Journal of Autonomous Agents and Multi-Agent Systems manuscript No.
(will be inserted by the editor)

Comparing Variants of Strategic Ability

How Uncertainty and Memory Influence General Properties of Games

Nils Bulling · Wojciech Jamroga

Received: date / Accepted: date

Abstract Alternating-time temporal logic (**ATL**) is a modal logic that allows to reason about agents' abilities in game-like scenarios. Semantic variants of **ATL** are usually built upon different assumptions about the kind of game that is played, including capabilities of agents (perfect vs. imperfect information, perfect vs. imperfect memory, etc.). **ATL** has been studied extensively in previous years; however, most of the research focused on model checking. Studies of other decision problems (e.g., satisfiability) and formal meta-properties of the logic (like axiomatization or expressivity) have been relatively scarce, and mostly limited to the basic variant of **ATL** where agents possess perfect information and perfect memory. In particular, a comparison between different semantic variants of the logic is largely left untouched.

In this paper, we show that different semantics of ability in **ATL** 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 strategic logics. Secondly, we show that different concepts of ability induce different general properties of games. Thirdly, the study can be seen as the first systematic step towards satisfiability-checking algorithms for **ATL** with imperfect information. We introduce sophisticated unfoldings of models and prove invariance results that are an important technical contribution to formal analysis of strategic logics.

Keywords Alternating-time temporal logic · Validity and satisfiability · Properties of games · Games with imperfect information

Nils Bulling
Department of Informatics,
Clausthal University of Technology, Germany
bulling@in.tu-clausthal.de

Wojciech Jamroga
Computer Science and Communication Research Unit,
and Interdisciplinary Centre on Security, Reliability and Trust
University of Luxembourg
wojtek.jamroga@uni.lu

1 Introduction

Alternating-time temporal logic (**ATL**) [7,9] is a temporal logic that incorporates some basic game theoretic notions. In **ATL** we can for instance express that a group of agents is able to *bring about* φ, i.e., the agents in the group are able to enforce that property φ holds whatever the other agents might do. 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 have a way of enforcing 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 game model. 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. In this paper, we settle the issue and show that most "classical" semantic variants of **ATL** are indeed different, and we characterize the relationship between their sets of validities.

The question is important for several reasons. First, many logicians identify a logic with the set of sentences that are true in the logic; a 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. Secondly, 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. All the variants studied here are defined over the same class of models (*imperfect information concurrent game structures*) that generalizes extensive form games. The difference between games "induced" by different semantics lies in the available strategies and the winning conditions for them.

Thirdly, 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 the logical analogue of game solving, satisfiability corresponds naturally to mechanism design. A systematic study on the abstract level is the first step towards algorithms that solve the problem.

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. Ultimately, we show that what agents can achieve is 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 temporal evolution of systems – may cease to be universally true if we change the strategic setting. 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 former kind of properties is especially important for practical purposes, since fixpoint equivalences provide the basis for most model checking and satisfiability checking algorithms. Finally, we introduce sophisticated unfoldings of models to show invariance results with respect to memoryless and perfect recall strategies. The unfoldings form an important technical contribution of this article. We believe that their impact goes beyond **ATL**, as they can probably be applied to other strategy logics. For example, it would be interesting to see which unfoldings preserve the truth values of formulae when imperfect information is combined with strategic commitment [66], or when explicit quantification over strategies is allowed [20,51,50].

The paper is structured as follows. We begin by presenting the relevant variants of **ATL** in Sections 2 and 3. Then we define several unfoldings of **ATL** models, and show that they preserve truth of **ATL** formulae under appropriate assumptions in Section 4. This is the most technical part of the paper, and readers interested only in the main *conceptual* contribution are advised to skip it and proceed to the next section. In Section 5, we show the formal relationships between validity sets for different variants of **ATL**. Summary of the main results and some conclusions are presented in Section 6.

About this article. Preliminary versions of this paper appeared in [40,41]. The journal version adds proofs, new results, examples, and more extensive discussions. This applies in particular to Section 4 where we stress the importance of tree-like unfoldings and provide a sophisticated construction as well as full proofs. We have also extended the results from [40,41] (formulated mainly for the restricted language of **ATL**) to the more general language of **ATL***.

1.1 Related Work

ATL has been studied extensively in the last 15 years. The research can be roughly divided into the computational and conceptual strands. The former has been focused on the way in which **ATL** and its extensions can be used for verification of multi-agent systems, in particular what is the complexity of model checking, and how one can overcome the inherent difficulties. An interested reader is referred to [13] for an overview, and to [9,57,59,42,48,17,24,20] for more specific results; some attempts at taming the complexity were proposed e.g. in [46,38,23,47,18]. Studies on other decision problems than model checking were much less frequent, though satisfiability of the basic variant of **ATL** has been investigated in [31,64,54,30].

The conceptual strand originally emerged in quest of the “right” semantics for strategies under uncertainty. **ATL** was combined with epistemic logic [60,61,1,62, 2,39], and several semantic variants were defined that match various possible interpretations of ability [35,45,57,43,39]. Moreover, many conceptual extensions have been considered, e.g., with explicit reasoning about strategies, rationality assumptions and solution concepts [63,58,65,19,20], agents with bounded resources [5,14, 6,15], coalition formation and negotiation [12], opponent modeling and action in stochastic environments [37,16,56,55] and reasoning about irrevocable plans and interplay between strategies of different agents [3,11].

In the rich literature on the conceptual virtues of alternating-time temporal logic, formal analysis is relatively scarce. Axiomatization of the basic variant of **ATL** was proposed in [31], and its expressivity was addressed in [9,48]. Axiomatization of a particular variant of imperfect information was proposed in [32]. For comparative studies, invariance of the basic semantics with respect to a couple of classes of models was proven in [27], and the correspondence between abstract and concrete models of strategic logics was the object of study in [52,29,28]. Surprisingly, relationships between the "classical" semantic variants of **ATL** (as defined e.g. in [57]) have not yet been studied, though analogous results exist that compare more sophisticated variations to a more standard variant (cf. [36] for undominated play, [3] for irrevocable strategies, [4] for agents with bounded memory, and [18] for recomputable strategies under uncertainty). That means in particular that formal properties of strategic ability under imperfect information are largely left untouched. We are trying to fill in the gap, and start a more systematic charting of the landscape.

2 Reasoning about Strategic Abilities

Alternating-time temporal logic **ATL** [7,9] is a temporal logic that incorporates some basic game-theoretical notions. Essentially, **ATL** generalizes the branching time logic **CTL** [21] by replacing path quantifiers E, A with *strategic 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"), "$\Diamond$" ("now or sometime in the future"), and $\mathcal{U}$ ("until"). Since **ATL** offers no way of representing agents' uncertainty in its models, and no operators to refer to agents' (lack of) knowledge in the object language, it allows to reason only about abilities of agents with perfect information about the current global state of the system.

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.

Definition 1 (Language of ATL*) 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:

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

The "sometime" and "always" operators can be defined as $\Diamond\gamma \equiv \top\,\mathcal{U}\,\gamma$ and $\Box\gamma \equiv \neg\Diamond\neg\gamma$.

Formulae φ are called *state formulae*, and γ *path formulae* of **ATL***. A path formula is *simple* if it consists of a temporal operator followed immediately by a state subformula and in the case of "until" the operator is also immediately preceded by a state subformula. In other words, temporal operators have to be applied to state subformulae.

The best known syntactic variant of alternating time temporal logic is **ATL** in which every occurrence of a strategic modality is immediately coupled with a temporal operator, i.e., we have coupled operators of the form $\langle\!\langle A\rangle\!\rangle\bigcirc$, $\langle\!\langle A\rangle\!\rangle\Box$, and $\langle\!\langle A\rangle\!\rangle\,\mathcal{U}$. The language of $\mathbf{ATL}^{+}$ sits between $\mathbf{ATL}^{*}$ and **ATL**: it allows strategic modalities to be followed by *a Boolean combination* of simple temporal subformulae.

Formally, formulae of **ATL** are defined be the following grammar:[1]

$$\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi\wedge\varphi \mid \langle\!\langle A\rangle\!\rangle\bigcirc\varphi \mid \langle\!\langle A\rangle\!\rangle\Box\varphi \mid \langle\!\langle A\rangle\!\rangle\varphi\,\mathcal{U}\,\varphi$$

and $\mathbf{ATL}^{+}$ formulae by:

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

Example 1 The **ATL** formula $\langle\!\langle jamesbond, octopussy\rangle\!\rangle\Diamond\mathsf{kiss}$ says that James Bond and Octopussy can eventually kiss, no matter how the other agents act. On the other hand, $\langle\!\langle jamesbond, jaws\rangle\!\rangle(\Box\neg\mathsf{crash}\wedge\Diamond\mathsf{land})$ (James Bond and Jaws can prevent the space ship from crashing and make it eventually land) is a formula of $\mathbf{ATL}^{+}$ but not **ATL**. Finally, $\langle\!\langle jamesbond\rangle\!\rangle\Box\Diamond\mathsf{deadBlofeld}\wedge\neg\langle\!\langle jamesbond\rangle\!\rangle\Diamond\Box\mathsf{deadBlofeld}$ is an $\mathbf{ATL}^{*}$ formula (which clearly belongs to neither **ATL** nor $\mathbf{ATL}^{+}$) which states that agent 007 can kill Ernst Stavro Blofeld infinitely many times, but he cannot kill Blofeld once and forever.

2.2 Basic Models of ATL

In [9], 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.

Definition 2 (Concurrent game structure) A *concurrent game structure*[2] (**CGS**) is a tuple $\mathfrak{M} = \langle\mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o\rangle$ which includes a nonempty finite set of all agents $\mathbb{A}\mathrm{gt} = \{1,\dots,k\}$, a nonempty (possibly infinite) set of states St, a set of atomic propositions Π and their valuation $\pi : \Pi \to 2^{St}$, and a nonempty 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 $\alpha_i \in d(i,q)$ that can be executed by $\mathbb{A}\mathrm{gt}$ in q.

Thus, we assume that all the agents execute their actions synchronously; the combination of the actions, together with the current state, determines the next transition of the system.

In the rest of the paper, we will write $d_i(q)$ instead of $d(i,q)$, and we will denote the set of collective choice of group A at state q by $d_A(q) = \prod_{i\in A} d_i(q)$.

We will sometimes use the term *pointed* **CGS** for a pair $(\mathfrak{M}, q)$ of a concurrent game structure and a state in it.

[1] Note that "always" is not definable from "until" in **ATL** [48], and has to be added explicitly to the language.

[2] We would like to note that it is essential for this work that we do not require a finite set of states or actions. We give a more detailed discussion in Section 2.3.

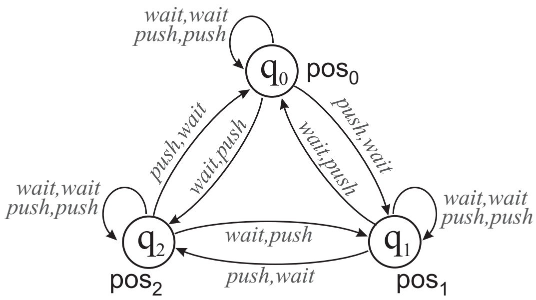

Fig. 1 Two robots and a carriage: concurrent game structure M_0

Definition 3 (Path) 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 $\Lambda_{\mathfrak{M}}(q)$.

Example 2 (Robots and Carriage) Consider the scenario depicted in Figure 1. 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.

To make our model of the domain discrete, we identify 3 different positions of the carriage, and associate them with states q_0, q_1, and q_2. We label the states with propositions $\mathsf{pos}_0, \mathsf{pos}_1, \mathsf{pos}_2$, respectively, to allow for referring to the current position of the carriage in the object language.

2.3 Finite vs. Infinite CGS

In our definition of **CGS** (Def. 2.2) we have not put up any requirement of finiteness with respect to the set of states and actions. The only requirement is that the set of agents must be finite. In particular, we allow for infinitely many states in a model; we also allow for infinitely branching models. In this section we shall discuss this choice in more detail.

We begin by reviewing the literature and showing that both types of **CGS** – finite and infinite ones – have been considered by other authors. The semantics of **ATL** in concurrent game models was originally proposed for finite structures only [9].[3] Many follow-up papers also adopted the assumption of finite models, for example [43, 57]

[3] An interested reader may observe that the preliminary versions of the semantics (in alternating transition systems) did not assume models explicitly to be finite [7,8]. However, the authors *de facto* considered only finite models since they were solely interested in the model checking problem, where the input must be finite.

that studied variants of **ATL** with imperfect information, [11] which extended **ATL** with persistent strategies, in [48] the expressive power of **ATL** is investigated, etc.

On the other hand, other authors did not restrict their analysis to the finite case, beginning with the work on coalition logic [52,53], through comparative studies of different semantics of **ATL** [26,27], the interplay between knowledge and strategies [2,39], strategic play in the presence of intentions and commitment [44,3], and so on. Also, different formalisms extending **ATL*** with explicit quantification over strategies follow different assumptions: on one hand, the strategy logic by Chatterjee et al. [20] assumes models to be finite; on the other, the strategy logic recently proposed by Mogavero et al. [51,51] only requires states and actions to be countable.

As we have already stated, we assume neither St nor Act to be finite (or even enumerable). *How does that affect our work?* First of all, for the new results in this article, it is especially important that some existing technical results can be applied to infinite models. This concerns in particular the axiomatization of **ATL** from [31] which was shown sound and complete for finite as well as infinite concurrent game models. To be more precise, the authors of [31] allow for infinitely many states, but assume that, at any state, there are only finally many outgoing transitions. However, their results extend to the case of infinite branching in a straightforward way. We use the axiomatization as a source of "standard" validities (like the fixpoint characterization for $\langle\langle A\rangle\rangle\Diamond$), and to show that the semantics of "perfect information memoryless **ATL**" and "perfect information perfect recall **ATL**" coincide also for infinite models (Proposition 1). Moreover, the notion of model equivalence for **ATL** (alternating bisimulation alias strategic bisimulation), while originally proposed for finite models only [10], was extended to the unrestricted case and proved correct in [3]. We use and extend the concept to prove invariance results for tree-like unfoldings in Section 4.

Secondly, *all the results proposed in this paper are proved to hold if the semantics of* ***ATL*** *and* ***ATL**** *does not restrict the class of models to finite ones*. More precisely, it may be possible that one of our inclusion results between the validity sets of two logics, $Val(\mathsf{L}_1) \subseteq Val(\mathsf{L}_2)$, requires the existence of an infinite model. This does *not* mean that the theorems that we present do *not* hold in the class of finite models. The latter issue, albeit interesting, is outside of the scope of the paper. Essentially, showing that our results hold in the finite semantics would require establishing finite model properties for the logics that we consider. To the best of our knowledge, such properties have only been proven for the "perfect information/perfect recall" variant of **ATL** [31] and **ATL*** [54]. Proving (or disproving) the finite model property for the other variants of **ATL**/**ATL*** is undoubtedly important, and we would like to study it further in the future.[4]

In summary:

1. Our inclusion results rely on the fact that we define the semantics of **ATL** and **ATL*** in both finite and infinite models; and
2. whenever a finite model property holds for two logics under consideration, our results comparing the two logics apply also when the semantics is restricted to finite models.

[4] We thank an anonymous JAAMAS reviewer for suggesting this.

2.4 Strategies and Abilities in Basic Semantics of **ATL**

ATL modalities refer to the outcome of strategic play for a given coalition. Following the tradition of extensive form games in game theory, a *strategy* of agent a is understood as a plan that specifies what a is going to do in each situation. In the standard version of **ATL** [7,9], strategies are represented by functions $s_a : St^+ \to Act$. Thus, it is implicitly assumed that agents have *perfect information* (at each moment, they can precisely recognize the current global state of the system) and *perfect recall* (they can base their choices on the whole history of the game so far, not just the last state). Alternatively, one can assume that agents have no memory beyond what is already contained in the current state, which gives rise to the notion of *memoryless* (or *positional*) strategy. As we explain more systematically in Section 3, we will use the notation from [57] where *i* (resp. *I*) stands for *imperfect* (resp. *perfect*) *information*, and *r* (resp. *R*) for *imperfect* (resp. *perfect*) *recall*.

Definition 4 (*IR*- and *Ir*-strategies) Let $\mathfrak{M}$ be a $\mathbf{CGS}$ over states St and actions Act. A *perfect information perfect recall strategy* for agent a in $\mathfrak{M}$ (*IR-strategy* in short) is a function $s_a : St^+ \to Act$ such that $s_a(q_0q_1 \dots q_n) \in d_a(q_n)$. The set of such strategies is denoted by Σ_a^{IR}.

A *perfect information memoryless strategy* for agent a in $\mathfrak{M}$ (*Ir-strategy* in short) is a function $s_a : St \to Act$ where $s_a(q) \in d_a(q)$. The set of such strategies is denoted by Σ_a^{Ir}.

A *collective strategy* for a group of agents $A = \{a_1, \dots, a_r\}$ is simply a tuple of individual strategies $s_A = \langle s_{a_1}, \dots, s_{a_r} \rangle$. The set of such strategies is denoted by Σ_A^{IR} (for *IR* strategies) and Σ_A^{Ir} (for *Ir* strategies, respectively).[5]

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. Let $a \in A$; by $s_A|_a$, we denote agent a's part s_a of the collective strategy s_A.

Definition 5 (Outcome) The *outcome* $out_{\mathfrak{M}}(q, s_A)$ *of an IR-strategy* s_A *from state* q *in model* $\mathfrak{M}$ is the set of all paths $\lambda = q_0q_1q_2 \dots$ such that $q_0 = q$ and for each $i = 1, 2, \dots$ there exists a tuple of agents' decisions $\langle \alpha_{a_1}^{i-1}, \dots, \alpha_{a_k}^{i-1} \rangle$ such that: *(i)* $\alpha_a^{i-1} \in d_a(q_{i-1})$ for every $a \in \mathbb{A}\mathrm{gt}$, *(ii)* $\alpha_a^{i-1} = s_A|_a(q_0q_1 \dots q_{i-1})$ for every $a \in A$, and *(iii)* $o(q_{i-1}, \alpha_{a_1}^{i-1}, \dots, \alpha_{a_k}^{i-1}) = q_i$.

The *outcome* $out_{\mathfrak{M}}(q, s_A)$ of an *Ir*-strategy s_A is defined analogously but $s_A|_a(q_0q_1 \dots q_{i-1})$ is replaced by $s_A|_a(q_{i-1})$.

Often, we will omit the subscript "$\mathfrak{M}$" if it is clear from the context.

Let $\mathfrak{M}$ be a $\mathbf{CGS}$, q a state, and λ a path in $\mathfrak{M}$. Now, the semantics of **ATL*** and its sublanguages can be defined by the following clauses [9]:

$\mathfrak{M}, q \models \mathsf{p}$ iff $q \in \pi(\mathsf{p})$, for $\mathsf{p} \in \Pi$;
$\mathfrak{M}, q \models \neg\varphi$ iff $\mathfrak{M}, q \not\models \varphi$;

[5] As commonly done we assume an implicit order on the elements in $\mathbb{A}\mathrm{gt}$ allowing to conveniently represent collective strategies as tuples. In our setting where agents are represented by natural numbers the order is apparent. In the general case, a collective strategy for A is a function that associates individual strategies to the agents in A.

$\mathfrak{M}, q \models \varphi_1 \wedge \varphi_2$ iff $\mathfrak{M}, q \models \varphi_1$ and $\mathfrak{M}, q \models \varphi_2$;

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

$\mathfrak{M}, \lambda \models \varphi$ iff $\mathfrak{M}, \lambda[0] \models \varphi$;
$\mathfrak{M}, \lambda \models \neg\gamma$ iff $\mathfrak{M}, \lambda \not\models \gamma$;
$\mathfrak{M}, \lambda \models \gamma_1 \wedge \gamma_2$ iff $\mathfrak{M}, \lambda \models \gamma_1$ and $\mathfrak{M}, \lambda \models \gamma_2$;
$\mathfrak{M}, \lambda \models \bigcirc \gamma$ iff $\mathfrak{M}, \lambda[1, \infty] \models \gamma$; and
$\mathfrak{M}, \lambda \models \gamma_1 \,\mathcal{U}\, \gamma_2$ iff there is an $i \in \mathbb{N}_0$ such that $\mathfrak{M}, \lambda[i, \infty] \models \gamma_2$ and $\mathfrak{M}, \lambda[j, \infty] \models \gamma_1$ for all $0 \leq j < i$.

