validities of **ATL** capture general properties of games under consideration: if, e.g., two variants of **ATL** generate the same valid sentences then the underlying notions of ability induce the same kind of games. Finally, the satisfiability problem for **ATL**, though far less studied than model checking, is not necessarily less important. While model checking **ATL** can be seen as a modal logic analogue of game solving, satisfiability corresponds naturally to mechanism design. A systematic study on the abstract level is the first step towards devising algorithms that solve the problem.

So, do different notions of ability induce different strategic logics (in the traditional sense)? Informal discussions revealed that the issue is by no means obvious. While it was clear for some researchers that we should get different validity sets, for others it was equally clear that validities should stay the same regardless of the actual notion of strategy. In this paper, we settle the issue and show that most of the known semantic variants of **ATL** are indeed different, and we characterize the relationship between their sets of validities.

Our results are relevant also outside the logical context. As already mentioned, by looking at validity sets we study general properties of strategic ability under various semantic assumptions. And since much of our study concerns ability under imperfect information, one can claim that it regards abilities of agents in *realistic* multi-agent systems. Ultimately, we show that what agents can achieve is much more sensitive to the strategic model of an agent (and a precise notion of achievement) than it was generally realized. No less importantly, our study reveals that some natural properties – usually taken for granted when reasoning about action – may cease to be universally true if we change the strategic setting. Prime examples include fixpoint characterizations of temporal/strategic operators (that enable incremental synthesis and iterative execution of strategies) and the duality between necessary and obtainable outcomes in a game.

The paper is structured as follows. We begin with presenting the relevant variants of **ATL** in Section 2. Then, we investigate the relationships between their validity sets in Section 3. Finally, we present some conclusions in Section 4.

2 Reasoning about Abilities

ATL [2] generalizes the branching time logic **CTL** by replacing path quantifiers E, A with *cooperation modalities* $\langle\langle A\rangle\rangle$. Informally, $\langle\langle A\rangle\rangle\gamma$ expresses that the group of agents A has a collective strategy to enforce temporal property γ. **ATL** formulae include temporal operators: "$\bigcirc$" ("in the next state"), "$\Box$" ("always from now on") and $\mathcal{U}$ ("until"). The additional operator "$\Diamond$" ("now or sometime in the future") can be defined as $\Diamond\gamma \equiv \top\,\mathcal{U}\,\gamma$.

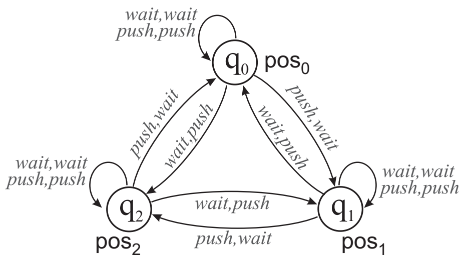

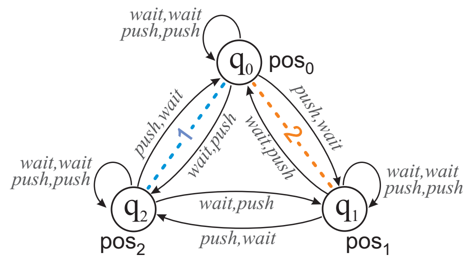

Fig. 1. Two robots and a carriage: (A) Perfect information model M_0 (left); (B) Imperfect information **CGS** M'_0 (right). Dashed lines represent indistinguishability relations between states.

2.1 Syntax of ATL

In the rest of the paper we assume that Π is a nonempty set of *proposition symbols* and $\mathbb{A}\mathrm{gt}$ a nonempty and finite set of *agents*. Alternating-time temporal logic comes in several syntactic variants, of which **ATL*** is the broadest. Formally, the language of **ATL*** is given by formulae φ generated by the grammar below, where $A \subseteq \mathbb{A}\mathrm{gt}$ is a set of agents, and $\mathsf{p} \in \Pi$ is an atomic proposition:

$$\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi \land \varphi \mid \langle\!\langle A \rangle\!\rangle \gamma, \qquad \gamma ::= \varphi \mid \neg\gamma \mid \gamma \land \gamma \mid \bigcirc\gamma \mid \gamma\,\mathcal{U}\,\gamma.$$

The best known syntactic variant of alternating time temporal logic is "**ATL** without star" (or "vanilla" **ATL**) in which every occurrence of a cooperation modality is uniquely coupled with a temporal operator. **ATL**$^+$ sits between **ATL*** and "vanilla" **ATL**: it allows cooperation modalities to be followed by *a Boolean combination* of simple temporal subformulae. We will use the acronym **ATL** to refer to the "**ATL** without star" when no confusion can arise.

2.2 Basic Models of ATL

In [2], the semantics of alternating-time temporal logic is defined over a variant of transition systems where transitions are labeled with combinations of actions, one per agent. Formally, a *concurrent game structure* (**CGS**) is a tuple $M = \langle \mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o \rangle$ which includes a nonempty finite set of all agents $\mathbb{A}\mathrm{gt} = \{1, \dots, k\}$, a nonempty set of states St, a set of atomic propositions Π and their valuation $\pi : \Pi \to 2^{St}$, and a nonempty finite set of (atomic) actions Act. Function $d : \mathbb{A}\mathrm{gt} \times St \to 2^{Act}$ defines nonempty sets of actions available to agents at each state, and o is a (deterministic) transition function that assigns the outcome state $q' = o(q, \alpha_1, \dots, \alpha_k)$ to state q and a tuple of actions $\langle \alpha_1, \dots, \alpha_k \rangle$ for $\alpha_i \in d(i, q)$ and $1 \le i \le k$, that can be executed by $\mathbb{A}\mathrm{gt}$ in q.

A *path* $\lambda = q_0 q_1 q_2 \dots$ is an infinite sequence of states such that there is a transition between each q_i, q_{i+1}. We use $\lambda[i]$ to denote the ith position on path λ (starting from $i = 0$) and $\lambda[i, \infty]$ to denote the subpath of λ starting from i.

The set of paths starting in q is denoted by $\Lambda_M(q)$. Moreover, we define $\Lambda_M^{fin}(q)$ as the set of all finite prefixes of $\Lambda_M(q)$ and we define $\Lambda_M^{fin} := \bigcup_{q \in St} \Lambda_M^{fin}(q)$.

Example 1 (Robots and Carriage). Consider the scenario depicted in Figure 1A. Two robots push a carriage from opposite sides. As a result, the carriage can move clockwise or anticlockwise, or it can remain in the same place. We assume that each robot can either push (action *push*) or refrain from pushing (action *wait*). Moreover, they both use the same force when pushing. Thus, if the robots push simultaneously or wait simultaneously, the carriage does not move. When only one of the robots is pushing, the carriage moves accordingly.

2.3 Basic Semantics of ATL

In the standard version of **ATL** [2], strategies are represented by functions $s_a : St^+ \to Act$. A *collective strategy* for a group of agents $A = \{a_1, \dots, a_r\}$ is simply a tuple of individual strategies $s_A = \langle s_{a_1}, \dots, s_{a_r} \rangle$. Let $a \in A$; by $s_A|_a$, we will denote agent a's part s_a of the collective strategy s_A. The "outcome" function $out(q, s_A)$ returns the set of all paths that may occur when agents A execute strategy s_A from state q onward:

$out(q, s_A) = \{\lambda = q_0 q_1 q_2 \dots \mid q_0 = q$ and for each $i = 1, 2, \dots$ there exists a tuple of all agents' decisions $\langle \alpha_1^{i-1}, \dots, \alpha_k^{i-1} \rangle$ such that $\alpha_a^{i-1} \in d_a(q_{i-1})$ for every $a \in \mathbb{A}\mathrm{gt}$, and $\alpha_a^{i-1} = s_A|_a(q_0 q_1 \dots q_{i-1})$ for every $a \in A$, and $o(q_{i-1}, \alpha_1^{i-1}, \dots, \alpha_k^{i-1}) = q_i\}$.

Let M be a **CGS**, q a state, and λ a path in M. The semantics of **ATL*** and its sublanguages can be defined by the standard clauses for Boolean and temporal operators, plus the following clause for $\langle\langle A \rangle\rangle$ (cf. [2] for details):

$M, q \models \langle\langle A \rangle\rangle \gamma$ iff there is a strategy s_A for agents A such that for each path $\lambda \in out(s_A, q)$, we have $M, \lambda \models \gamma$.

Example 2 (Robots and Carriage, ctd.). Since the outcome of each robot's action depends on the current action of the other robot, no agent can make sure that the carriage moves to any particular position. So, we have for example that $M_0, q_0 \models \neg \langle\langle 1 \rangle\rangle \Diamond \mathsf{pos}_1$. On the other hand, the robots can cooperate to move the carriage. For instance, it holds that $M_0, q_0 \models \langle\langle 1, 2 \rangle\rangle \Diamond \mathsf{pos}_1$ (example strategy: robot 1 always pushes and robot 2 always waits).

2.4 Some Important Formulae

The following fixpoint properties are valid in the standard semantics of **ATL** (that follows from the correctness of the model checking algorithm proposed in [2]):

$$\langle\langle A \rangle\rangle \Box \varphi \leftrightarrow \varphi \wedge \langle\langle A \rangle\rangle \bigcirc \langle\langle A \rangle\rangle \Box \varphi$$
$$\langle\langle A \rangle\rangle \varphi_1 \,\mathcal{U}\, \varphi_2 \leftrightarrow \varphi_2 \vee \varphi_1 \wedge \langle\langle A \rangle\rangle \bigcirc \langle\langle A \rangle\rangle \varphi_1 \,\mathcal{U}\, \varphi_2.$$

Moreover, the path quantifiers A, E of **CTL** can be expressed in **ATL** with $\langle\langle\emptyset\rangle\rangle, \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle$ respectively. As a consequence, the **CTL** duality axioms can be rewritten in **ATL**, and become validities in the standard semantics:

$$\neg\langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Diamond\varphi \leftrightarrow \langle\langle\emptyset\rangle\rangle\Box\neg\varphi,$$
$$\neg\langle\langle\emptyset\rangle\rangle\Diamond\varphi \leftrightarrow \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Box\neg\varphi.$$

2.5 Between Uncertainty and Recall

It makes sense, from a conceptual and computational point of view, to distinguish between two types of strategies: an agent may base its decision on the current state or on the whole history of events that have happened. Also, the agent may have complete or incomplete knowledge about the current global state of the system. To distinguish between those cases, a natural taxonomy of 4 strategy types was introduced in [10] and labeled as follows: i (resp. I) stands for *imperfect* (resp. *perfect*) *information*, and r (resp. R) refers to *imperfect* (resp. *perfect*) *recall*. Then, the semantics of **ATL** can be parameterized with the strategy type that we use – this way we obtain 4 different semantic variants of the logic, labeled accordingly ($\mathbf{ATL}_{\mathrm{IR}}$, $\mathbf{ATL}_{\mathrm{Ir}}$, $\mathbf{ATL}_{\mathrm{iR}}$, and $\mathbf{ATL}_{\mathrm{ir}}$). In this paper, we extend the taxonomy slightly by adding the distinction between *objective* and *subjective* abilities under imperfect information, denoted by $\mathrm{i_s}$ and $\mathrm{i_o}$, respectively.

The semantic variants are defined as follows: models, *imperfect information concurrent game structures* (**iCGS**), can be seen as concurrent game structures augmented with a family of indistinguishability relations $\sim_a \subseteq St \times St$, one per agent $a \in \mathbb{A}\mathrm{gt}$. The relations describe agents' uncertainty: $q \sim_a q'$ means that agent a cannot distinguish between states q and q' of the system. Each $\sim_a$ is assumed to be an equivalence relation. It is also required that agents have the same choices in indistinguishable states: if $q \sim_a q'$ then $d(a, q) = d(a, q')$.

We define a *history* to be a sequence of states. Two histories $h = q_0 q_1 \ldots q_n$ and $h' = q'_0 q'_1 \ldots q'_{n'}$ are *indistinguishable for agent* a ($h \approx_a h'$) iff $n = n'$ and $q_i \sim_a q'_i$ for $i = 1, \ldots, n$. Concatenation of h and h' is denoted by $h \circ h'$ or simply hh'. We also use $last(h)$ to refer to the last state on history h. Additionally, for any equivalence relation $\mathcal{R}$ over a set X we use $[x]_{\mathcal{R}}$ to denote the equivalence class of x. Moreover, we will use the abbreviations $\sim_A := \bigcup_{a \in A} \sim_a$ and $\approx_A := \bigcup_{a \in A} \approx_a$. We also write $\sim^M_A, \approx^M_A$ if the model is not clear from the context. Note that relations $\sim_A, \approx_A$ implement the "everybody knows" type of collective knowledge (i.e., q and q' are indistinguishable for group A iff there is at least one agent in A for whom q, q' look the same).