Example 3 (Robots and Carriage, ctd.) Since the outcome of each robot's action depends on the current action of the other robot, no robot 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). Moreover, single robots can play strategically to *avoid* the carriage entering a particular position: $M_0, q_0 \models \langle\!\langle 1 \rangle\!\rangle \Box \neg \mathsf{pos}_1$ (the strategy: wait in q_0 and push in q_2).

Note that the semantics does not address the issue of coordination between members of the coalition [25,34]: if there exist several successful joint strategies for A, the agents in A are assumed to somehow choose between them.

Finally, validity and satisfiability in **ATL** are defined in the standard way.

Definition 6 (Validity and satisfiability) Formula φ is *valid in model* $\mathfrak{M}$ iff it holds in every state of $\mathfrak{M}$, i.e., $\mathfrak{M}, q \models \varphi$ for every $q \in St_{\mathfrak{M}}$. The formula is valid in a class of models $\mathcal{C}$ iff it is valid in every model from $\mathcal{C}$.

Dually, φ is satisfiable in a class of models $\mathcal{C}$ iff there exists $\mathfrak{M} \in \mathcal{C}$ and a state q in $\mathfrak{M}$ such that $\mathfrak{M}, q \models \varphi$.

2.5 Some Important Validites

We recall that the following fixpoint properties are valid in the basic semantics of **ATL** presented in Section 2.4:

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

Validity of these formulae was demonstrated in [9] for finite models, and in [31] for finitely branching models (with possibly infinitely many states). It is easy to check that the argument extends to models with infinite branching.

The intuitive meaning of the fixpoint equivalences is that planning for a long-term goal (like achieving φ eventually, or maintaining φ persistently) can be decomposed into finding a good opening move and a suitable follow-up. Such properties are crucial for building model checking and satisfiability checking algorithms, and in particular they allow for incremental iterative synthesis of strategies.

Moreover, the path quantifiers A, E of $\mathbf{CTL}$ can be expressed in the standard semantics of **ATL** with $\langle\langle\emptyset\rangle\rangle, \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle$ respectively. Again, checking this is routine, even for models with infinitely many states and infinite branching. As a consequence, the $\mathbf{CTL}$ duality axioms can be rewritten in **ATL**, and become validities in the basic semantics:

$$\begin{aligned}
\neg\langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\bigcirc\varphi &\leftrightarrow \langle\langle\emptyset\rangle\rangle\bigcirc\neg\varphi\\
\neg\langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Diamond\varphi &\leftrightarrow \langle\langle\emptyset\rangle\rangle\Box\neg\varphi,\\
\neg\langle\langle\emptyset\rangle\rangle\Diamond\varphi &\leftrightarrow \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Box\neg\varphi.
\end{aligned}$$

We observe that all the properties presented in this subsection are *schemes*, rather than single formulae, and allow for uniform substitution. More precisely, φ can be replaced by any state formula of **ATL**, and the resulting formula is a validity of **ATL**. Moreover, φ in the duality axioms can be replaced by any state or path formula of **ATL***, and the resulting formula is a validity of **ATL***.

3 Semantic Variants: Uncertainty and Recall

As we already pointed out in Section 2.4, one can distinguish between two types of strategies: an agent may base its decision on the current state or on the whole history of states that have occurred. Also, the agent may have complete or incomplete knowledge about the current global state of the system throughout the game. A number of semantic variations have been proposed for **ATL**, e.g. [35,45,57,43,39,3,4]. In this paper, we study the "canonical" variants as proposed in [57]. There, a natural taxonomy of four strategy types was introduced and labeled as follows: *R* (resp. *r*) stands for perfect (resp. imperfect) *recall*, and *I* (resp. *i*) refers to perfect (resp. imperfect) *information*. The semantics of **ATL**, $\mathbf{ATL}^+$ and **ATL*** can be parameterized with the strategy type – yielding four different semantic variants of the logic, labeled accordingly ($\mathbf{ATL}_{IR}$, $\mathbf{ATL}_{Ir}$, $\mathbf{ATL}_{iR}$, $\mathbf{ATL}_{ir}$, etc.).

In this paper, we extend the taxonomy with a distinction between *objective* and *subjective* abilities under imperfect information, denoted by i_o and i_s, respectively; the distinction can be traced back to [35,45,43,39]. Intuitively, subjective ability to bring about γ means that the agents are able to both identify and execute the right strategy, i.e., they not only can play to achieve γ; they also *know* how to do it. Objective ability is a weaker property: the agents could execute the right strategy, but they do not necessarily know which one works out, and they might be even unaware that such a strategy exists. Examples of agents who have objective but not subjective ability to achieve their goals include: garbage collecting robots that execute a strategy (program) provided by the producer, a Master's student executing a strategy hinted by his/her supervisor, etc.

The distinction between perfect and imperfect recall (*R* vs. *r*) is reflected in the type of function used to represent strategies ($St^+ \to Act$ vs. $St \to Act$). The distinction between perfect and imperfect information (I vs. i) yields constraints on the set of functions that represent "feasible" strategies. The additional refinement of the imperfect information case (i_o vs. i_s) determines which outcome paths will be taken into account when evaluating the success of a strategy.

3.1 Imperfect Information Models and Strategies

Models, *imperfect information concurrent game structures* ($\mathbf{iCGS}$) [61,57], are 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. Note that $\mathbf{CGS}$'s can be seen as the subclass of $\mathbf{iCGS}$'s where all $\sim_a$ are the minimal reflexive relations. Formally, we have:

Definition 7 ($\mathbf{iCGS}$) An *imperfect information concurrent game structure* ($\mathbf{iCGS}$) is given by

$$\mathfrak{M} = \langle \mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o, \{\sim_a \mid a \in \mathbb{A}\mathrm{gt}\}\rangle$$

where $\langle \mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o\rangle$ is a $\mathbf{CGS}$, each $\sim_a \subseteq St \times St$ is an equivalence relation, and if $q \sim_a q'$ then $d(a, q) = d(a, q')$.

Definition 8 (Histories, $h \approx h'$, $last(h)$, $\circ$, $|h|$) A *history* is a finite sequence of states that can be effected by subsequent transitions. Two histories $h = q_0 q_1 \dots q_n$ and $h' = q'_0 q'_1 \dots q'_{n'}$ are *indistinguishable for agent* a ($h \approx_a h'$) iff $n = n'$ and $q_i \sim_a q'_i$ for $i = 0, \dots, n$. 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, and $|h|$ to denote the length of h (i.e., the number of states in h). As for paths, we use $h[i, j]$, $i \leq j$, $i < |h|$, to refer to the subhistory $h[i] \dots h[min(j, |h|-1)]$. We do also allow $j = \infty$. We define $\Lambda^{fin}_{\mathfrak{M}}(q)$ as the set of all histories in $\mathfrak{M}$ starting from q, i.e., all the finite prefixes of paths in $\Lambda_{\mathfrak{M}}(q)$. Moreover, $\Lambda^{fin}_{\mathfrak{M}} := \bigcup_{q \in St} \Lambda^{fin}_{\mathfrak{M}}(q)$ denotes the set of all histories in $\mathfrak{M}$.

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 use the abbreviations $\sim_A := \bigcup_{a \in A} \sim_a$ and $\approx_A := \bigcup_{a \in A} \approx_a$. We also write $\sim^{\mathfrak{M}}_A$ and $\approx^{\mathfrak{M}}_A$ if the model is not clear from the context. Note that relations $\sim_A$ and $\approx_A$ 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 and q' look the same).

Definition 9 (*ir*-, *iR*-strategies) An *imperfect information memoryless strategy* (*ir*-strategy in short) is an *Ir*-strategy satisfying the following additional constraint: if $q \sim_a q'$ then $s_a(q) = s_a(q')$.

An *imperfect information perfect recall strategy* (*iR*-strategy in short) of agent a is an *IR*-strategy such that, if $h \approx_a h'$, then $s_a(h) = s_a(h')$.

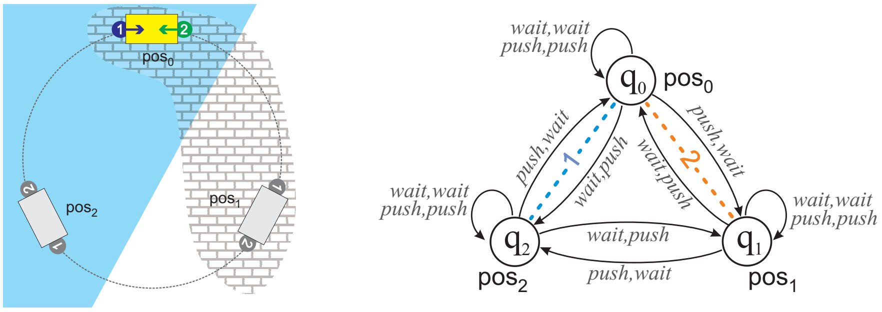

Fig. 2 Two robots and a carriage: (A) schematic view; (B) an imperfect information concurrent game structure M_0' that models the scenario. Dashed lines represent indistinguishability relations between states.

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).

Example 4 (Robots and Carriage, ctd.) We refine the scenario from Example 2 by restricting perception of the robots. Namely, we assume that robot 1 is only able to observe the color of the surface on which it is standing, and robot 2 perceives only the texture (cf. Figure 2). As a consequence, 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).

Note that the strategy from Example 3 cannot be used anymore because it is not uniform (indeed, the strategy tells robot 1 to wait in q_0 and push in q_2 but both states look the same to the robot).

3.2 Subjective Epistemic Outcome

Assumptions about agents' (un)certainty (i.e. the distinction between I and i) and recall (i.e. the distinction between R and r) are encoded in the mathematical structures that are used to represent strategies. However, if agent a is to make sure that a strategy s_a enforces property γ, it is not sufficient to consider only the paths from $out(q, s_a)$ because a does not necessarily know that q is the current state. To know that s_a guarantees γ, agent a should also check the outcome paths starting from states indistinguishable from q. From a conceptual point of view it makes sense to define two types of ability under imperfect information. *Objective* ability (i_o) means that a has an executable winning strategy, but the agent may be unaware of that, or be unable to identify the strategy on her own. *Subjective* ability (i_s) requires that a has a winning strategy *and* that a can identify such a strategy, i.e., the agents *knows how to play* and not only that a good way of playing exists.

On the semantic side, this is reflected by the set of paths that are taken into account. Objective ability refers to the outcome paths that can objectively happen, while subjective ability builds on the outcome paths that are subjectively possible according to a's available information.[6]

Definition 10 (Subjective epistemic outcome, x-outcome) Let $\mathfrak{M}$ be an $\mathbf{iCGS}$, q a state in it and s_A a collective strategy for group $A \subseteq \mathbb{A}\mathrm{gt}$. Let $x \in \{i_s, i_o, I\}$. The *x-outcome* $out^x_{\mathfrak{M}}(q, s_A)$ is defined as follows:

$$out^x_{\mathfrak{M}}(q, s_A) = \begin{cases} \bigcup_{q \sim_A q'} out_{\mathfrak{M}}(q', s_A) & \text{if } x = i_s; \\ out_{\mathfrak{M}}(q, s_A) & \text{else.} \end{cases}$$

Again, we omit $\mathfrak{M}$ if it is clear from context.

Example 5 (Robots and Carriage) In the scenario from Example 4, a possible uniform strategy of robot 1 is to push in q_0 and q_2, and wait in q_1. If the starting state is q_0 then the strategy *objectively* makes sure that the system will never move to q_2. However, robot 1 does not know that the strategy is successful in avoiding q_2 since he must take into account also the outcome paths starting from q_2 which trivially violate the path property $\Box\neg\mathsf{pos}_2$. Thus, 1 has the objective, but not the subjective, ability to enforce $\Box\neg\mathsf{pos}_2$ in state q_0.

In order to ensure a uniform notation, we introduce *xy-strategies* for $x \in \{i_s, i_o, I\}$ and $y \in \{r, R\}$ as follows:

IR: $s_a : St^+ \to Act$ such that $s_a(q_0 \dots q_n) \in d(a, q_n)$ for all $q_0, \dots, q_n$;
Ir: $s_a : St \to Act$ such that $s_a(q) \in d(a, q)$ for all q;
i_or, *i_sr*: like *Ir*, with the additional constraint that $q \sim_a q'$ implies $s_a(q) = s_a(q')$;
i_oR, *i_sR*: like *IR*, with the additional constraint that $h \approx_a h'$ implies $s_a(h) = s_a(h')$.

As before, collective xy-strategies s_A are tuples of individual xy-strategies s_a, one per $a \in A$. We emphasize that $i_s y$- and $i_o y$-strategies are defined in the very same way, only the notion of outcome is different. Note also 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. $\approx_A$).

3.3 Unified Setting: xy-Semantics of ATL

Finally, we put the pieces together and define the semantics of $\mathbf{ATL}_{xy}$, $\mathbf{ATL}^+_{xy}$, and $\mathbf{ATL}^*_{xy}$ for $x \in \{i_s, i_o, I\}$ and $y \in \{r, R\}$ by changing the clause for $\langle\!\langle A \rangle\!\rangle\gamma$ from Section 2.4 in the following way:

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

Note, again, that the I and i_o semantics look only at outcome paths starting from the current global state of the system. In contrast, the 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

[6] The issue is closely related to epistemic feasibility of plans, cf. e.g. [22,49].

that A can identify a single strategy that succeeds from all the states they consider possible. We follow [57] by taking the "everybody knows" interpretation of collective uncertainty. More general settings were proposed in [43, 39]. We believe that the results in this paper carry over to the other cases of "knowing how to play", too.

Example 6 (Robots and Carriage, ctd.) Consider the modified robots scenario from Example 4 (Figure 2). With observational capabilities of the robots restricted in this way, no agent knows how to make the carriage reach or avoid any selected state singlehandedly from q_0, i.e., $M_0', q_0 \models_{i_s y} \neg\langle\langle i\rangle\rangle\Diamond \mathsf{pos_j}$ and $M_0', q_0 \models_{i_s y} \neg\langle\langle i\rangle\rangle\Box\neg \mathsf{pos_j}$ for all $y \in \{\mathrm{r}, \mathrm{R}\}, i \in \{1, 2\}, j \in \{1, 2, 3\}$. Note in particular that the strategy from Example 3 cannot be used here because it is not uniform, and the strategy from Example 5 does not succeed because of outcome paths from indistinguishable states. Still, the latter strategy can be used to demonstrate that robot 1 has the objective ability to avoid q_2 (though not q_1 anymore): $M_0', q_0 \models_{i_o y} \langle\langle 1\rangle\rangle\Box\neg \mathsf{pos_2} \wedge \neg\langle\langle i\rangle\rangle\Box\neg \mathsf{pos_1}$.

The robots can keep the carriage away from $\mathsf{pos_1}$ together, but only in the objective sense: $M_0', q_0 \models_{i_o y} \langle\langle 1, 2\rangle\rangle\Box\neg \mathsf{pos_1}$. However, they cannot identify a strategy which guarantees that: $M_0', q_0 \models_{i_s y} \neg\langle\langle 1, 2\rangle\rangle\Box\neg \mathsf{pos_1}$ (when in q_0, robot 2 considers it possible that the system is already in the "bad" state q_1). So, do the robots know how to play to achieve anything? Yes, for example they know how to make the carriage reach a given state eventually: $M_0', q_0 \models_{i_s y} \langle\langle 1, 2\rangle\rangle\Diamond \mathsf{pos_1}$ etc. – it suffices that one of the robots pushes all the time and the other waits all the time.

For the above properties the type of robots' recall does not matter (they hold in both memoryless and perfect recall strategies). $\langle\langle 1, 2\rangle\rangle\Diamond\Box \mathsf{pos_1}$ is an example formula that distinguishes between the two sets of strategies. Note that $M_0', q_0 \models_{i_o r} \neg\langle\langle 1, 2\rangle\rangle\Diamond\Box \mathsf{pos_1}$: the robots have no memoryless strategy to bring the carriage to $\mathsf{pos_1}$ and keep it there forever, even in the objective sense. Still, they have a successful perfect recall strategy for that, and are able to identify it: $M_0', q_0 \models_{i_s R} \langle\langle 1, 2\rangle\rangle\Diamond\Box \mathsf{pos_1}$. The right strategy is that one robot pushes and the other waits for the first 3 steps. After that, they know their current position exactly, and can go straight to position 1 and stay there.

3.4 Folk Result: Memory Does Not Matter for Perfect Information

We observe that the basic semantics of **ATL*** from [9] corresponds exactly to **ATL**$^*_{IR}$. A folk result states that, in the restricted language of **ATL** both semantics for perfect information coincide. That is, exactly the same formulae characterize models and their states in **ATL**$_{IR}$ and **ATL**$_{Ir}$.

Proposition 1 *For every* i**CGS** $\mathfrak{M}$*, state* q*, and* ***ATL*** *formula* φ*, we have that* $\mathfrak{M}, q \models_{IR} \varphi$ *iff* $\mathfrak{M}, q \models_{Ir} \varphi$*.*

Proof. For finite models, equivalence of the semantics has been observed in [57], and follows from correctness of the model checking algorithm presented in [9]. It is not entirely obvious, however, that the result should extend to the infinite case. We present our own proof sketch below.

First, we observe that **ATL**$_{IR}$ and **ATL**$_{Ir}$ have the same validities. This follows from the results in [3] showing that: (1) perfect recall strategies in a **CGS** correspond to memoryless strategies in its tree unfolding, (2) every pointed **CGS** is strategically bisimilar to its tree unfolding, and (3) the same formulae of $\mathbf{ATL}_{Ir}$ hold in strategically bisimilar models (cf. also a more detailed exposition in Section 4.1).

Now we can prove the equivalence of $\mathfrak{M}, q \models_{IR} \varphi$ and $\mathfrak{M}, q \models_{Ir} \varphi$ by induction over the structure of φ. Cases $\varphi \equiv p, \neg\psi, \psi_1 \wedge \psi_2, \langle\!\langle A \rangle\!\rangle \bigcirc \psi$ are straightforward. For $\varphi \equiv \langle\!\langle A \rangle\!\rangle \Box \psi$, we take the axiom schemes $(\mathbf{FP}_\Box)$ and $(\mathbf{GFP}_\Box)$ from [31]. It was proved in [31] that all their instances are validities of $\mathbf{ATL}_{IR}$.[7] By the previous observation, all the instances of schemes $(\mathbf{FP}_\Box)$ and $(\mathbf{GFP}_\Box)$ are validities of $\mathbf{ATL}_{Ir}$ too. But that means that $\langle\!\langle A \rangle\!\rangle \Box$ is the greatest fixpoint of *the same monotone transformer of state sets* in both semantics $\models_{IR}$ and $\models_{Ir}$. Thus, the set of states satisfying $\langle\!\langle A \rangle\!\rangle \psi$ in $\mathfrak{M}$ is the same in both semantics.

The proof for $\langle\!\langle A \rangle\!\rangle \psi_1 \,\mathcal{U}\, \psi_2$ is analogous, by showing that its extension in $\models_{IR}$ and $\models_{Ir}$ is the least fixpoint of the same monotone transformer of state subsets from $\mathfrak{M}$. ■

Note that the *IR* and *Ir* semantics coincide only for the restricted syntactic variant **ATL**. For $\mathbf{ATL}^*$, and even $\mathbf{ATL}^+$, there are formulae that distinguish the two semantics, as we demonstrate in Section 5.1.