The following types of strategies are used in the respective semantic variants:

- Ir: $s_a : St \to Act$ such that $s_a(q) \in d(a, q)$ for all q;
- IR: $s_a : St^+ \to Act$ such that $s_a(q_0 \ldots q_n) \in d(a, q_n)$ for all $q_0, \ldots, q_n$;
- ir (i.e., $\mathrm{i_s}$r or $\mathrm{i_o}$r): like Ir, with the additional constraint that $q \sim_a q'$ implies $s_a(q) = s_a(q')$;
- iR (i.e., $\mathrm{i_s}$R or $\mathrm{i_o}$R): like IR, with the additional constraint that $h \approx_a h'$ implies $s_a(h) = s_a(h')$.

That is, strategy s_a is a conditional plan that specifies a's action in each state of the system (for memoryless agents) or for every possible history of the system evolution (for agents with perfect recall). Moreover, imperfect information strategies specify the same choices for indistinguishable states (resp. histories).

A collective xy-strategy s_A is a tuple of individual xy-strategies s_a, one per $a \in A$. Note that the constraints in collective strategies refer to individual choices and individual relations $\sim_a$ (resp. $\approx_a$), and not to collective choices and the derived relations $\sim_A$ (resp. $\approx_A$). The set of outcomes of a strategy is redefined as follows:

– $out^{xy}(q, s_A) = out(q, s_A)$ for $x \in \{\mathrm{I}, \mathrm{i_o}\}$ and $y \in \{\mathrm{r}, \mathrm{R}\}$;
– $out^{xy}(q, s_A) = \bigcup_{q \sim_A q'} out(q', s_A)$ for $x = \mathrm{i_s}$ and $y \in \{\mathrm{r}, \mathrm{R}\}$.

We obtain the semantics for $\mathbf{ATL}_{\mathrm{xy}}$ by changing the clause for $\langle\langle A\rangle\rangle\gamma$ from Section 2.3 in the following way:

$M, q \models_{xy} \langle\langle A\rangle\rangle\gamma$ iff there is an xy-strategy s_A for agents A such that for each path $\lambda \in out^{xy}(q, s_A)$, we have $M, \lambda \models_{xy} \gamma$.

Note that the I and $\mathrm{i_o}$ semantics of **ATL** look only at outcome paths starting from the current global state of the system. In other words, they formalize the properties which agents can enforce *objectively* (but, in case of uncertainty about the current state, they may be unaware of the fact). In contrast, the $\mathrm{i_s}$ semantics of $\langle\langle A\rangle\rangle\gamma$ refers to all outcome paths starting from states that look the same as the current state for coalition A. Hence, it formalizes the notion of A *knowing how to play* in the sense that group A can identify a single strategy that succeeds from all the states they consider possible. We follow [10] by taking the "everybody knows" interpretation of collective uncertainty. A more general setting was proposed in [6]; we believe that the results presented in this paper carry over to the other cases of "knowing how to play", too.

Example 3 (Robots and Carriage, ctd.). We refine the scenario from Example 1 by assuming that the first robot can distinguish between position 0 and position 1, but positions 0 and 2 look the same to it. Likewise, the second robot can distinguish between positions 0 and 2, but not 0 and 1 (cf. Figure. 1B). Now, we have that $M_0', q_0 \models_{xy} \neg\langle\langle 1\rangle\rangle\Box\neg\mathsf{pos}_1$ for all $x \in \{\mathrm{i_s}, \mathrm{i_o}\}$, $y \in \{\mathrm{r}, \mathrm{R}\}$ (that is, for all variants with imperfect information). Note in particular that the strategy from Example 1 cannot be used here because it is not uniform. The robots can achieve the task together, but only in the objective sense: $M_0', q_0 \models_{\mathrm{i_o r}} \langle\langle 1, 2\rangle\rangle\Box\neg\mathsf{pos}_1$. However, they cannot identify a strategy which guarantees that: $M_0', q_0 \models_{\mathrm{i_s r}} \neg\langle\langle 1, 2\rangle\rangle\Box\neg\mathsf{pos}_1$ (when in q_0, robot 2 considers it possible that the current state of the system is q_1, in which case all hope is gone).

So, do the robots know how to play to achieve anything? Yes, for example they know how to make the carriage reach a particular state eventually: $M_0', q_0 \models_{\mathrm{i_s r}} \langle\langle 1, 2\rangle\rangle\Diamond\mathsf{pos}_1$ etc. – it suffices that one of the robots pushes all the time and the other waits all the time.

It is important to note that in "vanilla" **ATL** both semantics for perfect information coincide:

Proposition 1 ([2, 10]). *For every* **iCGS** M*, state* q*, and* ***ATL*** *formula* φ*, we have that* $M, q \models_{\mathbf{IR}} \varphi$ *iff* $M, q \models_{\mathbf{Ir}} \varphi$*.*

3 Comparing Validities for Variants of ATL

In this section we present a formal comparison of the semantic variants defined in Section 2. As stated in the introduction, we compare the variants on the level of their validity sets (or, equivalently, satisfiable sentences). In most cases, the variants turn out to be different. Also in most cases, we can show that one variant is a refinement of the other in the sense that its set of validities strictly subsumes the validities induced by the other variant.

In what follows, we write $Val(\mathbf{ATL}_{\mathrm{sem}})$ to denote the set of **ATL** validities under semantics *sem*. Likewise, we write $Sat(\mathbf{ATL}_{\mathrm{sem}})$ for the set of **ATL** formulae satisfiable in *sem*.

3.1 Perfect vs. Imperfect Information

We begin by comparing models of perfect and imperfect information scenarios. That is, in the first class (I), agents recognize the current global state of the system by definition. In the latter (i), uncertainty of agents about states can be encoded.

Comparing $\mathbf{ATL_{ir}}$ vs. $\mathbf{ATL_{Ir}}$ First, we observe that perfect information can be seen as a special case of imperfect information.

Proposition 2. $Val(\boldsymbol{ATL}_{\mathrm{i_s r}}) \subseteq Val(\boldsymbol{ATL}_{\mathrm{Ir}})$ *and* $Val(\boldsymbol{ATL}_{\mathrm{i_o r}}) \subseteq Val(\boldsymbol{ATL}_{\mathrm{Ir}})$*.*

Proof. Since perfect information of agents can be explicitly represented in **iCGS** by fixing all relations $\sim_a$ as the minimal reflexive relations ($q \sim_a q'$ iff $q = q'$), we have that $\varphi \in Sat(\mathbf{ATL}_{\mathrm{Ir}})$ implies $\varphi \in Sat(\mathbf{ATL}_{\mathrm{i_s r}})$ and $\varphi \in Sat(\mathbf{ATL}_{\mathrm{i_o r}})$. Thus, dually, $Val(\mathbf{ATL}_{\mathrm{i_s r}}) \subseteq Val(\mathbf{ATL}_{\mathrm{Ir}})$ and $Val(\mathbf{ATL}_{\mathrm{i_o r}}) \subseteq Val(\mathbf{ATL}_{\mathrm{Ir}})$.