4 Perfect Recall and Tree-Like Unfoldings

Now we can turn to the original contribution of this paper. We begin by preparing the formal ground for our comparison of **ATL** validities under different semantics. In this section, we define several tree-like unfoldings of models, and show that they preserve truth of **ATL** formulae provided appropriate assumptions about agents' uncertainty and notion of success. This is the most technical part of the paper, needed mostly to prove the inclusion results in Section 5.1. However, its importance goes beyond technicalities. The unfoldings uncover some of the conceptual structure that underlies **ATL**. In particular, they expose a "forgetting" phenomenon in the semantics of **ATL**: even agents with perfect recall are assumed to forget their past observations when proceeding to a subtask specified by a nested subformula (like in $\langle\!\langle a \rangle\!\rangle \Diamond \underline{\langle\!\langle a \rangle\!\rangle \Box \mathsf{p}}$). In a way, one can talk about two variants of perfect recall: the "almost perfect recall" where agents use perfect recall strategies but abandon their previous observations when trying to enforce a nested strategic formula, and "truly perfect recall" where their hitherto observations carry over to the nested strategic task. On the other hand, our invariance theorems show that alternating-time temporal logic (even in its broadest syntactic variant $\mathbf{ATL}^*$) is too poor to distinguish between the two kinds of recall.

We believe that this section is of interest to readers who are intrigued by intricacies of game logics or search for tools that can be used to prove similar invariance results. On the other hand, readers interested only in the main *conceptual* contribution of this paper (i.e., the comparison of validities for variants of **ATL**) are advised to skip this part and proceed to Section 5.

[7] The proof in [31] was for the class of finitely branching **CGS** (with possibly infinite state spaces) but it extends to the case of infinite branching in a straightforward way.

Plan of Section 4. A *tree-like unfolding* of an $\mathbf{iCGS}$ is an (infinite) model in which nodes correspond to finite sequences of states (i.e., histories) in the original $\mathbf{iCGS}$. It is easy to see that the underlying transition structure of such an unfolding is a tree or a forest. The advantage of these structures is that perfect recall strategies and memoryless strategies coincide in tree-like unfoldings. Moreover, each perfect recall strategy in the original model corresponds to a memoryless strategy in the unfolding yielding an equivlant outcome, and vice versa. Both properties are rather standard in the perfect information setting. For imperfect information, however, the constructions are more involved due to the specialities of the *iR*-semantics; more precisely, the knowledge of agents is "reset" whenever a nested strategic modality is evaluated.

For each of the three semantic settings of:

– perfect information,
– imperfect information with the objective semantics,
– imperfect information with the subjective semantics,

we proceed as follows:

1. We characterise appropriate *tree-like structures* and show that memoryless and perfect recall strategies coincide on them;
2. We define appropriate *unfoldings* and show that they result in tree-like structures;
3. We show that the unfoldings are *truth-preserving* (i.e. a formula which is true in the original model is also true in the tree-like unfolding and vice versa).

4.1 Perfect Information

We begin with tree unfoldings of perfect information $\mathbf{CGS}$'s. We draw inspiration from the proof of [3, Theorem 8.3].

Definition 11 (Tree-like $\mathbf{CGS}$, $\rho_{\mathfrak{M}}(q_1, q_2)$) Let $\mathfrak{M}$ be a $\mathbf{CGS}$. $\mathfrak{M}$ is called *tree-like* iff there is a state q_0 (the *root*) such that for every q there is a unique history leading from q_0 to q.

Let q_1 and q_2 be states in $\mathfrak{M}$. If q_2 is reachable from q_1 then we use $\rho_{\mathfrak{M}}(q_1, q_2)$ to refer to the unique history from state q_1 to q_2; otherwise, if q_2 is not reachable from q_1 we set $\rho_{\mathfrak{M}}(q_1, q_2) = \epsilon$. Moreover, we use $\rho_{\mathfrak{M}}(q)$ as a shortcut for $\rho_{\mathfrak{M}}(q_0, q)$ (we will omit the subscript if clear from context). We note that $\rho_{\mathfrak{M}}(q_0) = q_0$.

Every state q in a tree-like $\mathbf{CGS}$ uniquely determines the path that leads from the root to q. Hence, perfect recall is already included in the states of the model. This is formally shown in the following proposition.

Proposition 2 (Recall invariance for tree-like $\mathbf{CGS}$) *For every tree-like $\mathbf{CGS}$ $\mathfrak{M}$, state q in $\mathfrak{M}$, and **ATL***-formula φ, we have: $\mathfrak{M}, q \models_{Ir} \varphi$ iff $\mathfrak{M}, q \models_{IR} \varphi$.*

Proof. The proof is done by induction over the structure of φ.
<u>Base cases</u>:

Propositional case: Straightforward.

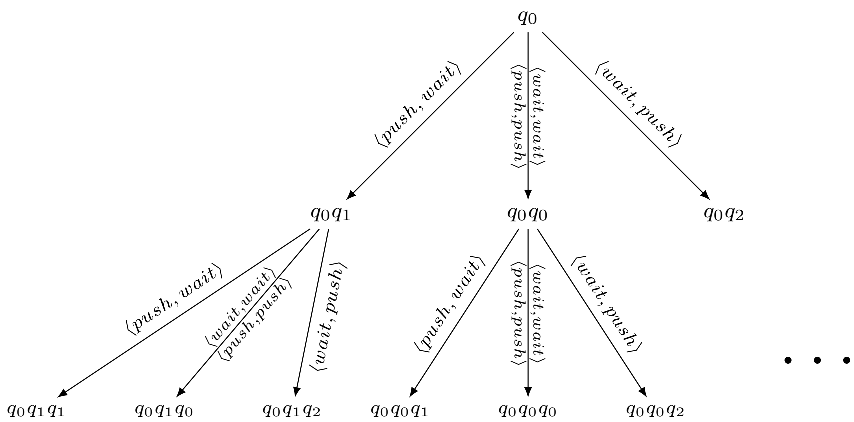

Fig. 3 Tree unfolding $T(M_0, q_0)$ for the robots and carriage $\mathbf{CGS}$ from Figure 1

Case: $\varphi \equiv \langle\!\langle A\rangle\!\rangle\gamma$ where γ contains no nested strategic modalities. The left-to-right direction is obvious. Now suppose that $\mathfrak{M}, q \models_{IR} \langle\!\langle A\rangle\!\rangle\gamma$ and let s_A be a collective *IR*-strategy for A such that for all $\lambda \in out(q, s_A)$ it holds that $\mathfrak{M}, \lambda \models_{IR} \gamma$. We define $t_A(q') = s_A(\rho(q, q'))$ for each state q' in $\mathfrak{M}$ reachable from q. Then, t_A is a well-defined *Ir*-strategy with $out(q, t_A) = out(q, s_A)$. Hence, we have $\mathfrak{M}, q \models_{Ir} \langle\!\langle A\rangle\!\rangle\gamma$.

Induction step:

Case: $\varphi \equiv \psi_1 \wedge \psi_2$. Straightforward.
Case: $\varphi \equiv \neg\psi$. $\mathfrak{M}, q \models_{Ir} \neg\psi$ iff not $\mathfrak{M}, q \models_{Ir} \psi$ iff (by induction hypothesis) not $\mathfrak{M}, q \models_{IR} \psi$ iff $\mathfrak{M}, q \models_{IR} \neg\psi$.
Case: $\varphi \equiv \langle\!\langle A\rangle\!\rangle\gamma$. We observe that each state q' at which a state subformula ψ of γ is evaluated forms the root of a tree-like $\mathbf{CGS}$. Then, by induction, ψ has the same truth value in q' according to the *IR*- and *Ir*-semantics and can be replaced by a new atomic proposition with the appropriate valuation. This yields formula $\varphi' \equiv \langle\!\langle A\rangle\!\rangle\gamma'$ with no nested strategic modalities, to which we apply the same argument as above.

■

A natural question is whether every model has an equivalent tree-like $\mathbf{CGS}$. By "equivalent" we mean that the sets of formulae which hold at corresponding states are always the same.

Definition 12 (Tree unfolding) Let $\mathfrak{M} = (\mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o)$ be a $\mathbf{CGS}$ and q be a state in it. The *(perfect information) tree unfolding of the pointed model* $(\mathfrak{M}, q)$ denoted $T(\mathfrak{M}, q)$ is defined as $(\mathbb{A}\mathrm{gt}, St', \Pi, \pi', Act, d', o')$ where

- $St' := \Lambda^{fin}_{\mathfrak{M}}(q)$,
- $d'(a, h) := d(a, last(h))$,
- $o'(h, \boldsymbol{\alpha}) := h \circ o(last(h), \boldsymbol{\alpha})$, and
- $\pi'(h) := \pi(last(h))$.

The node q in the unfolding is called the *root* of $T(\mathfrak{M}, q)$.

An example tree unfolding is shown in Figure 3. It is important to note that histories in $\mathfrak{M}$ are states in $T(\mathfrak{M}, q)$ and that each tree unfolding is tree-like:

Proposition 3 *The tree unfolding of a pointed* **CGS** $(\mathfrak{M}, q)$ *is tree-like.*

We now show that satisfaction of **ATL***-formulae is invariant under tree unfoldings and that memory is not needed in the tree unfolding.

Theorem 1 *For every* **CGS** $\mathfrak{M}$*, state* q *in* $\mathfrak{M}$*, and* ***ATL******-formula* φ *we have:*

$$\mathfrak{M}, q \models_{IR} \varphi \text{ iff } T(\mathfrak{M}, q), q \models_{IR} \varphi \text{ iff } T(\mathfrak{M}, q), q \models_{Ir} \varphi.^{8}$$

Proof. The second equivalence follows from Propositions 2 and 3. To prove the first equivalence we show that, for all $h \in \Lambda^{fin}_{\mathfrak{M}}(q)$, we have $\mathfrak{M}, last(h) \models_{IR} \varphi$ iff $T(\mathfrak{M}, q), h \models_{IR} \varphi$ (by induction over the structure of φ). Note that for each state q' reachable from q in $\mathfrak{M}$ there is a history h such that $last(h) = q'$.
Base cases:

Propositional case: Straightforward.
Case: $\varphi \equiv \langle\!\langle A \rangle\!\rangle \gamma$ where γ does not contain any strategic modalities.

"$\Rightarrow$": Suppose that $\mathfrak{M}, last(h) \models_{IR} \varphi$. Then, there is an *IR*-strategy s_A such that for all $\lambda \in out(last(h), s_A)$ we have that $\mathfrak{M}, \lambda \models \gamma$. Now let t_A be an *Ir*-strategy defined as follows: $t_A(hh') = s_A(last(h)h')$, and arbitrary otherwise. By definition of the tree unfolding and the construction of t_A we have that

$$\begin{array}{l} last(h)q_1q_2\cdots \in out_{\mathfrak{M}}(last(h), s_A) \text{ iff} \\ (h)(hq_1)(hq_1q_2)\cdots \in out_{T(\mathfrak{M},q)}(h, t_A). \end{array}$$

Since the valuation of propositions only depends on the final state of a history and since *Ir*-strategies can be seen as special cases of *IR*-strategies, we have also that $T(\mathfrak{M}, q), h \models_{IR} \varphi$.

"$\Leftarrow$": Suppose that $T(\mathfrak{M}, q), h \models_{IR} \varphi$. Then, by Propositions 2 and 3, there is an *Ir*-strategy s_A such that for all $\lambda \in out(h, s_A)$ we have that $T(\mathfrak{M}, q), \lambda \models \gamma$. We define the following *IR*-strategy t_A:

$$t_A(h') = \begin{cases} s_A(h(h'[1,\infty])) & \text{if } h'[0] = last(h) \\ \alpha & \text{else, for some arbitrary } \alpha \in d_A(last(h')) \end{cases}$$

The first case of the definition t_A applies if $h'[1, \infty]$, i.e. h' without the initial state, is a possible extension of history h. The history $h(h'[1, \infty])$ is the extension of h with h' where the last state of h or the initial state of h' has to be removed as it occurs twice. Again, we have

$$\begin{array}{l} last(h)q_1q_2\cdots \in out_{\mathfrak{M}}(last(h), t_A) \text{ iff} \\ (h)(hq_1)(hq_1q_2)\cdots \in out_{T(\mathfrak{M},q)}(h, s_A) \end{array}$$

[8] The equivalence of $\mathfrak{M}, q \models_{IR} \varphi$ and $T(\mathfrak{M}, q), q \models_{IR} \varphi$ follows also from the results on alternating bisimulation, cf. [10] for the bisimulation in finite models, and [3] for the general case. We present the construction nevertheless, as it will be adapted in the following sections to the case of imperfect information.

and thus $\mathfrak{M}, last(h) \models_{IR} \varphi$.

Induction step:

Case: $\varphi \equiv \psi_1 \wedge \psi_2$. Straightforward.
Case: $\varphi \equiv \neg\psi$. Straightforward.
Case: $\varphi \equiv \langle\langle A \rangle\rangle \gamma$. We observe that for each state q' in $\mathfrak{M}$ reachable from q at which a state-subformula φ' of γ is evaluated there is a history h such that $T(\mathfrak{M}, q)$ contains a state hq' at which the very subformula holds (by induction hypothesis). Then we apply the same reasoning as for the case with no nested strategic modalities.

■

4.2 Imperfect Information: Objective Ability

Unlike in the perfect information case, tree unfoldings for imperfect information must also take into account the indistinguishability relations. We construct our argument for the i_o case similarly to Section 4.1. The notion of tree-like imperfect information $\mathbf{CGS}$ has to include suitable constraints on the epistemic relations – otherwise we would not get truth invariance with respect to recall. To handle the issue, we introduce *objective epistemic tree unfoldings under perfect recall*, or *i_oR-tree unfoldings* in short.

Definition 13 (Tree structure of i$\mathbf{CGS}$) Let $\mathfrak{M}$ be an **i**$\mathbf{CGS}$. $\mathfrak{M}$ has *tree structure* iff the underlying $\mathbf{CGS}$ of $\mathfrak{M}$ (i.e., $\mathfrak{M}$ without epistemic relations) is tree-like. As in Definition 11 we use $\rho_{\mathfrak{M}}(q_1, q_2)$ to refer to the *unique history* between q_1 and q_2 in $\mathfrak{M}$ if it exists and set $\rho_{\mathfrak{M}}(q) = \rho_{\mathfrak{M}}(q_0, q)$ where q_0 is the root in $\mathfrak{M}$. Again, we omit $\mathfrak{M}$ from $\rho_{\mathfrak{M}}(\cdot)$ if clear from context.

Definition 14 (*i_oR*-tree-like) Let $\mathfrak{M}$ be an **i**$\mathbf{CGS}$ with tree structure. $\mathfrak{M}$ is called *i_oR-tree-like* iff for all $a \in \mathbb{A}\mathrm{gt}$ and all $q_1, q_2 \in St$ we have $q_1 \sim_a^{\mathfrak{M}} q_2$ iff $\rho(q_1) \approx_a^{\mathfrak{M}} \rho(q_2)$. (We note that $\rho(q_1) \approx_a^{\mathfrak{M}} \rho(q_2)$ implies $q_1 \sim_a^{\mathfrak{M}} q_2$ by definition of $\approx_a^{\mathfrak{M}}$.)

In other words, in an i_oR-tree-like structure the information sets in a game can only be more precise when the game already follows some previous interaction. The next proposition is analogous to Proposition 2.

Proposition 4 (Recall invariance for *i_oR*-tree-like models) *For every i_oR-tree-like* ***i*$\mathbf{CGS}$** *$\mathfrak{M}$, state q in $\mathfrak{M}$, and* ***ATL******-formula φ, we have that $\mathfrak{M}, q \models_{i_or} \varphi$ iff $\mathfrak{M}, q \models_{i_oR} \varphi$.*

Proof. Induction over the structure of φ.
Base cases:

Propositional case: Straightforward.

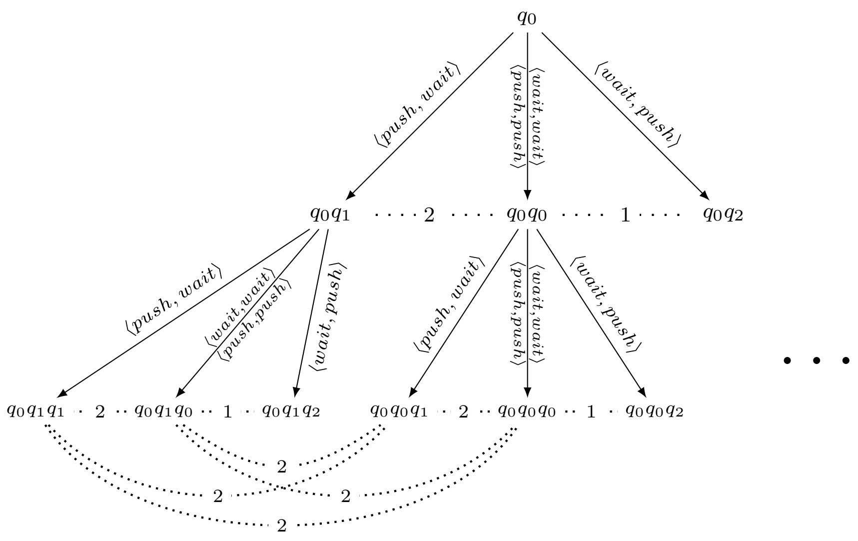

Fig. 4 *i_oR*-tree unfolding $T_o(M_0', q_0)$ of the robots with limited information from Figure 2 (we omitted reflexive epistemic links).

Case: $\varphi \equiv \langle\langle A\rangle\rangle\gamma$ where γ contains no strategic modalities. The left-to-right direction is obvious. Now suppose that $\mathfrak{M}, q \models_{i_oR} \langle\langle A\rangle\rangle\gamma$ and let s_A be a collective *iR*-strategy for A such that for all $\lambda \in out(q, s_A)$ it holds that $\mathfrak{M}, \lambda \models_{i_oR} \gamma$. We define $t_a(q') = s_a(\rho(q, q'))$ for each state q' in $\mathfrak{M}$ which is reachable from q. Then, we have $out(q, t_A) = out(q, s_A)$. Moreover, we have for all states q_1 and q_2 and all agents $a \in \mathbb{A}\mathrm{gt}$ that $q_1 \sim_a q_2$ iff $\rho(q_1) \approx_a \rho(q_2)$ because: (a) the right-to-left direction is clear from the definition of $\approx$, and (b) the left-to-right direction follows because $\mathfrak{M}$ is *i_oR*-tree-like. Hence, t_A is a well-defined *i_or*-strategy with $out(q, t_A) = out(q, s_A)$ and thus: $\mathfrak{M}, q \models_{i_or} \langle\langle A\rangle\rangle\gamma$.

Induction step:

Case: $\varphi \equiv \psi_1 \wedge \psi_2$. Straightforward.
Case: $\varphi \equiv \neg\psi$. $\mathfrak{M}, q \models_{i_or} \neg\psi$ iff not $\mathfrak{M}, q \models_{i_or} \psi$ iff (by induction hypothesis) not $\mathfrak{M}, q \models_{i_oR} \psi$ iff $\mathfrak{M}, q \models_{i_oR} \neg\psi$.
Case: $\varphi \equiv \langle\langle A\rangle\rangle\gamma$. We observe that each state q' at which a state subformula ψ of γ is evaluated forms the root of a *i_oR*-tree-like $\mathbf{iCGS}$. Then, by induction, ψ has the same truth value in q' according to the *i_oR*- and *i_or*-semantics and can be replaced by a new atomic proposition with the appropriate valuation. This yields formula $\varphi' \equiv \langle\langle A\rangle\rangle\gamma'$ with no nested strategic modalities, to which we apply the same argument as above.

■

Now, the *i_oR-tree unfolding* is defined as standard tree unfolding for the perfect information case extended with indistinguishability relations between histories of the model (which are nodes of the unfolding).

Definition 15 (i_oR-tree unfolding) Given an $\mathbf{iCGS}$ $\mathfrak{M}$ and a state q in it, we define the *i_oR-tree unfolding of* $(\mathfrak{M}, q)$, denoted $T_o(\mathfrak{M}, q)$, as $T(\mathfrak{M}, q)$ from Definition 12 extended with epistemic relations $\sim_a^{T_o(\mathfrak{M},q)}$ reflecting indistinguishability of histories in $\mathfrak{M}$; that is, $h \sim_a^{T_o(\mathfrak{M},q)} h'$ iff $h \approx_a^{\mathfrak{M}} h'$ where h and h' start in q.

As an example, the i_oR-tree unfolding of the robots and carriage $\mathbf{iCGS}$ is presented in Figure 4.