Proposition 3. $Val(\boldsymbol{ATL}_{\mathrm{Ir}}) \not\subseteq Val(\boldsymbol{ATL}_{\mathrm{i_s r}})$*.*

Proof. We show that by presenting a validity for $\mathbf{ATL}_{\mathrm{Ir}}$ which is not valid in $\mathbf{ATL}_{\mathrm{i_s r}}$. Consider the formula that captures the right-to-left direction in the fixpoint characterization of $\langle\!\langle a \rangle\!\rangle\Diamond$:

$$\Phi_1 \equiv \mathsf{p} \lor \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p} \rightarrow \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$$

Φ_1 is Ir-valid (cf. Section 2.4). To see its invalidity in the $\mathrm{i_s r}$ semantics, consider model M_1 from Figure 2.[3] Indeed, for $\mathsf{p} \equiv \mathsf{shot}$, we get $M_1, q_0 \models_{\mathrm{i_s r}} \mathsf{p} \lor \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$ and $M_1, q_0 \not\models_{\mathrm{i_s r}} \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$, which formally concludes our proof.

[3] The story behind Figure 2 is as follows. A man wants to shoot down a duck in a shooting gallery. (With all respect to animal rights, this is just a yellow rubber duck,

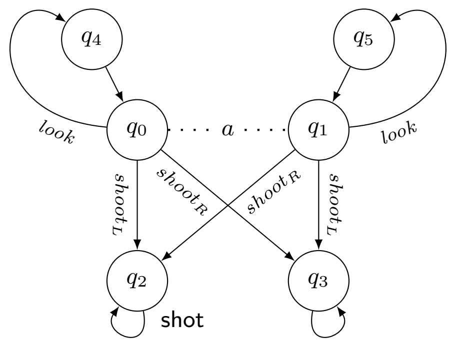

Fig. 2. Model M_1: shooting the poor duck. The model includes only one player (a), and transitions are labeled with a's actions. Automatic transitions (i.e., such that there is only one possible transition from the starting state) are left unlabeled.

Proposition 4. *$Val(\boldsymbol{ATL}_{\mathrm{Ir}}) \not\subseteq Val(\boldsymbol{ATL}_{\mathrm{i_o r}})$.*

Proof. It is sufficient to show that $\Phi_1 \equiv \mathsf{p} \lor \langle\langle a\rangle\rangle \bigcirc \langle\langle a\rangle\rangle \Diamond \mathsf{p} \to \langle\langle a\rangle\rangle \Diamond \mathsf{p}$ is invalid in the $\mathrm{i_o r}$ semantics. Take model M_2 in Figure 3 and $\mathsf{p} \equiv \mathsf{shot}$. Now we have that $M_2, q_0' \models_{\mathrm{i_o r}} \mathsf{p} \lor \langle\langle a\rangle\rangle \bigcirc \langle\langle a\rangle\rangle \Diamond \mathsf{p}$ and $M_2, q_0' \not\models_{\mathrm{i_o r}} \langle\langle a\rangle\rangle \Diamond \mathsf{p}$, which concludes the proof.

Corollary 1. *$Val(\boldsymbol{ATL}_{\mathrm{i_s r}}) \subsetneq Val(\boldsymbol{ATL}_{\mathrm{Ir}})$ and $Val(\boldsymbol{ATL}_{\mathrm{i_o r}}) \subsetneq Val(\boldsymbol{ATL}_{\mathrm{Ir}})$.*

Comparing $\mathbf{ATL_{iR}}$ vs. $\mathbf{ATL_{IR}}$ First, we observe that for $\mathbf{ATL}_{\mathrm{i_o R}}$ vs. $\mathbf{ATL}_{\mathrm{IR}}$ we can employ the same reasoning as for for $\mathbf{ATL}_{\mathrm{i_o r}}$ vs. $\mathbf{ATL}_{\mathrm{Ir}}$. Abilities under perfect information can be still seen as a special case of imperfect information abilities, and we can use the same model M_2 to invalidate the same formula Φ_1. Thus, analogously to Corollary 1 we get:

Corollary 2. *$Val(\boldsymbol{ATL}_{\mathrm{i_o R}}) \subsetneq Val(\boldsymbol{ATL}_{\mathrm{IR}})$.*

By the same reasoning as above, $Val(\mathbf{ATL}_{\mathrm{i_s R}}) \subseteq Val(\mathbf{ATL}_{\mathrm{IR}})$. To settle the other direction, we need to use another counterexample, though.

so there is no reason to worry about it.) The man knows that the duck is in one of the two cells in front of him, but he does not know in which one. Moreover, this has been a long party, and he is very tired, so he is only capable of using memoryless strategies at the moment. Does he have a memoryless strategy which he knows will achieve the goal? No. He can either decide to shoot to the left, or to the right, or reach out to the cells and look what is in (note also that the cells close in the moment after being opened). In each of these cases the man risks that he will fail (at least from his subjective point of view). Does he have an opening strategy that he knows will guarantee his knowing how to shoot the duck in the next moment? Yes. The opening strategy is to look; if the system proceeds to q_4 then the second strategy is to shoot to the left, otherwise the second strategy is to shoot to the right.

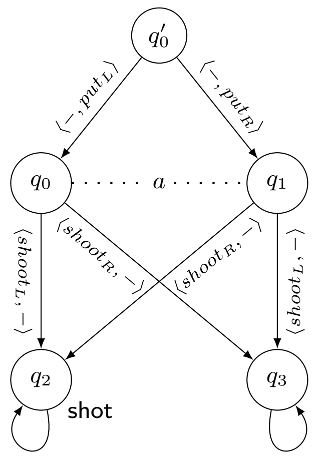

Fig. 3. M_2: a variant of the "poor duck problem" with 2 agents a, b. This time, we explicitly represent the agent (b) who puts the duck in one of the cells.

Proposition 5. *$Val(\boldsymbol{ATL}_{\mathrm{IR}}) \not\subseteq Val(\boldsymbol{ATL}_{\mathrm{i_sR}})$.*

Proof. This time we consider the other direction of the fixpoint characterization for $\langle\!\langle a \rangle\!\rangle \Diamond$: $\Phi_2 \equiv \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p} \to \mathsf{p} \vee \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$. Φ_2 is IR-valid, but it is not valid in $\mathrm{i_s R}$. Consider the "poor duck model" M_1 from Figure 2 and $\mathsf{p} \equiv \mathsf{shot}$. We have that $M_1, q_4 \models_{\mathrm{i_sR}} \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$, but $M_1, q_4 \not\models_{\mathrm{i_sR}} \mathsf{p} \vee \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$, which concludes the proof.

Corollary 3. *$Val(\boldsymbol{ATL}_{\mathrm{i_sR}}) \subsetneq Val(\boldsymbol{ATL}_{\mathrm{IR}})$ and $Val(\boldsymbol{ATL}_{\mathrm{i_sR}}) \subsetneq Val(\boldsymbol{ATL}_{\mathrm{Ir}})$.*

3.2 Perfect Recall vs. Memoryless Strategies

Now we proceed to examine the impact of perfect vs. imperfect *recall* on the general strategic properties of agent systems.

Comparing $\mathbf{ATL_{Ir}}$ vs. $\mathbf{ATL_{IR}}$ We have already mentioned that, in "vanilla" **ATL**, the Ir and IR semantics coincide (Proposition 1). As a consequence, they induce the same vailidities: $Val(\mathbf{ATL}_{\mathrm{Ir}}) = Val(\mathbf{ATL}_{\mathrm{IR}})$. Thus, regardless of the type of their recall, perfect information agents possess the same abilities with respect to winning conditions that can be specified in "vanilla" **ATL**. An interesting question is: does it carry over to more general classes of winning conditions, or are there (broader) languages that can discern between the two types of ability? The answer is: yes, there are. The Ir- and IR-semantics induce different validity sets for $\mathbf{ATL}^*$, and in fact the distinction is already present in $\mathbf{ATL}^+$. Moreover, it turns out that perfect recall can be seen as a special case of imperfect recall in the sense of their general properties.

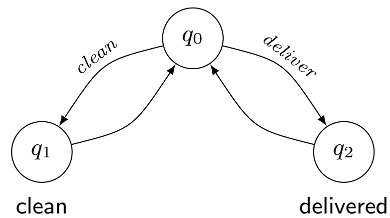

Fig. 4. Model M_3: robot with multiple tasks

Our proof of Proposition 6 draws inspiration from the proof of [1, Theorem 8.3]. We start with some additional notions and two useful lemmata.

Definition 1 (Tree-like CGS). *Let M be a* **CGS** *and q be a state in it. M is called* tree-like *iff there is a state q (the* root*) such that for every q' there is a unique finite sequence of states leading from q to q'.*

Definition 2 (Tree unfolding). *Let $M = (Agt, St, \Pi, \pi, Act, d, o)$ be a* **CGS** *and q be a state in it. The tree-unfolding of M starting from state q denoted $T(M,q)$ is defined as $(\mathbb{A}\mathrm{gt}, St', \Pi, \pi', Act, d', o')$ where $St' := \Lambda_M^{fin}(q)$, $d'(a,h) := d(a, last(h))$, $o'(h, \boldsymbol{\alpha}) := h \circ o(last(h), \boldsymbol{\alpha})$, and $\pi'(h) := \pi(last(h))$.*

Lemma 1. *For every tree-like* **CGS** *M, state q in M, and $\boldsymbol{ATL}^*$ formula φ, we have: $M, q \models_{\mathbf{Ir}} \varphi$ iff $M, q \models_{\mathbf{IR}} \varphi$.*

Proof (sketch). Induction over the structure of φ. The main case is $\varphi \equiv \langle\!\langle A \rangle\!\rangle \gamma$; to see that the proof goes through, observe that the subtree of M starting from q is also a tree-like **CGS**, and on tree-like **CGS**'s Ir and IR strategies coincide.

Lemma 2. *For every* **CGS** *M, state q in M, and $\boldsymbol{ATL}^*$ formula φ, $M, q \models_{\mathbf{IR}} \varphi$ iff $T(M,q), q \models_{\mathbf{IR}} \varphi$.*

Proof (sketch). Induction over the structure of φ. The main case is again $\varphi \equiv \langle\!\langle A \rangle\!\rangle \gamma$ for which it is sufficient to observe that (i) IR strategies in M, q uniquely correspond to IR strategies in $T(M,q), q$; (ii) $out_M(q, s_A) = out_{T(M,q)}(q, s_A)$ for every IR strategy s_A.

Proposition 6. *$Val(\boldsymbol{ATL}^*_{\mathrm{Ir}}) \subseteq Val(\boldsymbol{ATL}^*_{\mathrm{IR}})$*

Proof. Let an **ATL*** formula φ be Ir-valid in **CGS**'s, then it is also Ir-valid in tree-like **CGS**'s, and by Lemma 1 also IR-valid in tree-like **CGS**'s. Thus, by Lemma 2, it is IR-valid in arbitrary **CGS**'s.

In particular, the subsumption holds for formulae of **ATL**$^+$. Moreover:

Proposition 7. *$Val(\boldsymbol{ATL}^+_{\mathrm{IR}}) \nsubseteq Val(\boldsymbol{ATL}^+_{\mathrm{Ir}})$.*

Proof. Consider formula

$$\Phi_3 \equiv \langle\langle a\rangle\rangle(\Diamond \mathsf{p}_1 \wedge \Diamond \mathsf{p}_2) \leftrightarrow \langle\langle a\rangle\rangle\Diamond(\mathsf{p}_1 \wedge \langle\langle a\rangle\rangle\Diamond \mathsf{p}_2 \vee \mathsf{p}_2 \wedge \langle\langle a\rangle\rangle\Diamond \mathsf{p}_1).$$

The formula is valid in $\mathbf{ATL}^+_{\mathrm{IR}}$ [5]. On the other hand its right-to-left part is not valid in $\mathbf{ATL}^+_{\mathrm{Ir}}$. To see this, we take the single-agent **CGS** M_3 from Figure 4 where agent a (the robot) can either do the cleaning or the delivery of a package. Then, for $\mathsf{p}_1 \equiv \mathsf{clean}, \mathsf{p}_2 \equiv \mathsf{delivered}$, we have $M_3, q_0 \models_{\mathbf{Ir}} \langle\langle a\rangle\rangle\Diamond(\mathsf{p}_1 \wedge \langle\langle a\rangle\rangle\Diamond \mathsf{p}_2 \vee \mathsf{p}_2 \wedge \langle\langle a\rangle\rangle\Diamond \mathsf{p}_1)$ but also $M_3, q_0 \not\models_{\mathbf{Ir}} \langle\langle a\rangle\rangle(\Diamond \mathsf{p}_1 \wedge \Diamond \mathsf{p}_2)$.