Proposition 5 *Let $(\mathfrak{M}, q)$ be a pointed $\mathbf{iCGS}$. The i_oR-tree unfolding of $(\mathfrak{M}, q)$ is i_oR-tree-like.*

Proof. Clearly, the unfolding has tree structure and is i_oR-tree-like by definition of the indistinguishability relations in the i_oR-tree-unfolding. ■

Analogously to Theorem 1 we have that i_oR-tree unfoldings are truth preserving and that memory does not matter in these unfolded models.

Theorem 2 *For every $\mathbf{iCGS}$ $\mathfrak{M}$, state q in $\mathfrak{M}$, and $\mathbf{ATL}^*$-formula φ we have:*

$$\mathfrak{M}, q \models_{i_oR} \varphi \text{ iff } T_o(\mathfrak{M}, q), q \models_{i_oR} \varphi \text{ iff } T_o(\mathfrak{M}, q), q \models_{i_or} \varphi.$$

The proof is given in the appendix on page 40.

4.3 Imperfect Information: Subjective Ability

The case for the subjective semantics (i_s) cannot be proven in the same way by using i_oR-tree unfoldings. Obviously, when constructing an unfolding of $(\mathfrak{M}, q)$ for the i_sR-semantics one has to take into account paths starting from states indistinguishable from q. A first naive approach could be to define the i_sR-unfolding as a structure consisting of i_oR-tree unfoldings, one for each epistemic alternative, and to connect the root nodes of all these unfoldings. Unfortunately, this simple idea is not sufficient as illustrated in Example 7.

Example 7 (First naive approach to i_sR-tree unfoldings) We consider the $\mathbf{iCGS}$ M_1 shown in Figure 5. The story is as follows. A man wants to shoot down a yellow rubber duck in a shooting gallery. 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. He can shoot to the left (action $shoot_L$) or to the right ($shoot_R$). Alternatively, he can reach out and open one of the cells for a moment (action $look$), thus removing his uncertainty.

Let us take the i_oR-tree unfoldings $T_o(M_1, q_0)$ and $T_o(M_1, q_1)$, and interconnect their nodes by epistemic links whenever the corresponding histories are indistinguishable in the original model. The resulting model is shown in Figure 6 (we will call the model T_1). Unfortunately, this construction is *not* truth-preserving. That is because if a state-subformula is evaluated in states 040 and 151 of T_1 the agent will know where the game is – which is not consistent with the i_s semantics: only the last state of each history should be considered.

To be more precise, let us consider formula $\varphi = \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \bigcirc \mathsf{shot}$. Clearly, we have $M_1, q_0 \not\models_{i_sR} \varphi$. On the other hand, we have $T_1, q_0 \models_{i_sR} \varphi$.

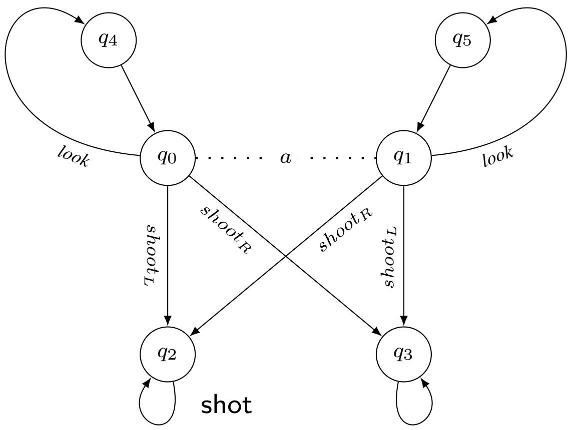

Fig. 5 "Poor duck model" M_1 with one player (a) and transitions labeled with a's actions. The dotted line depicts a's indistinguishability relation. Automatic transitions (i.e., such that there is only one possible transition from the starting state) are left unlabeled.

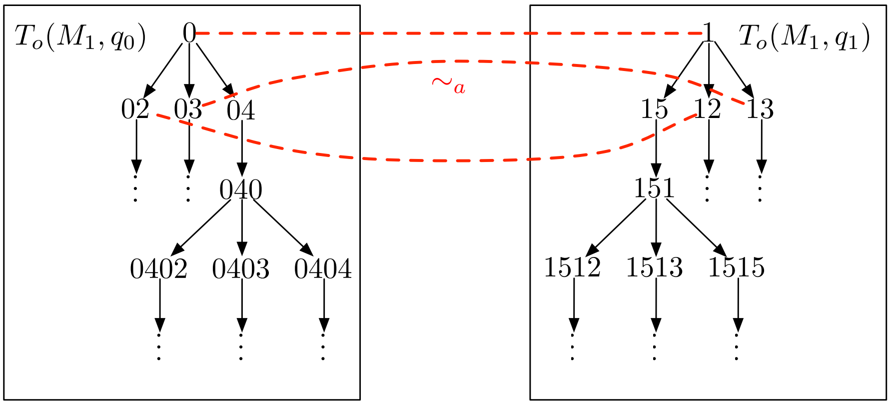

Fig. 6 Two i_oR-tree unfoldings connected by epistemic links. Each label $i_1 i_2 \ldots$ refers to the history $q_{i_1} q_{i_2} \ldots$.

In order to improve the naive approach one may simply add an epistemic link between states 040 and 151. Unfortunately, this does not work either. Such a link indicates that the states 040 and 151 are indistinguishable for a; on the other hand, player a can distinguish the histories which lead to these states. This contradicts the conceptual idea in which states are associated with histories. Moreover, it is easy to construct a concrete counterexample.

To make the observation in Example 7 more formal, suppose hq is some node in the i_oR-tree unfolding $T_o(\mathfrak{M}, q_0)$ and that in this node a formula $\langle\!\langle a \rangle\!\rangle\gamma$ is evaluated.

Then, $\langle\langle a\rangle\rangle\gamma$ holds iff agent a has a successful *iR*-strategy not only for all paths starting from hq, but also for paths starting from nodes $h'q'$ such that $q \sim_a^{\mathfrak{M}} q'$. In the i_oR-tree unfolding, however, these nodes are usually not linked via an epistemic transition. On the other hand, we cannot simply introduce the link $hq \sim_a^{T_o(\mathfrak{M},q_0)} h'q'$ as we would loose soundness of the construction (in general, h and h' do not need to be indistinguishable). This observation makes it necessary to introduce a more sophisticated construction for the subjective epistemic tree-like unfoldings under perfect recall, or *i_sR-pando unfoldings* in short.[9]

Firstly, we discuss when an **iCGS** should be considered *i_sR-pando-like*. The idea of a set of connected i_oR-tree-like models (like in Example 7) seems to come close. However, we should also account for the "forgetting" of the history of the play when a nested strategic operator is evaluated. This is because if a state subformula (like $\langle\langle a\rangle\rangle\gamma$) is evaluated in a history h, only the last state of h is relevant. The rest of h is "lost" as it does not influence the truth of $\langle\langle a\rangle\rangle\gamma$ in h. We deal with it by adding appropriate "hanging" trees with roots q that are indistinguishable from $last(h)$ in the original models. The new trees are connected to histories in the "basic" tree by appropriate epistemic links. We must also make sure that there are no epistemic links between such trees apart from the ones just explained.

Definition 16 (i_sR-pando-like, $\rho_{\mathfrak{M}}(q, q')$) An **iCGS**

$$\mathfrak{M} = \langle \mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o, \{\sim_a\}_{a\in\mathbb{A}\mathrm{gt}}\rangle$$

is *i_sR-pando-like* iff it consists of submodels

$$\mathfrak{M}_i = \langle \mathbb{A}\mathrm{gt}, St_i, \Pi, \pi_i, Act, d_i, o_i, \{\sim_a^{\mathfrak{M}_i}\}_{a\in\mathbb{A}\mathrm{gt}}\rangle$$

for $i \in I$ and some index set $I \subseteq \mathbb{N}$, where:

- each $\mathfrak{M}_i$ is an i_oR-tree-like **iCGS**,
- $St = \biguplus_{i\in I} St_i$ (i.e. the states of the $\mathfrak{M}_i$'s form a partition of $\mathfrak{M}$),
- $\pi : St \to 2^{\Pi}$ with $\pi(q) = \pi_i(q)$ for $q \in St_i$,
- $d : St \times \mathbb{A}\mathrm{gt} \to 2^{Act}$ with $d(q, a) = d_i(q, a)$ for $q \in St_i$,
- $o : St \times Act^{|\mathbb{A}\mathrm{gt}|} \to St$ with $o(q, \alpha) = o_i(q, \alpha)$ for $q \in St_i$, and
- $\sim_a \subseteq St \times St$ with $\sim_a := \left(\bigcup_{i\in I} \sim_a^{\mathfrak{M}_i}\right) \cup \hat{\sim}_a$ where each $\hat{\sim}_a \subseteq \bigcup_{i\in I} St_i \times \bigcup_{i\in I} St_i$,

and the following conditions are satisfied:

1. the relation $\sim_a$ is transitive for every $a \in \mathbb{A}\mathrm{gt}$.
2. $\hat{\sim}_a$ is a symmetric relation for every $a \in \mathbb{A}\mathrm{gt}$.
3. for all $i \in I$ we have $\hat{\sim}_a \cap (St_i \times St_i) = \emptyset$ (the relation does only exist between different i_oR-tree like models).

[9] We thank an anonymous reviewer of JAAMAS for the excellent terminological suggestion. An i_sR unfolding is not a tree, as it usually consists of several transition trees. On the other hand, it is not a typical forest because the trees are not separate – they are intimately connected by epistemic relations. For the biological Pando, see for example http://en.wikipedia.org/wiki/Pando(tree).

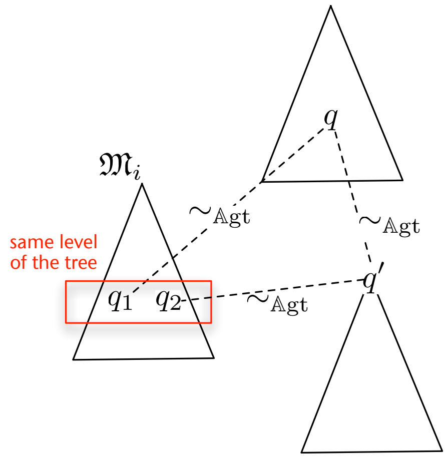

Fig. 7 Condition 5 of Definition 16.

4. for $q_1 \in St_i$ and $q_2 \in St_j$ with $i, j \in I$, $i \neq j$, we have that if $q_1 \hat{\sim}_a q_2$ then $\rho_{\mathfrak{M}_i}(q_1) \hat{\approx}_a^{\mathfrak{M}} \rho_{\mathfrak{M}_j}(q_2)$ or $\rho_{\mathfrak{M}_i}(q_1) = q_1$ or $\rho_{\mathfrak{M}_j}(q_2) = q_2$ where for two histories h and h' we have $h \hat{\approx}_a^{\mathfrak{M}} h'$ iff $|h| = |h'|$ and $h[l] \hat{\sim}_a h'[l]$ for all $l = 0, \ldots, |h| - 1$ (indistinguishable nodes in different $i_o R$-tree-like models must have indistinguishable histories or at least one of the nodes is a root node).
5. for $q_1, q_2 \in St_i$ and $\sim_{\mathbb{A}\mathrm{gt}} = \bigcup_{a \in \mathbb{A}\mathrm{gt}} \sim_a$, if $q_1 (\sim_{\mathbb{A}\mathrm{gt}})^* q_2$ then $|\rho_{\mathfrak{M}_i}(q_1)| = |\rho_{\mathfrak{M}_i}(q_2)|$ (nodes in the same tree indistinguishable for a group must be on the same level). The idea behind this condition is illustrated in Figure 7.

Moreover, we define $\rho_{\mathfrak{M}}(q_1, q_2)$ as $\rho_{\mathfrak{M}_i}(q_1, q_2)$ for $q_1, q_2 \in St_i$ and set $\rho_{\mathfrak{M}}(q_1, q_2) = \epsilon$ if $q_1 \in St_i$ and $q_2 \in St_j$ for $i \neq j$.

Remark 1 We would like to note that it is possible, due to condition 4 of Definition 16, to weaken condition 5 of Definition 16 to the following: Let $q_1 \in St_i$, $q_2 \in St_j$, $i, j \in I$, $i \neq j$, $q_1 (\sim_{\mathbb{A}\mathrm{gt}})^* q_2$, and $\rho_{\mathfrak{M}_j}(q_2) = q_2$ where $\sim_A = \bigcup_{a \in A} \sim_a$ for $A \subseteq \mathbb{A}\mathrm{gt}$. If there is an $q_1' \in St_i$, with $q_2 (\sim_{\mathbb{A}\mathrm{gt}})^* q_1'$; then, $|\rho_{\mathfrak{M}_i}(q_1)| = |\rho_{\mathfrak{M}_i}(q_1')|$.

Before we give an intuitive example we show that the concept of $i_s R$-pando-like **iCGS** is well-defined.

Proposition 6 *Let $\mathfrak{M}$ be an $i_s R$-pando like* **iCGS** *as defined in Definition 16. Each relation $\sim_a$ is an equivalence relation for $a \in \mathbb{A}\mathrm{gt}$.*

Proof. By definition, each $\sim_a$ is transitive. Symmetry follows from the symmetry of $\hat{\sim}_a$ and of $\sim_a^{\mathfrak{M}_i}$. Reflexivity of $\sim_a^{\mathfrak{M}_i}$ does also imply reflexivity of $\sim_a$.

■

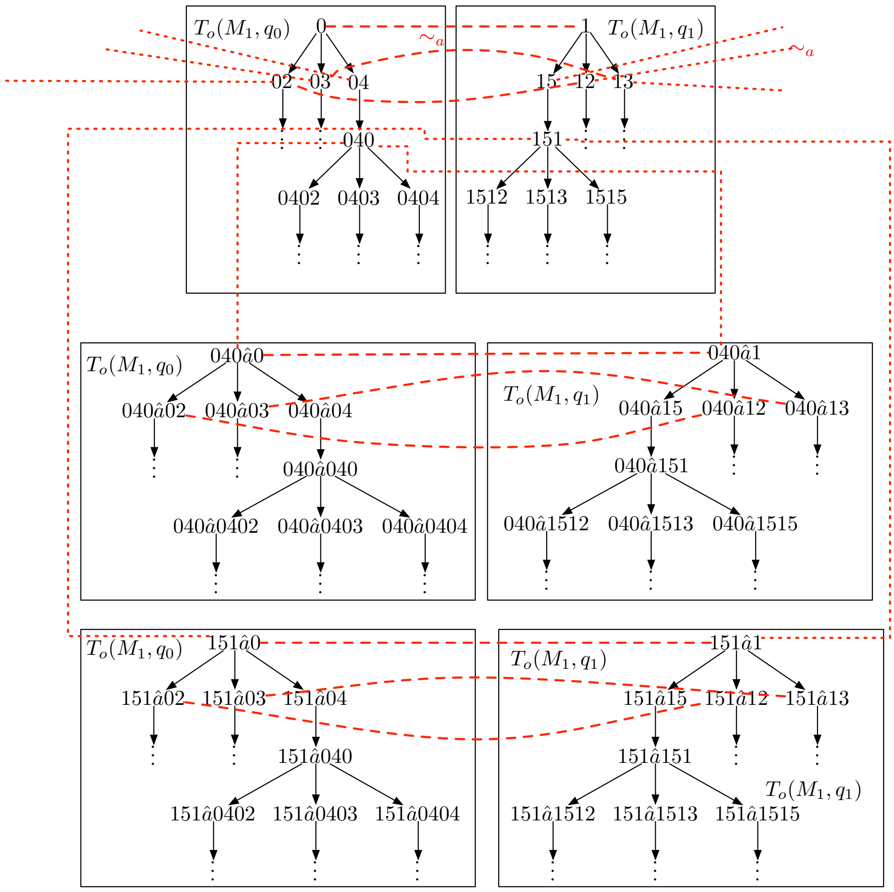

Fig. 8 The figure shows the i_sR-pando unfolding of (M_1, q_0) from Figure 5. All dotted and dashed lines denote the indistinguishability relation of agent a. Dashed links encode indistinguishability between nodes in trees that have roots on the same level; dotted links connect nodes in trees from different levels. Again, for the sake of readability, reflexive and transitive connections are omitted.

Example 8 Figure 8 depicts an i_sR-pando-like **iCGS**. In fact, the model shows a suitable unfolding of the pointed **iCGS** (M_1, q_0) from Figure 5. We will formally introduce i_sR-pando unfoldings in Definition 17.

In the spirit of Propositions 2 and 4 we have that memory is not needed in i_sR-pando-like models. The proof for the left-to-right direction is obvious. The sophisticated step is to construct an i_sr-strategy from an i_sR-strategy. For the sake of readability we have moved the technical part in the appendix (Lemma 3 on page 41).

Proposition 7 (Recall invariance for i_sR-pando-like models) *For every i_sR-pando-like **iCGS** $\mathfrak{M}$, state q in $\mathfrak{M}$, and **ATL***-formula φ, we have that $\mathfrak{M}, q \models_{i_sr} \varphi$ iff $\mathfrak{M}, q \models_{i_sR} \varphi$.*

Proof. Firstly, we recall that all the "subpandos" which form an i_oR-tree-like **iCGS** are not interconnected by transitions and thus the path to each state is unique. The proof is done analogously to Proposition 4; we only consider the important base case where $\varphi \equiv \langle\langle A\rangle\rangle\gamma$ and γ contains no strategic modalities. The left-to-right direction is obvious.

For the right-to-left direction suppose that $\mathfrak{M}, q \models_{i_sR} \langle\langle A\rangle\rangle\gamma$ and let s_A be a collective *iR*-strategy for A such that for all $\lambda \in out^{i_s}(q, s_A)$ it holds that $\mathfrak{M}, \lambda \models_{i_sR} \gamma$. Let $q' \in St_j$ be a state reachable from $\hat{q} \in St_j$ with $\hat{q} \sim_a^{\mathfrak{M}} q$ for some $a \in A$. Then, we define the memoryless strategy t_a as follows: $t_a(q') = s_a(\rho_{\mathfrak{M}_j}(\hat{q}, q'))$.
We proceed like this for all states q', $\hat{q}$ and define the strategies t_a arbitrarily but in a uniform way for all other states in $\mathfrak{M}$. (Note, that these are all states which are not reachable from any epistemic alternative of q for some agent in A.) Firstly, we observe that each t_a is well-defined as each $\mathfrak{M}_j$ is i_oR-tree-like and thus the path $\rho_{\mathfrak{M}_j}(\hat{q}, q')$ to a state q' is unique.

In order to show that each t_a is uniform and that $out^{i_s}(q, t_A) = out^{i_s}(q, s_A)$ we have to prove that for any two states q_1 reachable from $\hat{q}_1$, and q_2 reachable from $\hat{q}_2$ with $q_1 \sim_a^{\mathfrak{M}} q_2$ and $\hat{q}_1, \hat{q}_2 \in \{q' \in St \mid q \sim_A q'\}$ (i.e. $\hat{q}_1 \sim_b^{\mathfrak{M}} q \sim_c^{\mathfrak{M}} \hat{q}_2$ for some $b, c \in A$) we also have $\rho(\hat{q}_1, q_1) \approx_a^{\mathfrak{M}} \rho(\hat{q}_2, q_2)$. This part is shown in Lemma 3 in the appendix on page 41. ■

The basic idea of the *subjective epistemic pando unfolding under perfect recall* (*i_sR-pando unfolding* in short) is to create copies of the tree starting in q', one for each epistemic alternative. Then, we can link hq with these new root nodes q' of the "copies" of the trees starting in q' (cf. Figure 8 and take e.g. $h = 04$, $q = 0$, and $q' = 040\hat{a}1$; the new node is named $040\hat{a}1$ to ensure that the name is unique as explained below). It is easy to see that these "new" subtrees can only be reached if a formula $\langle\langle a\rangle\rangle\gamma$ is evaluated in hq (or some other state $h''q''$ with $hq \sim_a^{T_s(\mathfrak{M},q_0)} h''q''$ by transitivity). As mentioned above all nodes in these new subtrees must have unique names. This is the reason why we have to prefix each node h'' in the new tree by $hq\hat{a}$ where hq is the history in the "current tree" and $\hat{a}$ encodes that we have used a's indistinguishability relation to reach the "new" tree.