Corollary 4. *$Val(\boldsymbol{ATL}^+_{\mathrm{Ir}}) \subsetneq Val(\boldsymbol{ATL}^+_{\mathrm{IR}})$.*

Comparing $\mathbf{ATL}_{\mathbf{i_o r}}$ vs. $\mathbf{ATL}_{\mathbf{i_o R}}$ Now we compare the memoryless and perfect recall semantics under uncertainty. The basic idea is similar to the one behind Proposition 6. First, we define *tree-like **i**CGS's* analogously to Definition 1. Tree unfoldings of **iCGS**'s are also similar – this time, however, we must take into account the indistinguishability relations.

Definition 3 ($\mathrm{i_o}$R-tree unfolding). *Given an **i**CGS M and a state q in it, we define the* $\mathrm{i_o}$R-tree unfolding *of M, q, denoted $T_{\mathrm{i_o}R}(M,q)$ as $T(M,q)$ from Definition 2, and add epistemic relations $\sim_a^{T_{\mathrm{i_o}R}(M,q)}$ reflecting indistinguishability of histories in M: $h \sim_a^{T_{\mathrm{i_o}R}(M,q)} h'$ iff $h \approx_a^M h'$.*

Let $T = T_{\mathrm{i_o R}}(M,q)$, and h a node in it. We define the *h-subtree of T* as the tree-like **iCGS** obtained from T by removing all nodes except for h and its descendants. We have the following obvious lemma.

Lemma 3. *Let T_h be the h-subtree of $T := T_{\mathrm{i_o}R}(M,q_0)$, where $h = q_0 q_1 \dots q_n$. Then, for every node h' in T_h and every $a \in \mathbb{A}\mathrm{gt}$ if we have $h' \sim_a^{T_h} h''$ then $hh' \sim_a^T hh''$.*

In other words, the information sets in a game can only be more precise when the game already follows some previous interaction. Analogously to Lemma 1 we have the following result.

Lemma 4. *For every tree-like **i**CGS M, state q in M, and $\boldsymbol{ATL}^*$ formula φ, we have that $M, q \models_{\mathbf{i_o r}} \varphi$ iff $M, q \models_{\mathbf{i_o R}} \varphi$.*

Lemma 5. *For every node h in $T_{\mathrm{i_o}R}(M,q_0)$ it holds that $T_{\mathrm{i_o}R}(M,q_0), h \models_{\mathrm{i_o}R} \varphi$ iff $M, last(h) \models_{\mathrm{i_o}R} \varphi$.*

Proof. The proof follows by induction on the structure of φ. Here, we only show it for $\varphi \equiv \langle\langle A\rangle\rangle\Box\psi$; the other cases are either analogous or straightforward. Let $T := T_{\mathrm{i_o R}}(M,q_0)$.

"$\Rightarrow$": We have $T, h \models_{i_oR} \langle\langle A\rangle\rangle\Box\psi$ iff $T, h \models_{i_or} \langle\langle A\rangle\rangle\Box\psi$ by Lemma 4. So there is an ir-strategy s_A such that $(\star)$ $\forall\lambda \in out_T(h, s_A)\forall j.T, \lambda[0,j] \models \psi$.

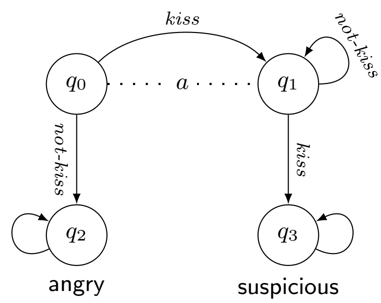

Fig. 5. Model M_4: dangers of marital life

We construct a witnessing perfect recall strategy s'_A in M as follows: $s'_a(\hat{h}) = s_a(hh')$ for every $a \in A$ and $\hat{h}$ such that $last(h)h' \sim_a^M \hat{h}$. We define s'_a arbitrarily for all other histories with the condition to assign the same actions to indistinguishable histories.

Strategy s'_A is uniform by construction. Moreover, by ($\star$), we have $\forall\lambda \in out_M(last(h), s'_A)\forall j.T, h\lambda[1..j] \models_{i_oR} \psi$ (where $h\lambda[1..j]$ is the history obtained by combining h and $\lambda[0..j]$, and extracting state $\lambda[0]$). Hence, by induction hypothesis: $\forall\lambda \in out_M(last(h), s'_A)\forall j.M, last(h\lambda[1..j]) \models_{i_oR} \psi$. Thus, we have $M, last(h) \models_{i_oR} \langle\langle A\rangle\rangle\Box\psi$.

"$\Leftarrow$": We have $M, last(h) \models_{i_oR} \langle\langle A\rangle\rangle\Box\psi$, so there is an iR-strategy s_A such that ($\star\star$) $\forall\lambda \in out_M(last(h), s_A)\forall i.M, \lambda[i] \models_{i_oR} \psi$. We construct the memoryless strategy s'_A in T as follows: $s'_a(\hat{h}h') = s_a(last(h)h')$ for every $a \in A$ and $\hat{h}$ such that $h \sim_a^T \hat{h}$. For all other histories h' we define $s'_a(h)$ arbitrarily but in a uniform way. The strategy s'_A is uniform by Lemma 3 and by construction. By ($\star\star$), $\forall\lambda \in out_T(h, s'_A)\forall i.M, last(\lambda[i]) \models_{\mathrm{i_oR}} \psi$. So, by induction hypothesis: $\forall\lambda \in out_T(h, s'_A)\forall i.T, \lambda[i] \models \psi$, and as a consequence $T, h \models_{\mathrm{i_oR}} \langle\langle A\rangle\rangle\Box\psi$.

Proposition 8. *$Val(\boldsymbol{ATL}_{\mathrm{i_or}}) \subseteq Val(\boldsymbol{ATL}_{\mathrm{i_oR}})$.*

Proof. We prove that $Sat(\mathbf{ATL}_{\mathrm{i_oR}}) \subseteq Sat(\mathbf{ATL}_{\mathrm{i_or}})$. Let $\varphi \in Sat(\mathbf{ATL}_{\mathrm{i_oR}})$. Then, there must be a pointed **iCGS** M, q such that $M, q \models_{\mathrm{i_oR}} \varphi$. By Lemma 5, $T_{\mathrm{i_oR}}(M, q), q \models_{\mathrm{i_oR}} \varphi$. But on iR-tree unfoldings, iR- and ir-strategies coincide (Lemma 4), so we get that $T_{\mathrm{i_oR}}(M, q), q \models_{\mathrm{i_or}} \varphi$, and as a consequence $\varphi \in Sat(\mathbf{ATL}_{\mathrm{i_or}})$.

The converse does not hold:

Proposition 9. *$Val(\boldsymbol{ATL}_{\mathrm{i_oR}}) \not\subseteq Val(\boldsymbol{ATL}_{\mathrm{i_or}})$*

Proof. To show this, we take the **ATL** embedding of the **CTL** duality between combinators $\mathsf{E}\Box$ and $\mathsf{A}\Diamond$. In fact, only one direction of the equivalence is important here: $\Phi_4 \equiv \neg\langle\langle\emptyset\rangle\rangle\Diamond\neg\mathsf{p} \to \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Box\mathsf{p}$ (note that the other direction is valid for all the semantics considered in this paper, and actually all the reasonable semantics of strategic ability that one can come up with).

First, we observe that: (i) $\neg\langle\langle\emptyset\rangle\rangle\Diamond\neg\mathsf{p}$ expresses (regardless of the actual type of ability being considered) that there is a path in the system on which p always holds; (ii) in the "objective" semantics the set $out(q, s_{\mathbb{A}\mathrm{gt}})$ always consists of exactly one path; (iii) for every path λ starting from q, there is an $\mathrm{i_oR}$-strategy $s_{\mathbb{A}\mathrm{gt}}$ such that $out(q, s_{\mathbb{A}\mathrm{gt}}) = \{\lambda\}$. From these, it is easy to see that Φ_4 is valid in $\mathbf{ATL}_{\mathrm{i_oR}}$.

Second, we consider model M_4 in Figure 5.[4] Let us take $\mathsf{p} \equiv \neg\mathsf{angry} \wedge \neg\mathsf{suspicious}$. Then, we have $M, q_0 \models_{\mathrm{i_or}} \neg\langle\langle\emptyset\rangle\rangle\Diamond\neg\mathsf{p}$ but also $M, q_0 \not\models_{\mathrm{i_or}} \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Box\mathsf{p}$, which demonstrates that Φ_4 is not valid in $\mathbf{ATL}_{\mathrm{i_or}}$.

Corollary 5. *$Val(\boldsymbol{ATL}_{\mathrm{i_or}}) \subsetneq Val(\boldsymbol{ATL}_{\mathrm{i_oR}})$.*

Proposition 10. *$Val(\boldsymbol{ATL}_{\mathrm{i_sR}}) \not\subseteq Val(\boldsymbol{ATL}_{\mathrm{i_sr}})$.*

Proof. We take

$$\Phi_5 \equiv \langle\langle a\rangle\rangle \bigcirc \langle\langle a\rangle\rangle\Diamond\mathsf{p} \rightarrow \langle\langle a\rangle\rangle\Diamond\mathsf{p}.$$

The formula is stating that if a has an opening move and a follow-up strategy to achieve p eventually, then these can be integrated into a single strategy achieving p already from the initial state. It is easy to see that Φ_5 is valid in $\mathbf{ATL}_{\mathrm{i_sR}}$, and that the single strategy is just a concatenation of the two strategies that we get on the left hand side of the implication. On the other hand, for the "poor duck model" M_1 and $\mathsf{p} \equiv \mathsf{shot}$, we get that $M_4, q_4 \models_{\mathrm{i_sr}} \langle\langle a\rangle\rangle \bigcirc \langle\langle a\rangle\rangle\Diamond\mathsf{p}$ but also $M_4, q_4 \models_{\mathrm{i_sr}} \langle\langle a\rangle\rangle\Diamond\mathsf{p}$, so Φ_5 is not valid in $\mathbf{ATL}_{\mathrm{i_sr}}$.

We leave the other direction ($Val(\mathbf{ATL}_{\mathrm{i_sr}}) \subseteq Val(\mathbf{ATL}_{\mathrm{i_sR}})$) for future research.

3.3 Between Subjective and Objective Ability

Finally, we compare validity sets for the semantic variants of **ATL** that differ on the outcome paths which are taken into account, i.e., whether only the paths representing the "objectively" possible courses of action are considered, or all the executions that are "subjectively" possible from the agents' perspective.

Proposition 11. *$\Phi_2 \equiv \langle\langle a\rangle\rangle\Diamond\mathsf{p} \rightarrow \mathsf{p} \vee \langle\langle a\rangle\rangle \bigcirc \langle\langle a\rangle\rangle\Diamond\mathsf{p}$ is valid in $\boldsymbol{ATL}_{\mathrm{i_oR}}$ and $\boldsymbol{ATL}_{\mathrm{i_or}}$, but invalid in $\boldsymbol{ATL}_{\mathrm{i_sR}}$ and $\boldsymbol{ATL}_{\mathrm{i_sr}}$.*

Proof. We first prove validity of Φ_2 in $\mathbf{ATL}_{\mathrm{i_or}}$, which implies also validity in $\mathbf{ATL}_{\mathrm{i_oR}}$ by Proposition 8. Suppose that $M, q \models_{\mathrm{i_or}} \langle\langle a\rangle\rangle\Diamond\mathsf{p}$, then there must be an ir-strategy s_A that enforces $\Diamond\mathsf{p}$ for every execution starting from q. But then, if p is not the case right at the beginning, s_A must lead to a next state from which it enforces $\Diamond\mathsf{p}$.

For the second part, invalidity of Φ_2 in $\mathbf{ATL}_{\mathrm{i_sR}}$ was already proved in Proposition 5. Thus, by Proposition 10, Φ_2 is not valid in $\mathbf{ATL}_{\mathrm{i_sr}}$, too.

[4] The example depicts some simple traps that await a married man if he happens to be absent-minded. If he doesn't kiss his wife in the morning, he is likely to make her angry. However, if he kisses her more than once, she might get suspicious. It is easy to see that the absent-minded (i.e., memoryless) husband does not have a strategy to survive safely through the morning, though a "safe" path through the model does exist ($\lambda = q_0 q_1 q_1 \ldots$).

Ir	i_sR	i_oR	i_sr	i_or	
$=^{(*)}$	$\subsetneq$	$\subsetneq$	$\subsetneq$	$\subsetneq$	IR
	$\subsetneq$	$\subsetneq$	$\subsetneq$	$\subsetneq$	Ir
		$\bowtie$	$\not\supseteq^{(**)}$	$\bowtie$	i_sR
			$\bowtie$	$\subsetneq$	i_oR
				$\bowtie$	i_sr

Fig. 6. Summary: comparison of validity sets induced by various semantics of **ATL**. The $\bowtie$ symbol denotes incomparable sets. (*) For $\mathbf{ATL}^+$ and $\mathbf{ATL}^*$: Ir $\subsetneq$ IR. (**) The other direction is still to be established.