Before we formally define the i_sR-pando unfolding, we introduce some additional notation. In the following, we consider words over $D := (St \cup St \circ \{\hat{a} \mid a \in \mathbb{A}\mathrm{gt}\} \circ St)^+$. Thus, D consists of finite sequences of states, possibly interleaved by references to some agents. We use elements from D to give names to nodes of the pando unfolding. Essentially, the name of a node shows how the node is reached from a root by following temporal paths and "jumping" between different trees by use of epistemic links (cf. Figure 8).

We also define auxiliary functions $rel : D \to St^+$, $ref : D \to D$, $lastr : D \to St$ and $jump : D \to \mathbb{A}\mathrm{gt} \cup \{\epsilon\}$ as follows:

$$rel(h) = \begin{cases} h & \text{, for } h \in St^+; \\ h'' & \text{, for } h = h'\hat{a}h'' \text{ and } h'' \in St^+ \text{ and } a \in \mathbb{A}\mathrm{gt}; \end{cases}$$

$$ref(h) = \begin{cases} h & \text{if } h \in St^+; \\ \hat{h} & \text{if } h = \hat{h}\hat{a}rel(h); \end{cases}$$

$$lastr(h) = last(rel(h));$$
$$jump(h) = \begin{cases} \epsilon & \text{if } h \in St^+; \\ a & \text{if } h = \hat{h}\hat{a}rel(h). \end{cases}$$

The intuition for these functions is as follows. Given an element $h \in D$, $rel(h)$ returns the "relevant" part of h, i.e., the subhistory at the end of h of maximal length that does not contain any $\hat{a}$ symbol for any $a \in \mathbb{A}\mathrm{gt}$. On the other hand, $ref(h)$ returns the "reference" node in the higher-level tree which was used to obtain h. Finally, $jump(h)$ returns the agent whose epistemic link was used to "jump" between the two trees. For example, $rel(q_1\hat{a}q_2q_3\hat{b}q_4) = q_4$, $ref(q_1\hat{a}q_2q_3\hat{b}q_4) = q_1\hat{a}q_2q_3$, $lastr(ref(q_1\hat{a}q_2q_3\hat{b}q_4)) = q_3$ and $jump(q_1\hat{a}q_2q_3\hat{b}q_4) = b$.

Let $\mathfrak{M} = (\mathbb{A}\mathrm{gt}, St, \Pi, \pi, \{\sim_a\}_{a\in\mathbb{A}\mathrm{gt}}, Act, d, o)$ be an **iCGS** and $q \in St$. We recursively define sets $\Delta^i_{\mathfrak{M}} \subseteq D$ which contain the nodes of the *i*$_s$*R*-pando unfolding:

$$\Delta^0_{\mathfrak{M}}(q) := \bigcup_{q' \sim^{\mathfrak{M}}_a q} \Lambda^{fin}_{\mathfrak{M}}(q'),$$
$$\Delta^{i+1}_{\mathfrak{M}}(q) := \{h\hat{a}h' \mid h \in \Delta^i_{\mathfrak{M}}(q),\ |rel(h)| \geq 2,\ a \in \mathbb{A}\mathrm{gt},\ \text{and } h' \in \Lambda^{fin}_{\mathfrak{M}}(q') \text{ for some } q' \sim^{\mathfrak{M}}_a lastr(h)\}.$$

We write $\Delta^i_{\mathfrak{M}}$ for $\Delta^i_{\mathfrak{M}}(q)$ if state q is clear from context. Note that each $h \in \Delta^i_{\mathfrak{M}}$ contains exactly i symbols of type $\hat{a}_j$ for $a_j \in \mathbb{A}\mathrm{gt}$ and $j = 1, \ldots, i$. Intuitively, $\hat{a}$ denotes that we took a $\sim_a$-relation step between different trees. Note also that, for instance, $q_0q_1\hat{a}q_2 \in \Delta^1(q_0)$ but $q_0\hat{a}q_2 \notin \Delta^1(q_0)$. This is because if a link to a new tree model is taken histories have to be "forgotten" and in cases in which the history consists of a single state (e.g. q_0 in $q_0\hat{a}q_2$) such a link is not necessary and also not desired due to technical reasons. Now, we are ready to define the i_sR-pando unfolding.

Definition 17 (i_sR-pando unfolding) Let $\mathfrak{M} = (\mathbb{A}\mathrm{gt}, St, \Pi, \pi, \{\sim_a\}_{a\in\mathbb{A}\mathrm{gt}}, Act, d, o)$ be an **iCGS** and $q \in St$. The *i_sR-pando unfolding* of $(\mathfrak{M}, q)$, denoted $T_s(\mathfrak{M}, q)$, is defined as $T_s := T_s(\mathfrak{M}, q) = (\mathbb{A}\mathrm{gt}, St', \Pi, \pi', \{\sim'_a\}_{a\in\mathbb{A}\mathrm{gt}}, Act, d', o')$ where $d'(a, h)$, $o'(h, \boldsymbol{\alpha})$, and $\pi'(h)$ are given as in Definition 12 and 15 where function "*last*" is replaced with "*lastr*" and (note that $\sim^{T_s}_a$ referes to relation $\sim'_a$):

1. $St' := \bigcup_{i=0}^{\infty} \Delta^i_{\mathfrak{M}}(q)$;
2. for all $a \in \mathbb{A}\mathrm{gt}$, $\sim^{T_s}_a \subseteq St_{\mathfrak{M}} \times St_{\mathfrak{M}}$ is the smallest reflexive relation such that $h \sim^{T_s}_a h'$ if:
   (a) $rel(h) \approx^{\mathfrak{M}}_a rel(h')$, for $h, h' \in \Delta^0_{\mathfrak{M}}(q)$, or
   (b) $rel(h) \approx^{\mathfrak{M}}_a rel(h')$ and
      i. $ref(h) \sim^{T_s}_a ref(h')$ and $jump(h) = a = jump(h')$, and $h, h' \in \Delta^i_{\mathfrak{M}}(q)$, $i > 0$, or
      ii. $jump(h) = b = jump(h')$ with $a \neq b$, and $h, h' \in \Delta^i_{\mathfrak{M}}(q)$, $i > 0$, or
   (c) $h \in \Delta^i_{\mathfrak{M}}$, $h' \in \Delta^{i+1}_{\mathfrak{M}}$, $jump(h') = a$, $ref(h') \sim^{T_s}_a h$, $lastr(ref(h')) \sim^{\mathfrak{M}}_a rel(h')$ or vice versa with the roles of h and h' switched.
   We note that this means that $h' = \hat{h}\hat{a}q$, $lastr(\hat{h}) \sim^{\mathfrak{M}}_a q$, and $\hat{h} \sim^{T_s}_a h$ for some $q \in St_{\mathfrak{M}}$ and $\hat{h} \in \Delta^i_{\mathfrak{M}}(q)$.

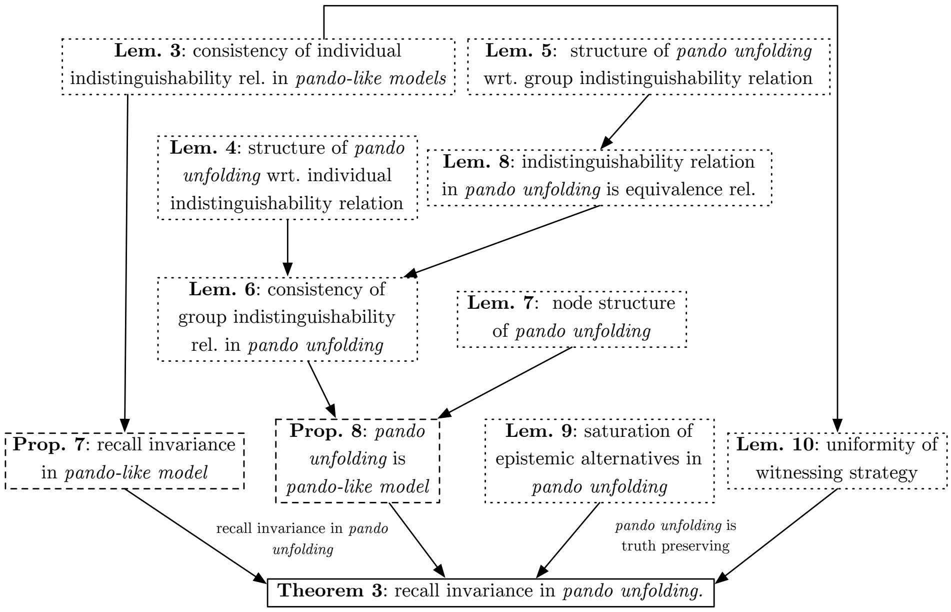

Fig. 9 Structure of the proofs of Propositions 7 and 8 and of Theorem 3. Full proofs of all results are given in Appendix A.2.

Remark 2 (i_sR-pando unfolding) We motivate points 2(a), 2(b), and 2(c) in Definition 17. Items 2(a) and (b) define indistinguishability between nodes of trees from the same set Δ^i. In this case, the "jump" must be obtained by the same epistemic relation and the final parts of the corresponding histories in the current trees (the "relevant" parts) must be indistinguishable; moreover, the "reference" nodes (in the trees one level up) must be indistinguishable for the "jump" agent (point 2(b)i.) in case we are concerned with epistemic alternatives of this very agent. This is needed to obtain transitivity of the epistemic relation in the resulting forest. Note that, in particular, the length of the relevant subhistories must be the same.

Item 2(c) defines the only way how nodes h and h' from *different* sets Δ^i and Δ^j, $i \neq j$, can be linked via an epistemic link. Firstly, it must be the case that $j = i+1$. Secondly, the relevant part of $h' \in \Delta^{i+1}$ must be a *single state* which is indistinguishable from the last state of the reference part of $h' \in \Delta^i$; moreover, the reference part of h' must also be linked to h. Note, that the relevant parts of h and h' do not have to have the same length. This models the "forgetting" if a new state-subformula is evaluated in h.

Example 9 (i_sR-pando unfolding) The i_sR-pando unfolding of model (M_1, q_0) from Figure 5 on page 22 is shown in Figure 8.

Similarly to Section 4.2 we can show that an i_sR-pando unfolding is i_sR-pando-like as expected. For example, it has to be shown that all nodes are disjunct, in order to obtain a tree-like structure, and that the epistemic relation $\sim_a^{T_s(\mathfrak{M},q)}$ is an equivalence relation for each agent $a \in \mathbb{A}\mathrm{gt}$. The proof of the following result is rather technical and is formally proven in the appendix on page 44. The structure of the proof of this proposition and also of our main result, Theorem 3, is outlined in Figure 9.

Proposition 8 *The i_sR-pando unfolding of a pointed* $\boldsymbol{i}\mathbf{CGS}$ *is i_sR-pando-like.*

Then, thanks to Proposition 7, we obtain that i_sR-pando unfoldings are truth-invariant under recall. Now we can state our main result for i_sR-pando unfoldings.

Theorem 3 *For every* $\boldsymbol{i}\mathbf{CGS}$ $\mathfrak{M}$*, state* q *in* $\mathfrak{M}$*, and* ***ATL*****-formula* φ*, it holds that*

$$\mathfrak{M}, q \models_{i_sR} \varphi \textit{ iff } T_s(\mathfrak{M}, q), q \models_{i_sR} \varphi \textit{ iff } T_s(\mathfrak{M}, q), q \models_{i_sr} \varphi.$$

Again, the proof is moved to the appendix, and can be found on page 45.

5 Comparing Validities for Variants of ATL

In this section we present a formal comparison of the semantic variants defined in Sections 2 and 3. As stated in the introduction, we compare the variants on the level of their validity sets (or, equivalently, satisfiable sentences). In most cases, they turn out to be different. Also, we can usually 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 $Val(\mathbf{ATL}_{sem})$ to denote the set of ***ATL*** *validities*, or the *theory of* ***ATL***, under semantics *sem*. Likewise, we write $Sat(\mathbf{ATL}_{sem})$ for the set of **ATL** formulae satisfiable in the semantics *sem*. Note that validity and satisfiability of formulae in all cases considered in this paper is defined over *the same class of models*, namely the class of imperfect information concurrent game structures. The conceptual reading of $Val(\mathbf{ATL}_{sem_1}) \subsetneq Val(\mathbf{ATL}_{sem_2})$ can be as follows: for "game boards" given by $\boldsymbol{i}\mathbf{CGS}$'s, we have that the "game rules" in the $\mathbf{ATL}_{sem_1}$ variant strictly refine the rules in $\mathbf{ATL}_{sem_2}$. Note also that $Val(\mathbf{ATL}_{sem_1}) \subsetneq Val(\mathbf{ATL}_{sem_2})$ is equivalent to $Sat(\mathbf{ATL}_{sem_2}) \subsetneq Sat(\mathbf{ATL}_{sem_1})$. Thus, an alternative reading is "$\mathbf{ATL}_{sem_1}$ admits reasoning about a larger variety of games than $\mathbf{ATL}_{sem_2}$".

We will always prove inclusion results for the broadest possible language (usually $\mathbf{ATL}^*$) and non-inclusion results for the narrowest one (usually **ATL**). Clearly, for languages $\mathcal{L} \subseteq \mathcal{L}'$, we have that $Val(\mathcal{L}'_{sem_1}) \subseteq Val(\mathcal{L}'_{sem_2})$ implies $Val(\mathcal{L}_{sem_1}) \subseteq Val(\mathcal{L}_{sem_2})$, and $Val(\mathcal{L}_{sem_1}) \nsubseteq Val(\mathcal{L}_{sem_2})$ implies $Val(\mathcal{L}'_{sem_1}) \nsubseteq Val(\mathcal{L}'_{sem_2})$.

Summary of the results. Figure 10 gives an overview of the results of Sections 5.3-5.6. We show that almost all the semantic variants discussed here are different on the level of validities, and that they show 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. Moreover, the type of information has more impact on the validities than the type of recall in the more restricted language of **ATL**. Interestingly, for the richer languages of $\mathbf{ATL}^+$ and $\mathbf{ATL}^*$ this is not the case anymore.

Note that if we reverse the subsumption signs in Figure 10 then the graphs describe the hierarchy of *satisfiable* sentences in different semantics of $\mathbf{ATL}/\mathbf{ATL}^*$.

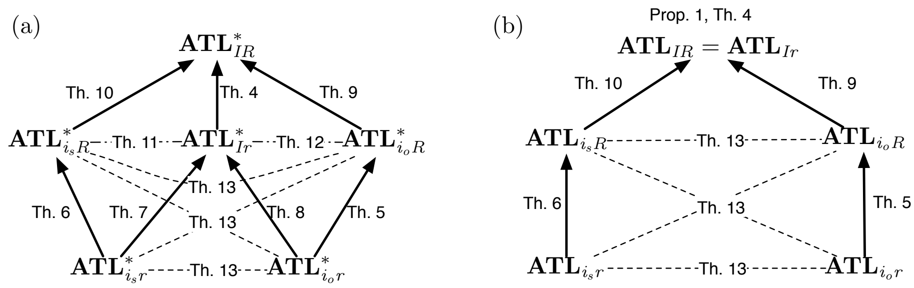

Fig. 10 Comparison of validity sets induced by various semantics of (a) $\mathbf{ATL}^*$, and (b) $\mathbf{ATL}$. Arrows depict strict subsumption of validity sets, e.g., "$\mathbf{ATL}^*_{Ir} \to \mathbf{ATL}^*_{IR}$" means that $Val(\mathbf{ATL}^*_{Ir}) \subsetneq Val(\mathbf{ATL}^*_{IR})$. Dotted lines connect semantic variants with incomparable validity sets. We do not include links that follow from transitivity of the subsumption relation. Note: the hierarchy for $\mathbf{ATL}^+$ is exactly the same as for $\mathbf{ATL}^*$.

Remark 3 It is important to observe that *comparing validities* is not the same as *comparing abilities*. For example, subjective ability to enforce γ always implies objective ability to enforce γ. Yet, as we show in Section 5.6, the set of validities for objective ability does not subsume the one for subjective ability. It is tempting to think that it should, because for every validity $\langle\langle A\rangle\rangle\gamma$ in the subjective semantics, $\langle\langle A\rangle\rangle\gamma$ must be also valid in the objective semantics. On the other hand, what about validities stating inability, i.e., $\neg\langle\langle A\rangle\rangle\gamma$? Should they adhere to the reverse subsumption? Either way, this line of reasoning is totally misleading.

The reason for that is simple. Almost *no* formulae of type $\langle\langle A\rangle\rangle\gamma$ or $\neg\langle\langle A\rangle\rangle\gamma$ are validities of $\mathbf{ATL}$ in any semantics that we study. There are only two exceptions: $\langle\langle A\rangle\rangle\top$ and $\neg\langle\langle A\rangle\rangle\bot$. Or, to be more precise, all formulae $\langle\langle A\rangle\rangle\gamma$ where γ is tautologically true (i.e., holds on all paths that can occur in any $\mathbf{CGS}$) and $\neg\langle\langle A\rangle\rangle\gamma$ where γ is tautologically false (i.e., fails on all paths in all $\mathbf{CGS}$'s). For a nontrivial ability (that is, one which refers to a temporal property that can, but does not have to be true), a valid formula can only connect it to another kind of ability. For example, $\langle\langle A\rangle\rangle\Diamond\mathsf{p} \to \langle\langle A\cup B\rangle\rangle\Diamond\mathsf{p}$ is valid in all the semantics considered in this paper.

5.1 Perfect Recall vs. Memoryless Play under Perfect Information (*IR* vs. *Ir*)

We first proceed to examine the impact of recall on the general strategic properties of agent systems under prefect information. The inclusion results follow naturally from the invariance theorems for tree-like unfoldings presented in Section 4. Non-inclusion will be demonstrated by appropriate formulae (that are valid in one semantics and not valid in another). We have already mentioned that, in $\mathbf{ATL}$, the *Ir*- and *IR*-semantics coincide (Proposition 1). As a consequence, they induce the same validities: $Val(\mathbf{ATL}_{Ir}) = Val(\mathbf{ATL}_{IR})$. 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 $\mathbf{ATL}$. An interesting question is: *Does it carry over to more general classes of winning conditions, or are there (broader) languages that*

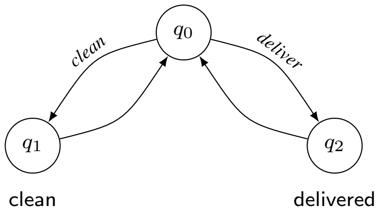

Fig. 11 Single-agent model M_5: robot with multiple tasks

can discern between the two types of ability? The answer is: no, it doesn't, and yes, there are. The *Ir*- and *IR*-semantics induce different validity sets for **ATL***, and in fact the distinction is already present in **ATL**$^+$. Moreover, it turns out that perfect recall can be seen as a special case of memoryless play in the sense of their general properties.

Proposition 9 $Val(\boldsymbol{ATL}^*_{Ir}) \subseteq Val(\boldsymbol{ATL}^*_{IR})$

Proof. Let an **ATL***-formula φ be *Ir*-valid in $\mathbf{iCGS}$'s, then it is also *Ir*-valid in tree-like $\mathbf{CGS}$'s, and by Proposition 2 also *IR*-valid in tree-like $\mathbf{CGS}$'s. Thus, by Theorem 1, it is *IR*-valid in arbitrary $\mathbf{CGS}$'s. Since indistinguishability relations do not influence the semantic relation $\models_{IR}$, we get that φ is *IR*-valid in $\mathbf{iCGS}$'s. ■

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

Proposition 10 $Val(\boldsymbol{ATL}^+_{IR}) \not\subseteq Val(\boldsymbol{ATL}^+_{Ir})$.