Proposition 12. *Let us define an additional operator* N *("now") as* $\mathsf{N}\varphi \equiv \varphi\,\mathcal{U}\,\varphi$. *Formula* $\Phi_6 \equiv \langle\langle a\rangle\rangle \mathsf{N} \langle\langle c\rangle\rangle \bigcirc \langle\langle a\rangle\rangle \bigcirc \mathsf{p} \rightarrow \langle\langle a, c\rangle\rangle \Diamond \mathsf{p}$ *is valid in* $\boldsymbol{ATL}_{\mathrm{i_sR}}$ *and* $\boldsymbol{ATL}_{\mathrm{i_sr}}$*, but invalid in* $\boldsymbol{ATL}_{\mathrm{i_oR}}$ *and* $\boldsymbol{ATL}_{\mathrm{i_or}}$*.*

Proof. Analogously to the proof of Proposition 11, we will prove the validity of Φ_6 in $\mathbf{ATL}_{\mathrm{i_sr}}$, and its invalidity in $\mathbf{ATL}_{\mathrm{i_oR}}$.

First, let $M, q \models_{\mathrm{i_sr}} \langle\langle a\rangle\rangle \mathsf{N} \langle\langle c\rangle\rangle \bigcirc \langle\langle a\rangle\rangle \bigcirc \mathsf{p}$. Then, for every state $q' \in [q]_{\sim_a}$, c has an action α_c that enforces $\langle\langle a\rangle\rangle \bigcirc \mathsf{p}$ from $[q']_{\sim_c}$. By collecting these actions into an ir-strategy s_c (we can do it since single actions are successful for whole indistinguishability classes of c), we have that s_c enforces $\bigcirc \langle\langle a\rangle\rangle \bigcirc \mathsf{p}$ from every state in $[q]_{\sim_{\{a,c\}}}$, regardless of what the other players do (in particular, regardless of what a does). But then, for every execution λ of s_c from $[q]_{\sim_{\{a,c\}}}$, a will have a choice to enforce $\bigcirc \mathsf{p}$ from $\lambda[1]$. Again, collecting these choices together yields an ir-strategy s_a (we can fix the remaining choices arbitrarily). By taking $s_{\{a,c\}} = (s_a, s_c)$, we get a strategy for $\{a, c\}$ that enforces that p will be true in two steps, from every state in $[q]_{\sim_{\{a,c\}}}$. Hence, also $M, q \models_{\mathrm{i_sr}} \langle\langle a, c\rangle\rangle \Diamond \mathsf{p}$.

For the invalidity, consider the "modified poor duck model" M_2 augmented with additional agent c that has no choice (i.e., at each state, it has only a single irrelevant action *wait* available). Let us denote the new **iCGS** by M_2'. If we identify p with shot, it is easy to see that $M_2', q_0' \models_{\mathrm{i_oR}} \langle\langle c\rangle\rangle \bigcirc \langle\langle a\rangle\rangle \bigcirc \mathsf{shot}$, and hence also $M_2', q_0' \models_{\mathrm{i_oR}} \langle\langle a\rangle\rangle \mathsf{N} \langle\langle c\rangle\rangle \bigcirc \langle\langle a\rangle\rangle \bigcirc \mathsf{shot}$. On the other hand, $M_2', q_0' \not\models_{\mathrm{i_oR}} \langle\langle a, c\rangle\rangle \Diamond \mathsf{shot}$, which concludes the proof.

Corollary 6. *For every* $\mathrm{y}, \mathrm{z} \in \{\mathrm{R}, \mathrm{r}\}$*, sets* $Val(\boldsymbol{ATL}_{\mathrm{i_sy}})$ *and* $Val(\boldsymbol{ATL}_{\mathrm{i_oz}})$ *are incomparable.*

4 Summary and Conclusions

In this paper, we compare validity sets for different semantic variants of alternating-time temporal logic. In other words, we compare the general properties of games induced by different notions of ability. It is clear that changing the notions of strategy and success in a game leads to a different game. The issue considered here is whether, given a *class* of games, such a change leads to a different class of games, too. And, if so, what is the precise relationship between the two classes.

A summary of the results is presented in Figure 6. The first, and most important, conclusion is that all the semantic variants discussed here are *different* on the level of general properties they induce; before our study, it was by no means obvious. Moreover, our results suggest a strong pattern: perfect information is a special case of imperfect information, perfect recall games are special case of memoryless games, and properties of objective and subjective abilities of agents are incomparable.

Acknowledgements. Wojciech Jamroga acknowledges the support of the FNR (National Research Fund) Luxembourg under project S-GAMES – C08/IS/03. The authors would like to express their thanks to Jan Broersen for the argument that triggered this study, and to Valentin Goranko for discussions on (non-)validity of fixpoint properties in various semantics of **ATL**.

References

1. T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In D. Samet, editor, *Proceedings of TARK XI*, pages 15–24, 2007.
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
3. V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004.
4. V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, 353(1):93–117, 2006.
5. A. Harding, M. Ryan, and P.-Y. Schobbens. Approximating ATL* in ATL. In *VMCAI '02: Revised Papers from the Third International Workshop on Verification, Model Checking, and Abstract Interpretation*, pages 289–301. Springer-Verlag, 2002.
6. W. Jamroga and T. Ågotnes. Constructive knowledge: What agents can achieve under incomplete information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.
7. W. Jamroga and J. Dix. Model checking abilities of agents: A closer look. *Theory of Computing Systems*, 42(3):366–410, 2008.
8. F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. *LMCS*, 4:7, 2008.
9. Sven Schewe. ATL* satisfiability is 2ExpTime-complete. In *Proceedings of ICALP 2008*, volume 5126 of *Lecture Notes in Computer Science*, pages 373–385. Springer-Verlag, 2008.
10. P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2), 2004.
11. W. van der Hoek, A. Lomuscio, and M. Wooldridge. On the complexity of practical ATL model checking. In P. Stone and G. Weiss, editors, *Proceedings of AAMAS'06*, pages 201–208, 2006.
12. D. Walther, C. Lutz, F. Wolter, and M. Wooldridge. ATL satisfiability is indeed EXPTIME-complete. *Journal of Logic and Computation*, 16(6):765–787, 2006.