Proof. Consider formula

$$\Phi_3 \equiv \langle\!\langle a \rangle\!\rangle(\Diamond \mathsf{p}_1 \land \Diamond \mathsf{p}_2) \leftrightarrow \langle\!\langle a \rangle\!\rangle \Diamond(\mathsf{p}_1 \land \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}_2 \lor \mathsf{p}_2 \land \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}_1).$$

The formula is valid in $\mathbf{ATL}^+_{IR}$ [33]. On the other hand, its right-to-left part is not valid in $\mathbf{ATL}^+_{Ir}$. To see this, we take the single-agent $\mathbf{CGS}$ M_5 from Figure 11 where agent a (the robot) can do the cleaning or deliver a package. Then, for $\mathsf{p}_1 \equiv \mathsf{clean}$, $\mathsf{p}_2 \equiv \mathsf{delivered}$, we have $M_5, q_0 \models_{Ir} \langle\!\langle a \rangle\!\rangle \Diamond(\mathsf{p}_1 \land \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}_2 \lor \mathsf{p}_2 \land \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}_1)$ but also $M_5, q_0 \not\models_{Ir} \langle\!\langle a \rangle\!\rangle(\Diamond \mathsf{p}_1 \land \Diamond \mathsf{p}_2)$. ■

Theorem 4 $Val(\boldsymbol{ATL}_{Ir}) = Val(\boldsymbol{ATL}_{IR})$. *However,* $Val(\boldsymbol{ATL}^+_{Ir}) \subsetneq Val(\boldsymbol{ATL}^+_{IR})$ *and* $Val(\boldsymbol{ATL}^*_{Ir}) \subsetneq Val(\boldsymbol{ATL}^*_{IR})$.

Proof. From Proposition 1 it follows that $Val(\mathbf{ATL}_{Ir}) = Val(\mathbf{ATL}_{IR})$. From Proposition 9 we know that $Val(\mathbf{ATL}^*_{Ir}) \subseteq Val(\mathbf{ATL}^*_{IR})$ and can also deduce that $Val(\mathbf{ATL}^+_{Ir}) \subseteq Val(\mathbf{ATL}^+_{IR})$ because the language of **ATL**$^+$ is just a syntactic restriction of the one of **ATL***. Finally, Proposition 10 proves that $Val(\mathbf{ATL}^+_{Ir}) \subsetneq Val(\mathbf{ATL}^+_{IR})$ and also that $Val(\mathbf{ATL}^*_{Ir}) \subsetneq Val(\mathbf{ATL}^*_{IR})$ because the formula given in the proof of the very proposition is in particular also an **ATL***-formula. ■

5.2 Perfect Recall vs. Memoryless Play under Imperfect Information (*iR* vs. *ir*)

Now we compare the memoryless and perfect recall semantics under uncertainty. We treat the case of objective and subjective ability separately.

5.2.1 Imperfect Information: Objective Ability

Proposition 11 $Val(\boldsymbol{ATL}^*_{i_or}) \subseteq Val(\boldsymbol{ATL}^*_{i_oR})$.

Proof. We prove that $Sat(\mathbf{ATL}^*_{i_oR}) \subseteq Sat(\mathbf{ATL}^*_{i_or})$. Let $\varphi \in Sat(\mathbf{ATL}^*_{i_oR})$. Then, there must be a pointed **iCGS** $(\mathfrak{M}, q)$ such that $\mathfrak{M}, q \models_{i_oR} \varphi$. By Theorem 2, $T_o(\mathfrak{M}, q), q \models_{i_oR} \varphi$. But on i_oR-tree unfoldings, *iR*- and *ir*-strategies coincide (Theorem 2), so we get that $T_o(\mathfrak{M}, q), q \models_{i_or} \varphi$, and as a consequence $\varphi \in Sat(\mathbf{ATL}^*_{i_or})$. ∎

The converse does not hold:

Proposition 12 $Val(\boldsymbol{ATL}_{i_oR}) \not\subseteq Val(\boldsymbol{ATL}_{i_or})$

Proof. To show this, we take the **ATL** embedding of the **CTL** duality between combinators $\mathsf{E}\Box$ and $\mathsf{A}\Diamond$ (see Section 2.5). In fact, only one direction of the equivalence is important here:

$$\Phi_4 \equiv \neg\langle\langle\emptyset\rangle\rangle\Diamond\neg\mathsf{p} \rightarrow \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Box\mathsf{p}$$

(note that the other direction is valid for all the semantics considered in this paper, and actually for 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 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 $\mathbf{ATL}_{i_oR}$.

Second, we consider model M_6 in Figure 12.[10] Let us take $\mathsf{p} \equiv \neg\mathsf{angry} \wedge \neg\mathsf{suspicious}$. Then, we have $M_6, q_0 \models_{i_or} \neg\langle\langle\emptyset\rangle\rangle\Diamond\neg\mathsf{p}$ but also $M_6, q_0 \not\models_{i_or} \langle\langle\mathbb{A}\mathrm{gt}\rangle\rangle\Box\mathsf{p}$, which demonstrates that Φ_4 is not valid in $\mathbf{ATL}_{i_or}$. ∎

Theorem 5 *$Val(\boldsymbol{ATL}_{i_or}) \subsetneq Val(\boldsymbol{ATL}_{i_oR})$, and similarly for $\boldsymbol{ATL}^+$ and $\boldsymbol{ATL}^*$.*

[10] 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 ($\lambda = q_0q_1q_1q_1\dots$).

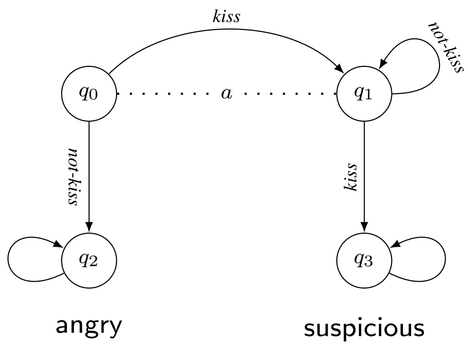

Fig. 12 Model M_6 with $\mathbb{A}\mathrm{gt} = \{a\}$: dangers of marital life

5.2.2 *Imperfect Information: Subjective Ability*

Proposition 13 $Val(\boldsymbol{ATL}^*_{i_s r}) \subseteq Val(\boldsymbol{ATL}^*_{i_s R})$.

Proof. Analogous to Proposition 11. ∎

Proposition 14 $Val(\boldsymbol{ATL}_{i_s R}) \not\subseteq Val(\boldsymbol{ATL}_{i_s r})$.

Proof. We take the formula Φ_5 which is a consequence of the fixpoint equivalence for $\langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$:

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

The formula states 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 $\mathbf{ATL}_{i_s R}$, 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 $M_1, q_0 \models_{i_s r} \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$ but also $M_1, q_0 \not\models_{i_s r} \langle\!\langle a \rangle\!\rangle \Diamond \mathsf{p}$, so Φ_5 is not valid in $\mathbf{ATL}_{i_s r}$. ∎

Theorem 6 *$Val(\boldsymbol{ATL}_{i_s r}) \subsetneq Val(\boldsymbol{ATL}_{i_s R})$, and similarly for $\boldsymbol{ATL}^+$ and $\boldsymbol{ATL}^*$.*

5.3 Perfect vs. Imperfect Information under Memoryless Play (*Ir* vs. *ir*)

We continue by comparing 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_s/i_o), uncertainty of agents about states constrains their choices. Firstly, we observe that perfect information can be seen as a special case of imperfect information.

Proposition 15 $Val(\boldsymbol{ATL}^*_{i_s r}) \subseteq Val(\boldsymbol{ATL}^*_{Ir})$ *and* $Val(\boldsymbol{ATL}^*_{i_o r}) \subseteq Val(\boldsymbol{ATL}^*_{Ir})$.

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

Proposition 16 *$Val(\boldsymbol{ATL}_{Ir}) \nsubseteq Val(\boldsymbol{ATL}_{i_s r})$.*

Proof. We show this by presenting a validity for $\mathbf{ATL}_{Ir}$ which is not valid in $\mathbf{ATL}_{i_s r}$. Consider the formula that captures the right-to-left direction in the fixpoint characterization of $\langle\langle A\rangle\rangle\Diamond\varphi$ for single-agent teams and atomic propositions:

$$\Phi_1 \equiv (\mathsf{p} \vee \langle\langle a\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\Diamond\mathsf{p}) \rightarrow \langle\langle a\rangle\rangle\Diamond\mathsf{p}$$

Φ_1 is *Ir*-valid (cf. Section 2.5). To see its invalidity in the $i_s r$-semantics, consider model M_1 from Figure 5 on page 22. We recall that the story behind the model is as follows. A man wants to shoot down a yellow rubber duck in a shooting gallery. 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). Can he identify an opening strategy that 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.

Indeed, for $\mathsf{p} \equiv \mathsf{shot}$, we get $M_1, q_0 \models_{i_s r} \mathsf{p} \vee \langle\langle a\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\Diamond\mathsf{p}$ and $M_1, q_0 \not\models_{i_s r} \langle\langle a\rangle\rangle\Diamond\mathsf{p}$, which formally concludes our proof. ∎

Proposition 17 *$Val(\boldsymbol{ATL}_{Ir}) \nsubseteq Val(\boldsymbol{ATL}_{i_o r})$.*

Proof. It is sufficient to show that $\Phi_1 \equiv (\mathsf{p} \vee \langle\langle a\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\Diamond\mathsf{p}) \rightarrow \langle\langle a\rangle\rangle\Diamond\mathsf{p}$ is invalid in the $i_o r$-semantics. Take model M_2 in Figure 13 and $\mathsf{p} \equiv \mathsf{shot}$. Now we have that $M_2, q_0' \models_{i_o r} \mathsf{p} \vee \langle\langle a\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\Diamond\mathsf{p}$ because a has a uniform strategy that *objectively* achieves $\Diamond\mathsf{p}$ in q_0 ($s_a(q) = shoot_L$ for every q) and another uniform strategy in q_1 ($s_a'(q) = shoot_R$ for every q). However, s_a and s_a' cannot be merged into a single uniform strategy, and indeed $M_2, q_0' \not\models_{i_o r} \langle\langle a\rangle\rangle\Diamond\mathsf{p}$, which concludes the proof. ∎

Note that, for $\mathbf{ATL}_{i_o r}$, formula Φ_1 *is* valid in single-agent models, so we really needed to add another agent to the picture.

The following theorems are straightforward consequences.

Theorem 7 *$Val(\boldsymbol{ATL}_{i_s r}) \subsetneq Val(\boldsymbol{ATL}_{Ir})$, $Val(\boldsymbol{ATL}^+_{i_s r}) \subsetneq Val(\boldsymbol{ATL}^+_{Ir})$, and $Val(\boldsymbol{ATL}^*_{i_s r}) \subsetneq Val(\boldsymbol{ATL}^*_{Ir})$.*

Theorem 8 *$Val(\boldsymbol{ATL}_{i_o r}) \subsetneq Val(\boldsymbol{ATL}_{Ir})$, and similarly for $\boldsymbol{ATL}^+$ and $\boldsymbol{ATL}^*$.*

By Proposition 1 and Theorems 4, 7, and 8, we get the following corollary:

Corollary 1 *$Val(\boldsymbol{ATL}_{i_s r}) \subsetneq Val(\boldsymbol{ATL}_{IR})$ and $Val(\boldsymbol{ATL}_{i_o r}) \subsetneq Val(\boldsymbol{ATL}_{IR})$, and similarly for $\boldsymbol{ATL}^+$ and $\boldsymbol{ATL}^*$.*

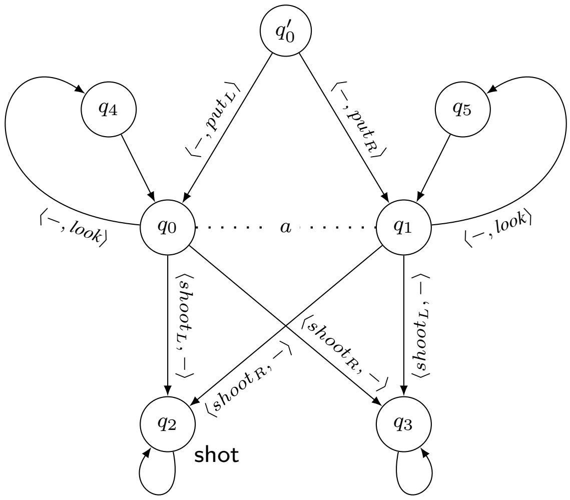

Fig. 13 Modified "poor duck" model M_2 with two agents a, b. This time, we explicitly represent the agent (b) who puts the duck in one of the cells.

5.4 Perfect vs. Imperfect Information under Perfect Recall Play (*IR* vs. *iR*)

First, we observe that for $\mathbf{ATL}_{i_oR}$ vs. $\mathbf{ATL}_{IR}$ we can employ the same reasoning as for for $\mathbf{ATL}_{i_or}$ vs. $\mathbf{ATL}_{Ir}$. 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 in $\mathbf{ATL}_{i_oR}$. Thus, analogously to Theorem 8 we get:

Theorem 9 *$Val(\mathbf{ATL}_{i_oR}) \subsetneq Val(\mathbf{ATL}_{IR})$, and similarly for $\mathbf{ATL}^+$ and $\mathbf{ATL}^*$.*

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

Proposition 18 *$Val(\mathbf{ATL}_{IR}) \not\subseteq Val(\mathbf{ATL}_{i_sR})$.*

Proof. This time we consider the other direction of the fixpoint characterization for $\langle\langle a\rangle\rangle\Diamond \mathsf{p}$:

$$\Phi_2 \equiv \langle\langle a\rangle\rangle\Diamond \mathsf{p} \rightarrow (\mathsf{p} \vee \langle\langle a\rangle\rangle\bigcirc \langle\langle a\rangle\rangle\Diamond \mathsf{p}).$$

Φ_2 is *IR*-valid, but it is not valid in i_sR. Consider a modification of the "poor duck model" in Figure 14 (the party goes on, and the man is not even able to reach out and look anymore; the cells are open initially but they will close in a moment). Take $\mathsf{p} \equiv \mathsf{shot}$. We have that $M_3, q_4 \models_{i_sR} \langle\langle a\rangle\rangle\Diamond \mathsf{p}$, but $M_3, q_4 \not\models_{i_sR} \mathsf{p} \vee \langle\langle a\rangle\rangle\bigcirc \langle\langle a\rangle\rangle\Diamond \mathsf{p}$, which concludes the proof. ∎

Theorem 10 *$Val(\mathbf{ATL}_{i_sR}) \subsetneq Val(\mathbf{ATL}_{IR})$, and similarly for $\mathbf{ATL}^+$ and $\mathbf{ATL}^*$.*

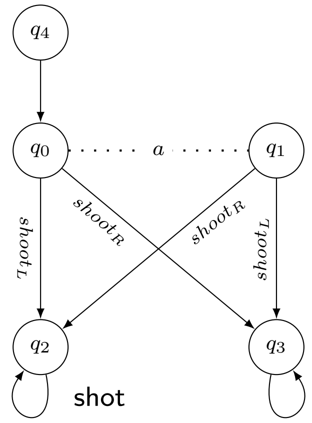

Fig. 14 Variant of "poor duck" after a particularly long party (model M_3)

5.5 Mixed Setting: Information vs. Memory (*Ir* vs. *iR*)

In this section we compare abilities if both dimensions change. For **ATL** we already know the complete picture because $\mathbf{ATL}_{Ir}$ and $\mathbf{ATL}_{IR}$ are the same logics, cf. Figure 10(b). For $\mathbf{ATL}^*$ it remains to compare $\mathbf{ATL}^*_{Ir}$, $\mathbf{ATL}^*_{i_oR}$, and $\mathbf{ATL}^*_{i_sr}$.

To facilitate proofs, we define an additional temporal operator N ("now") as $\mathsf{N}\varphi \equiv \varphi\,\mathcal{U}\,\varphi$. Note that $\mathfrak{M},\lambda \models \mathsf{N}\varphi$ iff $\mathfrak{M},\lambda \models \varphi$ in the semantics of $\mathbf{CTL}^*$ and any **ATL**-semantics that we have discussed in this paper. Moreover, we note that the formula $\langle\!\langle A\rangle\!\rangle\mathsf{N}\varphi$ expresses $E_A\varphi$ (everybody in A knows that φ) if $\langle\!\langle A\rangle\!\rangle$ is interpreted according to the subjective semantics for imperfect information (i.e., according to $\models_{i_sR}$ and $\models_{i_sr}$).

Theorem 11 *The sets* $Val(\boldsymbol{ATL}^*_{i_sR})$ *and* $Val(\boldsymbol{ATL}^*_{Ir})$ *are incomparable, and similarly for* $\boldsymbol{ATL}^+$.

Proof. We prove incomparability for $\mathbf{ATL}^+$. From this, incomparability for $\mathbf{ATL}^*$ follows immediately.

1. $Val(\mathbf{ATL}^*_{Ir}) \not\subseteq Val(\mathbf{ATL}^*_{i_sR})$. Suppose that $Val(\mathbf{ATL}^*_{Ir}) \subseteq Val(\mathbf{ATL}^*_{i_sR})$. This implies that $Val(\mathbf{ATL}_{Ir}) \subseteq Val(\mathbf{ATL}_{i_sR})$ and by Theorem 4 $Val(\mathbf{ATL}_{IR}) \subseteq Val(\mathbf{ATL}_{i_sR})$. But this contradicts Theorem 10.
2. $Val(\mathbf{ATL}^*_{i_sR}) \not\subseteq Val(\mathbf{ATL}^*_{Ir})$. For this case we consider the $\mathbf{ATL}^+$-formula

$$\Phi_6 = \langle\!\langle\mathbb{A}\mathrm{gt}\rangle\!\rangle\Diamond\langle\!\langle\mathbb{A}\mathrm{gt}\rangle\!\rangle\mathsf{N}(\mathsf{p}_1 \wedge \langle\!\langle\mathbb{A}\mathrm{gt}\rangle\!\rangle\Diamond\mathsf{p}_2) \rightarrow \langle\!\langle\mathbb{A}\mathrm{gt}\rangle\!\rangle(\Diamond\mathsf{p}_1 \wedge \Diamond\mathsf{p}_2)$$

which is a validity of $\mathbf{ATL}^+_{i_sR}$ but not of $\mathbf{ATL}^+_{Ir}$. The latter fact can be shown by the same counterexample as used in the proof of Proposition 10 (we have $M_5, q_0 \not\models_{Ir} \Phi_6$).
It remains to show that $\Phi_6 \in Val(\mathbf{ATL}^+_{i_sR})$. Suppose that $\mathfrak{M}, q \models_{i_sR} \langle\!\langle\mathbb{A}\mathrm{gt}\rangle\!\rangle\Diamond\langle\!\langle\mathbb{A}\mathrm{gt}\rangle\!\rangle\mathsf{N}(\mathsf{p}_1 \wedge \langle\!\langle\mathbb{A}\mathrm{gt}\rangle\!\rangle\Diamond\mathsf{p}_2)$. That is, there is an *iR*-strategy

$s_{\mathbb{A}\mathrm{gt}}$ such that for all $\lambda \in out^{i_s}(q, s_{\mathbb{A}\mathrm{gt}})$ there is an $i \geq 0$ and an *iR*-strategy $s'_{\mathbb{A}\mathrm{gt}}$ such that for all $\lambda' \in out^{i_s}(\lambda[i], s'_{\mathbb{A}\mathrm{gt}})$ we have that $\lambda'[0] \models \mathsf{p}_1$ and there is an *iR*-strategy $s''_{\mathbb{A}\mathrm{gt}}$ such that for all $\lambda'' \in out^{i_s}(\lambda'[0], s''_{\mathbb{A}\mathrm{gt}})$ it holds that $\mathfrak{M}, \lambda'' \models \Diamond \mathsf{p}_2$. Because we have that $\{q' \mid q' \sim_{\mathbb{A}\mathrm{gt}} \lambda[i]\} = \{q' \mid q' \sim_{\mathbb{A}\mathrm{gt}} \lambda'[0]\}$, we can take $s'_{\mathbb{A}\mathrm{gt}}$ as $s''_{\mathbb{A}\mathrm{gt}}$. Then, we have that $\mathfrak{M}, q \models_{i_sR} \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle \Diamond \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle \mathsf{N}(\mathsf{p}_1 \wedge \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle \Diamond \mathsf{p}_2)$ iff there is an *iR*-strategy $s_{\mathbb{A}\mathrm{gt}}$ such that for all $\lambda \in out^{i_s}(q, s_{\mathbb{A}\mathrm{gt}})$ there is an $i \geq 0$ and an *iR*-strategy $s'_{\mathbb{A}\mathrm{gt}}$ such that for all $\lambda' \in out^{i_s}(\lambda[i], s'_{\mathbb{A}\mathrm{gt}})$ we have that $\lambda'[0] \models \mathsf{p}_1$ and $\mathfrak{M}, \lambda' \models \Diamond \mathsf{p}_2$.
Now, it is easy to see that we can combine $s_{\mathbb{A}\mathrm{gt}}$ and each $s'_{\mathbb{A}\mathrm{gt}}$ to a single strategy $\hat{s}_{\mathbb{A}\mathrm{gt}}$ such that for all $\lambda \in out^{i_s}(q, \hat{s}_{\mathbb{A}\mathrm{gt}})$ it holds that $\mathfrak{M}, \lambda \models \Diamond \mathsf{p}_1 \wedge \Diamond \mathsf{p}_2$. This shows that $\mathfrak{M}, q \models \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle (\Diamond \mathsf{p}_1 \wedge \Diamond \mathsf{p}_2)$.

■

Apart from minor modifications, the next theorem, considering objective ability, is proven along the same lines.

Theorem 12 *The sets $Val(\boldsymbol{ATL}^*_{i_oR})$ and $Val(\boldsymbol{ATL}^*_{Ir})$ are incomparable, and similarly for $\boldsymbol{ATL}^+$.*

Proof. Again, we prove incomparability for $\mathbf{ATL}^+$. From this, incomparability for $\mathbf{ATL}^*$ follows immediately.

1. $Val(\mathbf{ATL}^*_{Ir}) \not\subseteq Val(\mathbf{ATL}^*_{i_oR})$. Suppose that $Val(\mathbf{ATL}^*_{Ir}) \subseteq Val(\mathbf{ATL}^*_{i_oR})$. This implies that $Val(\mathbf{ATL}_{Ir}) \subseteq Val(\mathbf{ATL}_{i_oR})$ and by Theorem 4 $Val(\mathbf{ATL}_{IR}) \subseteq Val(\mathbf{ATL}_{i_oR})$. But this contradicts Theorem 9.
2. $Val(\mathbf{ATL}^*_{i_oR}) \not\subseteq Val(\mathbf{ATL}^*_{Ir})$. For this case we consider the $\mathbf{ATL}^*$-formula
$$\varphi = \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle \Diamond (\mathsf{p}_1 \wedge \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle \Diamond \mathsf{p}_2) \rightarrow \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle (\Diamond \mathsf{p}_1 \wedge \Diamond \mathsf{p}_2)$$
which is a validity of $\mathbf{ATL}^*_{i_oR}$ but not of $\mathbf{ATL}^*_{Ir}$. The latter is shown by the same counterexample as used in the proof of Proposition 10 (we have $M_5, q_0 \not\models_{Ir} \varphi$).
Finally, it remains to show that $\varphi \in Val(\mathbf{ATL}^*_{i_oR})$. This part is proven following the same idea as in the proof of Theorem 11. We observe that every strategy $s_{\mathbb{A}\mathrm{gt}}$ of the grand coalition generates a *unique* path wrt. objective ability (because, in the objective semantics, possible paths starting from epistemic alternatives are not considered). This also means that uniformity of a strategy does not matter: there is no need to ever consider epistemic alternatives along a path. Hence, the two strategies witnessing $\langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle \Diamond (\mathsf{p}_1 \wedge \langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle \Diamond \mathsf{p}_2)$ can be combined to a single strategy witnessing $\langle\!\langle \mathbb{A}\mathrm{gt} \rangle\!\rangle (\Diamond \mathsf{p}_1 \wedge \Diamond \mathsf{p}_2)$.

■

5.6 Between Subjective and Objective Ability for Imperfect Information (i_s vs. i_o)

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 19 *Formula $\Phi_2 \equiv \langle\langle a\rangle\rangle\Diamond \mathsf{p} \to \mathsf{p} \vee \langle\langle a\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\Diamond \mathsf{p}$ is valid in* $\mathbf{ATL}_{i_oR}$ *and* $\mathbf{ATL}_{i_or}$*, but invalid in* $\mathbf{ATL}_{i_sR}$ *and* $\mathbf{ATL}_{i_sr}$*.*

Proof. We first prove validity of Φ_2 in $\mathbf{ATL}_{i_or}$, which implies also validity in $\mathbf{ATL}_{i_oR}$ by Proposition 11. Suppose that $\mathfrak{M}, q \models_{i_or} \langle\langle a\rangle\rangle\Diamond \mathsf{p}$, 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 $\mathbf{ATL}_{i_sR}$ was already proved in Proposition 18. Thus, by Proposition 13, Φ_2 is not valid in $\mathbf{ATL}_{i_sr}$, too. ■

In the next result we make use of the operator N introduced in Section 5.5.

Proposition 20 *Formula*

$$\Phi_8 \equiv \langle\langle a\rangle\rangle \mathsf{N} \langle\langle c\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\bigcirc \mathsf{p} \to \langle\langle a,c\rangle\rangle\Diamond \mathsf{p}$$

is valid in $\mathbf{ATL}_{i_sR}$ *and* $\mathbf{ATL}_{i_sr}$*, but invalid in* $\mathbf{ATL}_{i_oR}$ *and* $\mathbf{ATL}_{i_or}$*.*

Proof. Analogously to Proposition 19, we prove the validity of Φ_8 in $\mathbf{ATL}_{i_sr}$, and its invalidity in $\mathbf{ATL}_{i_oR}$.

First, let $\mathfrak{M}, q \models_{i_sr} \langle\langle a\rangle\rangle \mathsf{N} \langle\langle c\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\bigcirc \mathsf{p}$. Then, for every state $q' \in [q]_{\sim_a}$, c has an *ir*-strategy $s_c^{q'}$ that enforces $\bigcirc\langle\langle a\rangle\rangle\bigcirc \mathsf{p}$ from $[q']_{\sim_c}$. By combining all these strategies into an *ir*-strategy s_c (we can do it since the startegies $s_c^{q'}$ 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 $[\lambda[1]]_{\sim_a}$. 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 $\mathfrak{M}, q \models_{i_sr} \langle\langle a,c\rangle\rangle\Diamond \mathsf{p}$.

For the invalidity, consider the modified poor duck model M_2 from Figure 13 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_3', and let $\mathsf{p} \equiv \mathsf{shot}$. It is easy to see that $M_2', q_0' \models_{i_oR} \langle\langle c\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\bigcirc \mathsf{p}$, and hence also $M_2', q_0' \models_{i_oR} \langle\langle a\rangle\rangle \mathsf{N} \langle\langle c\rangle\rangle\bigcirc\langle\langle a\rangle\rangle\bigcirc \mathsf{p}$. On the other hand, $M_2', q_0' \not\models_{i_oR} \langle\langle a,c\rangle\rangle\Diamond \mathsf{p}$, which concludes the proof. ■

The following is an immediate consequence.

Theorem 13 *For every $y, z \in \{R, r\}$, the sets Val($\mathbf{ATL}_{i_sy}$) and Val($\mathbf{ATL}_{i_oz}$) are incomparable, and similarly for* $\mathbf{ATL}^+$ *and* $\mathbf{ATL}^*$

6 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 10 on page 30. The first and most important conclusion is that almost all the semantic variants discussed here are indeed different on the level of general properties they induce; before our study, it was by no means obvious. Moreover, our results show a very 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.

The relationships seem very natural, but they were surprisingly nontrivial to prove. This is best witnessed by Section 4 which comprises a third part of the paper only to construct appropriate tree-like unfoldings, and prove their equivalence to the original models. While embedding of perfect information in imperfect information is straightforward, the same cannot be said about embedding perfect recall in memoryless semantics – except when we disallow nested modalities. Consider e.g. the truth of formula $\langle\langle a\rangle\rangle\Box\langle\langle a,b\rangle\rangle\Diamond \mathsf{p}$ in a pointed **iCGS** $(\mathfrak{M}, q)$. Let s_a be a's strategy that enforces $\langle\langle a,b\rangle\rangle\Diamond \mathsf{p}$ to be always the case (suppose that such a strategy exists). After a history h, agent a has different information when executing s_a (because the agent has collected observations along h from the root until now) than when we evaluate $\langle\langle a,b\rangle\rangle\Diamond \mathsf{p}$ in the last state of h (here, the collecting of observations starts anew). In consequence, the "straightforward" unfolding of $(\mathfrak{M}, q)$ endows agents with too much information when nested strategic formulas are evaluated, and the correctness of the construction is not automatic. For objective abilities, we prove that the standard unfolding still works because path formulae of **ATL*** (that can be seen as "winning conditions" in the corresponding game) do not discern between the two epistemic positions. For subjective abilities, the unfolding does *not* work, but it can be recovered by a technical construction with "hanging" subtrees added to the basic tree. This construction is among the main contributions of this paper.

Technical subtleties aside, the most interesting contribution lies possibly in our *non-inclusion* results. First, they show that the language of **ATL** is sufficiently expressive to distinguish between the main notions of ability. Moreover, non-inclusion is demonstrated on formulae encoding intuitive and well known properties, like fixpoint characterizations of strategic/temporal modalities and the duality between necessary and obtainable outcomes. It is important to see in which semantics the formulae hold, and in which they do not hold. Finally, although the proofs of non-inclusion are very comprehensive (since they are based on counterexamples), finding the counterexamples required expertise and was not straightforward either.

Another interesting outcome of the study is that the type of information has strictly more impact on validities than the type of recall in the language of **ATL** *but not in* ***ATL****. In particular the validity sets of **ATL**$^*_{Ir}$ and **ATL**$^*_{iR}$ are incomparable. This suggests that **ATL*** allows to specify significantly subtler properties of strategic play than the more restricted language of **ATL**.

Acknowledgements. The authors would like to express their thanks to Jan Broersen for the argument that triggered this study, Valentin Goranko for discussions on (non-) validity of fixpoint properties in various semantics of **ATL**, and anonymous reviewers of EUMAS'10, IJCAI'11 and JAAMAS for their valuable remarks. Wojciech Jamroga acknowledges the support of the FNR (National Research Fund) Luxembourg under projects S-GAMES – C08/IS/03 and GALOT – INTER/DFG/12/06.

A Proofs

A.1 Proofs of Section 4.2

The following Lemma is obvious by the definition of i_oR-tree unfoldings. It states that that nodes group indistinguishable in the tree unfolding are also group indistinguishable in the model if interpreted as histories.

Lemma 1 *Let $\mathfrak{M}$ be an* $\boldsymbol{i}\mathbf{CGS}$*, h_1 and h_2 two nodes in its i_oR-tree unfolding, and $A \subseteq \mathbb{A}\mathrm{gt}$ a group of agents. If $h_1 \sim_A^{T_o(\mathfrak{M},q)} h_2$ then $h_1 \approx_A^{\mathfrak{M}} h_2$.*

Moreover, we have that all histories indistinguishable in the model are also indistinguishable in the tree if only states reachable from the current state are considered.

Lemma 2 *Let $\mathfrak{M}$ be an* $\boldsymbol{i}\mathbf{CGS}$*, $A \subseteq \mathbb{A}\mathrm{gt}$ and h a node in $T_o(\mathfrak{M}, q)$. Then, for all $h_1, h_2 \in \Lambda_{\mathfrak{M}}(last(h))$ we have that*

$$\textit{if } h_1 \approx_A^{\mathfrak{M}} h_2 \textit{ then } h(h_1[1,\infty]) \sim_A^{T_o(\mathfrak{M},q)} h(h_2[1,\infty]).$$

Theorem 2 ($\rightsquigarrow$ page 21). *For every* $\boldsymbol{i}\mathbf{CGS}$ *$\mathfrak{M}$, state q in $\mathfrak{M}$, and* ***ATL******-formula φ we have:*

$$\mathfrak{M}, q \models_{i_oR} \varphi \textit{ iff } T_o(\mathfrak{M}, q), q \models_{i_oR} \varphi \textit{ iff } T_o(\mathfrak{M}, q), q \models_{i_or} \varphi.$$

Proof. We show that, for every node h in $T_o(\mathfrak{M}, q)$, it holds that $\mathfrak{M}, last(h) \models_{i_oR} \varphi$ iff $T_o(\mathfrak{M}, q), h \models_{i_or} \varphi$. Then, the claim follows by Propositions 4 and 5 and for $h = q$. The proof is done by induction on the structure of φ.
<u>Base cases</u>:
Propositional case: Straightforward.
Case: $\varphi \equiv \langle\!\langle A\rangle\!\rangle\gamma$ where γ contains no nested strategic modalities.

"$\Rightarrow$": Suppose that $\mathfrak{M}, last(h) \models_{i_oR} \langle\!\langle A\rangle\!\rangle\gamma$. So, there is an i_oR-strategy s_A such that

$$(\star)\ \forall\lambda \in out_{\mathfrak{M}}(last(h), s_A) : \mathfrak{M}, \lambda \models_{i_oR} \gamma.$$

We construct the memoryless strategy s'_A in $T_o(\mathfrak{M}, q)$ as follows: $s'_a(\hat{h}h') = s_a(last(h)h')$ for every $a \in A$ and $\hat{h}$ such that $h \sim_A^{T_o(\mathfrak{M},q)} \hat{h}$. For all other histories h'' (which do not have the form $\hat{h}h'$) we define $s'_a(h'')$ arbitrarily but in a uniform way. It is easy to see that s'_A is uniform: For two histories $h_1 = \hat{h}'h'$ and $h_2 = \hat{h}''h''$ with $\hat{h}' \sim_A^{T_o(\mathfrak{M},q)} h$ and $\hat{h}'' \sim_A^{T_o(\mathfrak{M},q)} h$ and $h_1 \sim_A^{T_o(\mathfrak{M},q)} h_2$ we have $s'_A(h_1) = s'_A(h_2)$; for, $h_1 \sim_A^{T_o(\mathfrak{M},q)} h_2$ implies $h_1 \approx_A^{\mathfrak{M}} h_2$ (by Lemma 1) and thus $s_A(last(h)h') = s_A(last(h)h'')$.
By construction of s'_A we have that
$last(h)q_1q_2 \cdots \in out_{\mathfrak{M}}(last(h), s_A)$ iff $(h)(hq_1)(hq_1q_2) \cdots \in out_{T_o(\mathfrak{M},q)}(h, s'_A)$.
Since the valuation of propositions does only depend on the final state of a history and by $(\star)$ we have $T_o(\mathfrak{M}, q), h \models_{i_or} \langle\!\langle A\rangle\!\rangle\gamma$.

$\Leftarrow$: Suppose we have $T_o(\mathfrak{M}, q), h \models_{i_or} \langle\!\langle A\rangle\!\rangle\gamma$. So there is an i_or-strategy s_A such that

$$(\star)\ \forall\lambda \in out_{T_o(\mathfrak{M},q)}(h, s_A) : T_o(\mathfrak{M}, q), \lambda \models_{i_or} \gamma.$$

We construct a witnessing i_oR-strategy s'_A in $\mathfrak{M}$ as follows: $s'_a(\hat{h}) = s_a(hh')$ for every $a \in A$ and $\hat{h}$ such that $last(h)h' \approx_a^{\mathfrak{M}} \hat{h}$ and $last(h)h' \in \Lambda_{\mathfrak{M}}^{fin}(last(h))$. We define s'_a arbitrarily for all other histories with the condition to assign the same actions to indistinguishable histories in $\mathfrak{M}$. The definition of s'_a does only take into account the subtree starting at h. Then, by Lemma 2 we have that strategy s'_A is uniform by construction. Note, that it may differ from s_A but only for histories which are not realizable given that the initial state is $last(h)$.
By construction of s'_A, we also have

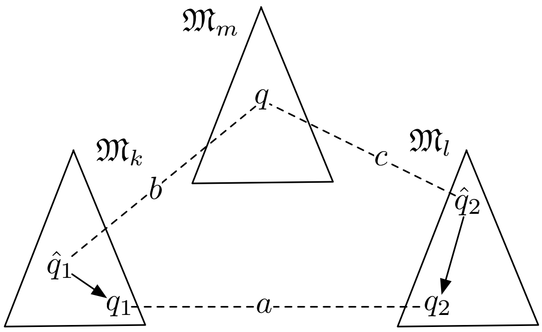

Fig. 15 General setting of the proof of Proposition 7.

$$(h)(hq_1)(hq_1q_2)\cdots \in out_{T_o(\mathfrak{M},q)}(h, s_A) \text{ iff } last(h)q_1q_2\cdots \in out_{\mathfrak{M}}(last(h), s'_A).$$

Since the valuation of propositions does only depend on the final state of a history and by $(\star)$ we have $\mathfrak{M}, last(h) \models_{i_oR} \langle\langle A\rangle\rangle\gamma$.

Induction step:

Case: $\varphi \equiv \psi_1 \wedge \psi_2$. Straightforward.

Case: $\varphi \equiv \neg\psi$. $\mathfrak{M}, last(h) \models_{i_oR} \neg\psi$ iff not $\mathfrak{M}, last(h) \models_{i_oR} \psi$ iff (by induction hypothesis) not $T_o(\mathfrak{M}, q), h \models_{i_oR} \psi$ iff $T_o(\mathfrak{M}, q), h \models_{i_oR} \neg\psi$.

Case: $\varphi \equiv \langle\langle A\rangle\rangle\gamma$. By induction hypothesis we have for each history h in $T_o(\mathfrak{M}, q)$ and each strict state-subformula φ' of γ that $\mathfrak{M}, last(h) \models_{i_oR} \varphi'$ iff $T_o(\mathfrak{M}, q), h \models_{i_or} \varphi'$. For any maximal strict subformula φ' of φ we label all states h in $T_o(\mathfrak{M}, q)$ and states $last(h)$ in $\mathfrak{M}$ with a new proposition $\mathsf{p}_{\varphi'}$ iff φ' holds in this very state. Then, we replace each φ' in φ with proposition $\mathsf{p}_{\varphi'}$. This yields a formula without nested modalities and the claim follows by induction.

■

A.2 Proofs of Section 4.3

In this section we give all details needed to prove Theorem 3. The structure of the proof is shown in Figure 9 on page 28.

The following lemma is essential to show that truth in i_sR-pando-like models is insensitive to the type of available strategies (memoryless vs. perfect recall). The lemma is needed to show that a uniform perfect recall strategy in the pando-like model gives rise to a uniform memoryless strategy. Therefore, we have to show that two states which are indistinguishable in the model give rise to indistinguishable histories.

Lemma 3 *Let $\mathfrak{M}$ be an i_sR-pando like* **iCGS** *and $q, q_1, \hat{q}_1, q_2, \hat{q}_2 \in St$ where q_i is reachable from $\hat{q}_i$, i.e. $\rho(\hat{q}_i, q_i) \neq \epsilon$, for $i = 1, 2$. Moreover, let $\hat{q}_1 \sim_b^{\mathfrak{M}} q \sim_c^{\mathfrak{M}} \hat{q}_2$ for some $b, c \in A$ and $q_1 \sim_a^{\mathfrak{M}} q_2$. Then, we have that $\rho(\hat{q}_1, q_1) \approx_a^{\mathfrak{M}} \rho(\hat{q}_2, q_2)$.*

Proof. The setting is illustrated in Figure 15. In the following we consider all possibilities how q, $\hat{q}_1$, $\hat{q}_2$, q_1, and q_2 can be located. We recall that $\rho_{\mathfrak{M}_k}(q') = q'$ means that q' is the root node of model $\mathfrak{M}_k$. We assume that $k, l, m \in I$ where I is the index set from Definition 16.

Case 1: $q_1 \hat{\sim}_a q_2$. Let $q_1 \in \mathfrak{M}_k$ and $q_2 \in \mathfrak{M}_l$, $k \neq l$.

Case 1.1: $\rho_{\mathfrak{M}_k}(q_1) = q_1$. That is, q_1 is the root node of $\mathfrak{M}_k$. We have $\hat{q}_1 = q_1 \sim_b^{\mathfrak{M}} q$. Then, by Definition 16.5 $|\rho(\hat{q}_2)| = |\rho(q_2)|$ which implies $\hat{q}_2 = q_2$. Hence, we have $\rho_{\mathfrak{M}_k}(\hat{q}_1, q_1) = q_1 \approx_a^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(\hat{q}_2, q_2) = q_2$ and are done.

Case 1.2 $\rho_{\mathfrak{M}_k}(q_2) = q_2$. Analogously to Case 1.1.

Case 1.3 $\rho_{\mathfrak{M}_k}(q_1) \hat{\approx}_a^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(q_2)$ and both q_1 and q_2 are not root nodes.

Case 1.3.1 $\hat{q}_1 \hat{\sim}_b^{\mathfrak{M}} q$. Let $q \in St_m$ with $m \neq k$.

Case 1.3.1.1 $\rho_{\mathfrak{M}_k}(\hat{q_1}) = \hat{q}_1$. Then, $\hat{q}_1 \hat{\sim}_a^{\mathfrak{M}} q'$ where $q' \in St_l$ is the root node of $\mathfrak{M}_l$. But then, by Definition 16.5 we have $|\rho(q')| = |\rho(\hat{q}_2)|$ and thus $\rho_{\mathfrak{M}_l}(\hat{q}_2) = \hat{q}_2$. This proves that $\rho_{\mathfrak{M}_k}(\hat{q}_1, q_1) \approx_a^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(\hat{q}_2, q_2)$.

Case 1.3.1.2 $\rho_{\mathfrak{M}_m}(q) = q$. Again, by Definition 16.5 following the same reasoning as in Case 1.3.1.1 we obtain $\rho_{\mathfrak{M}_l}(\hat{q}_2) = \hat{q}_2$. Showing that $\rho_{\mathfrak{M}_k}(\hat{q}_1, q_1) \approx_a^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(\hat{q}_2, q_2)$.

Case 1.3.1.3 $\rho_{\mathfrak{M}_k}(\hat{q}_1) \hat{\approx}_a^{\mathfrak{M}} \rho_{\mathfrak{M}_m}(q)$ and neither $\hat{q}_1$ nor q are root nodes. Then, we either have that $l = m$ and $\rho_{\mathfrak{M}_l}(\hat{q}_2) \approx_a^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(q)$ which implies that $\rho_{\mathfrak{M}_k}(\hat{q}_1, q_1) \hat{\approx}_a^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(\hat{q}_2, q_2)$. Or, $l \neq m$ and we have to distinguish again two cases. If $\hat{q}_2$ is the root node; then, it is connected with the root $q' \in St_k$ of $\mathfrak{M}_k$ by $\sim_a^{\mathfrak{M}}$. We have $q' \sim_a^{\mathfrak{M}} \hat{q}_2 \sim_c^{\mathfrak{M}} q \sim_b^{\mathfrak{M}} \hat{q}_1$ and by Definition 16.5 it must be the case that $|\rho(q')| = |\rho(\hat{q}_1)|$. Contradiction.

Hence, we can safely assume that $\hat{q}_2$ is not a root. Then, $\rho_{\mathfrak{M}_l}(\hat{q}_2) \approx_c^{\mathfrak{M}} \rho_{\mathfrak{M}_m}(q) \approx_b^{\mathfrak{M}} \rho_{\mathfrak{M}_k}(\hat{q}_1)$. So, all these states are on the same height level which implies that $\rho_{\mathfrak{M}_k}(\hat{q}_1, q_1) \approx_a^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(\hat{q}_2, q_2)$.

Case 1.3.2 $\hat{q}_2 \hat{\sim}_a^{\mathfrak{M}} q$. Analogously to Case 1.3.1

Case 1.3.3 $\hat{q}_1 \sim_a^{\mathfrak{M}_k} q$ or $\hat{q}_2 \sim_a^{\mathfrak{M}_l} q$. In each of these cases it means that either $q \notin St_k$ or $q \notin St_l$ as $k \neq l$. Case 1.3.1 or Case 1.3.2 applies.

Case 2: $q_1 \sim_a^{\mathfrak{M}_k} q_2$. Then, by definition $k = l$ and $\rho_{\mathfrak{M}_k}(q_1) \approx_a^{\mathfrak{M}} \rho_{\mathfrak{M}_k}(q_2)$.

Case 2.1: $q \in St_k$. We have $|\rho_{\mathfrak{M}_k}(\hat{q}_1)| = |\rho_{\mathfrak{M}_k}(q)| = |\rho_{\mathfrak{M}_k}(\hat{q}_2)|$ which follows from the assumption $\hat{q}_1 \sim_b^{\mathfrak{M}_k} q \sim_c^{\mathfrak{M}_k} \hat{q}_2$; hence also, $\rho_{\mathfrak{M}_k}(\hat{q}_1, q_1) \hat{\approx}_a^{\mathfrak{M}} \rho_{\mathfrak{M}_k}(\hat{q}_2, q_2)$.

Case 2.2: $q \in St_m,\ m \neq k$. Then, we have $\hat{q}_1 \hat{\sim}_b^{\mathfrak{M}} q \hat{\sim}_c^{\mathfrak{M}} \hat{q}_2$. Again we have to distinguish the different cases how q is connected to $\hat{q}_1$ and $\hat{q}_2$ respectively.

Case 2.1.1 $\rho_{\mathfrak{M}_m}(q) = q$. That is, we assume that q is a root node. By Definition 16.5 we have $|\rho(\hat{q}_1)| = |\rho(\hat{q_2})|$ and $\rho(\hat{q}_1, q_1) \approx_a^{\mathfrak{M}} \rho(\hat{q}_2, q_2)$ follows.

Case 2.1.2 $\rho_{\mathfrak{M}_m}(q) \neq q$. We have $\rho_{\mathfrak{M}_k}(\hat{q}_1) \hat{\approx}_b^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(q)$ and $\rho_{\mathfrak{M}_k}(\hat{q}_2) \hat{\approx}_c^{\mathfrak{M}} \rho_{\mathfrak{M}_l}(q)$ which implies $|\rho_{\mathfrak{M}_k}(\hat{q}_1)| = |\rho_{\mathfrak{M}_m}(q)| = |\rho_{\mathfrak{M}_k}(\hat{q}_2)|$ and hence $\rho_{\mathfrak{M}_k}(\hat{q}_1, q_1) \approx_a^{\mathfrak{M}} \rho_{\mathfrak{M}_k}(\hat{q}_2, q_2)$. ■

The next lemma analyses the structure of two indistinguishable nodes from subsequent tree levels.

Lemma 4 *Let $\mathfrak{M}$ be an* **iCGS***, q a state in it, $h_1 \in \Delta_{\mathfrak{M}}^i(q)$, $h_2 \in \Delta_{\mathfrak{M}}^{i+1}(q)$, and $h_1 \sim_a^{T_s(\mathfrak{M},q)} h_2$ for some $i \in \mathbb{N}_0$. Then, we have that $lastr(h_1) \sim_a^{\mathfrak{M}} rel(h_2)$.*

Proof. By definition, we have that $ref(h_2) \sim_a^{T_s(\mathfrak{M},q)} h_1$ and $lastr(ref(h_2)) \sim_a^{\mathfrak{M}} rel(h_2)$. Because $ref(h_2) \in \Delta_{\mathfrak{M}}^i(q)$ we also have $rel(h_1) \approx_a^{\mathfrak{M}} rel(ref(h_2))$, and hence $lastr(h_1) \sim_a^{\mathfrak{M}} lastr(ref(h_2))$. The claim follows because $lastr(ref(h_2)) \sim_a^{\mathfrak{M}} rel(h_2)$ and by the transitivity of $\sim_a^{\mathfrak{M}}$. ■

The next lemma states that nodes which are indistinguishable for a group of agents must be located on subsequent or the same level of the pando unfolding; moreover, it characterizes the structure of these nodes.

Lemma 5 *Let $\mathfrak{M}$ be an* **iCGS***, q a state in it, and $A \subseteq \mathbb{A}\mathrm{gt}$ be a group of agents. Then, for all $h \in St_{T_s(\mathfrak{M},q)}$ there is an $i \in \mathbb{N}_0$ such that for all $h' \in St_{T_s(\mathfrak{M},q)}$ with $h(\sim_A^{T_s(\mathfrak{M},q)})^* h'$ we have that $h, h' \in \Delta_{\mathfrak{M}}^i(q) \cup \Delta_{\mathfrak{M}}^{i+1}(q)$; moreover, if $h' \in \Delta_{\mathfrak{M}}^{i+1}(q)$ $h(\sim_A^{T_s(\mathfrak{M},q)})^* h'$ and there is an $h'' \in \Delta_{\mathfrak{M}}^i(q)$ with $h(\sim_A^{T_s(\mathfrak{M},q)})^* h''$ then $rel(h') \in St_{\mathfrak{M}}$ and $jump(h') \in A$.*

Proof. We write T_s for $T_s(\mathfrak{M},q)$ and Δ^i for $\Delta_{\mathfrak{M}}^i(q)$ and so on. We proceed by induction on the length of the epistemic path $h' = h_1 \sim_{a_1} \cdots \sim_{a_{l+1}} h_{l+1}$. We show the following: (i) if $h_j \in \Delta^i$ for all $j = 1, \dots, l$ then $h_{l+1} \in \Delta^{i-1} \cup \Delta^i \cup \Delta^{i+1}$; if this is not the case then (ii) if $h_j \in \Delta^i \cup \Delta^{i+1}$ for all $j = 1, \dots, l$ then $h_{l+1} \in \Delta^i \cup \Delta^{i+1}$ and for each $h_j \in \Delta^{i+1}$ we have that $rel(h_j) \in St$, $jump(h_j) \in A$, and $ref(h_j) \in \Delta^i$.

Base case: The case for $h \sim_a h'$ is clear by definition.
Induction step: Suppose we have $h' = h_1 \sim_{a_1} \cdots \sim_{a_l} h_l$ satisfying the assumption and assume that $h_l \sim_{a_{l+1}} h_{l+1}$. Firstly, assume that case (i) applies; that is, that all $h_j \in \Delta^i$ for $j = 1, \dots, l$. By definition $h_{l+1} \in \Delta^{i-1} \cup \Delta^i \cup \Delta^{i+1}$. Moreover, if $h_{l+1} \in \Delta^{i+1}$ then it must have the required form $h''\hat{a}q'$ by Definition 17 (c).

We consider case (ii). Firstly, suppose that $h_l \in \Delta^{i+1}$ and $h_l = h'\hat{a}_l q'$. We consider $h_l \sim_{a_{l+1}} h_{l+1}$. By definition $h_{l+1} \in \Delta^i \cup \Delta^{i+1} \cup \Delta^{i+2}$. If $h_{l+1} \in \Delta^{i+1}$ it also has the required form. For the sake of contradiction, suppose that $h_{l+1} \in \Delta^{i+2}$. Then, $h_{l+1} = h''\hat{a}_{l+1}q''$ for some $h'' \in \Delta^{i+1}$ with $h'' \sim_{a_{l+1}} h_l$. In this case, h'' does also have the form $h'' = h'''\hat{a}_l q'''$ and therewith $h_{l+1} = h'''\hat{a}_l q'''\hat{a}_{l+1}q''$ contradicting $h_{l+1} \in \Delta^{i+2}$ by definition of the sets Δ^j.

Secondly, if $h_l \in \Delta^i$ we have that $h_{l+1} \in \Delta^{i-1}\Delta^i \cup \Delta^{i+1}$ and that h_{l+1} has the required form if $h_{l+1} \in \Delta^{i+1}$ following the very same reasoning as in case (i). Moreover, it cannot be the case that $h_{l+1} \in \Delta^{i-1}$. To see this, we observe that there is some h_u with $1 \leq u < l+1$, $h_u \in \Delta^{i+1}$ and $h_{l+1}(\sim_A^{T_s})^* h_u$. Now the the reasoning of the previous case can be applied to obtain a contradiction. ■

The next lemma states that the relevant parts (i.e. histories) of two states of the pando unfolding are group indistinguishable in the model if the states are group indistinguishable in the pando unfolding and are located within the same tree (i.e. share the same root node).

Lemma 6 *Let $h_1(\sim_A^{T_s})^* h_2$. If there is an $i \in \mathbb{N}_0$ with $h_1, h_2 \in \Delta^i$ then $rel(h_1)(\approx_A^{\mathfrak{M}})^* rel(h_2)$.*

Proof. The poof is done by induction on the number of epistemic steps between h_1 and h_2. More precisely, we show that for all h' with $h_1(\sim_A^{T_s})^* h'$ we have that

(i) $rel(h_1)(\approx_A^{\mathfrak{M}})^* rel(h')$ if $h' \in \Delta^i$;
(ii) $ref(h')(\sim_A^{T_s})^* h_1$ and $lastr(h_1)(\sim_A^{\mathfrak{M}})^* lastr(ref(h'))$ if $h' \in \Delta^{i+1}$;
(iii) and $rel(h_1)(\sim_A^{\mathfrak{M}})^* lastr(h')$ if $h' \in \Delta^{i-1}$ (and $i > 0$).

The base cases are clear by definition. We assume that $h_1(\sim_A^{T_s})^* h'$, $i > 0$ and we show that h_2 with $h_1(\sim_A^{T_s})^* h' \sim_a^{T_s} h_2$ for $a \in A$ satisfies the property of the lemma.

Case: $h' \in \Delta^i$ and $h_2 \in \Delta^i$. By definition $rel(h') \approx_a^{\mathfrak{M}} rel(h_2)$ and by induction $rel(h_1)(\approx_A^{\mathfrak{M}})^* rel(h')$; hence, $rel(h_1)(\approx_A^{\mathfrak{M}})^* rel(h_2)$.
Case: $h' \in \Delta^i$ and $h_2 \in \Delta^{i+1}$. By induction, $rel(h_1)(\approx_A^{\mathfrak{M}})^* rel(h')$ and in particular, $lastr(h_1) \sim_A^* lastr(h')$. By Lemma 4 $lastr(h') \sim_a^{\mathfrak{M}} rel(h_2)$. This shows that, $lastr(h_1)(\sim_A^{\mathfrak{M}})^* rel(h_2) \sim_a^{\mathfrak{M}} lastr(ref(h_2))$.
Case: $h' \in \Delta^i$ and $h_2 \in \Delta^{i-1}$. By definition $h_2 \sim_A^{T_s} ref(h')$ and $rel(h') \sim_A^{\mathfrak{M}} lastr(h_2)$; hence, $rel(h') \in St_{\mathfrak{M}}$. Thus, by induction $rel(h_1)(\sim_A^{\mathfrak{M}})^* rel(h')$ and by Lemma 4 $rel(h') \sim_a^{\mathfrak{M}} lastr(h_2)$. This shows that $rel(h_1)(\sim_A^{\mathfrak{M}})^* lastr(h_2)$.
Case: $h' \in \Delta^{i-1}$ and $h_2 \in \Delta^{i-1}$. Follows immediately.
Case: $h' \in \Delta^{i-1}$ and $h_2 \in \Delta^i$. We have $rel(h_1)(\sim_A^{\mathfrak{M}})^* lastr(h')$. By Lemma 4 $rel(h_2) \sim_a^{\mathfrak{M}} lastr(h')$ and hence $rel(h_1)(\sim_A^{\mathfrak{M}})^* rel(h_2)$. The claim follows as $rel(h_1), rel(h_2) \in St_{\mathfrak{M}}$.

Cases where $h' \in \Delta^j$ and $h_2 \in \Delta^k$ with $|j - k| > 1$ are not possible due to Lemma 5. ■

Lemma 7 *For all q in $\mathfrak{M}$ and all $i, j \in \mathbb{N}_0$ with $i \neq j$ we have that $\Delta^i_{\mathfrak{M}}(q) \cap \Delta^j_{\mathfrak{M}}(q) = \emptyset$.*

Lemma 8 *Let $\mathfrak{M}$ be an **iCGS**, q a state in it, and $a \in \mathbb{A}\mathrm{gt}$. Every relation $\sim_a^{T_s(\mathfrak{M},q)}$ is an equivalence relation.*

Proof. We write T_s for $T_s(\mathfrak{M}, q)$. Reflexivity and symmetry of epistemic relations in T_s are clear from the definition of i_sR-pando unfoldings, but we need to prove transitivity. Suppose that $h_1 \sim_a^{T_s} h_2$ and $h_2 \sim_a^{T_s} h_3$. We have to show that $h_1 \sim_a^{T_s} h_3$. The proof is done by induction on the level Δ^i. The base case for Δ^0 is clear from the transitivity of the standard indistinguishability relation $\approx_a^{\mathfrak{M}}$. By Lemma 5 (and the symmetry of $\sim^{T_s}$) it is sufficient to consider the following cases (we assume that $i > 0$):

$h_1, h_2, h_3 \in \Delta^i$: Follows by the transitivity of $\approx_a^{\mathfrak{M}}$ (induction hypothesis).

$h_1, h_2 \in \Delta^i, h_3 \in \Delta^{i+1}$: From $h_1 \sim_a^{T_s} h_2$ it follows that $rel(h_1) \approx_a^{\mathfrak{M}} rel(h_2)$; and from $h_2 \sim_a^{T_s} h_3$ that $ref(h_3) \sim_a^{T_s} h_2$ and $lastr(ref(h_3)) \sim_a^{\mathfrak{M}} rel(h_3)$. Furthermore, because $ref(h_3) \in \Delta^i$ we can deduce from the transitivity of $\sim_a^{T_s}$ (and by induction) that $ref(h_3) \sim_a^{T_s} h_1$ and hence $h_3 \sim_a^{T_s} h_1$ by definition.

$h_1 \in \Delta^i, h_2, h_3 \in \Delta^{i+1}$: We have that $ref(h_2) \sim_a^{T_s} h_1$ and $lastr(ref(h_2)) \sim_a^{T_s} rel(h_2)$ and $jump(h_2) = a$. Then, we also have that $rel(h_2) \approx_a^{T_s} rel(h_3)$ and $jump(h_3) = jump(h_2) = a$ and $ref(h_2) \sim_a^{T_s} ref(h_3)$ proving that also $h_1 \sim_a^{T_s} h_3$ by induction.

$h_1, h_3 \in \Delta^i, h_2 \in \Delta^{i+1}$: We have that $ref(h_2) \sim_a^{T_s} h_1$ and $lastr(ref(h_2)) \sim_a^{T_s} rel(h_2)$ and $jump(h_2) = a$ and $ref(h_2) \sim_a^{T_s} h_3$ and thus $h_1 \sim_a^{T_s} h_3$ (by induction).

$h_1, h_3 \in \Delta^{i+1}, h_2 \in \Delta^i$: We have that $ref(h_1) \sim_a^{T_s} h_2$ and $lastr(ref(h_1)) \sim_a^{T_s} rel(h_1)$ and $jump(h_1) = a$ and $ref(h_3) \sim_a^{T_s} h_2$ $lastr(ref(h_3)) \sim_a^{T_s} rel(h_3)$ and $jump(h_3) = a$. But then by induction $ref(h_1) \sim_a^{T_s} ref(h_3)$ and $rel(h_1) \sim_a^{\mathfrak{M}} rel(h_3)$ which shows that $h_1 \sim_a^{T_s} h_3$.

■

Proposition 8 ($\rightsquigarrow$ page 29). *The $i_s R$-pando unfolding of a pointed* **iCGS** *is $i_s R$-pando-like.*

Proof. Let $\mathfrak{M}$ be an **iCGS** and T_s its $i_s R$-pando unfolding. For each $h \in St'$ with $|rel(h)| = 1$ we define St_h as the set of states/histories in St' reachable from h, i.e. $St_h = \{h' \in St' \mid \rho_{T_s}(h, h') \neq \epsilon\}$. Let $\mathfrak{M}_h$ denote the submodel of T_s which does only consist of states St_h and in which the domain of all elements is restricted to St_h. Moreover, we take $I = \{h \in St \mid |rel(h)| = 1\}$.
Claim: We have that $St' = \biguplus_{h \in I} St_h$ and each $\mathfrak{M}_h$ is $i_o R$-tree-like.
Proof of claim: Clearly, all sets St_h are mutually disjoint and each $h \in St$ has to occur in some St_h. It is also obvious that each $\mathfrak{M}_h$ has tree-structure. Now suppose $h_1, h_2 \in St_h$ with $h_1 \sim_a^{T_s} h_2$; then, by definition also $h_1 \approx_a^{\mathfrak{M}} h_2$. Q.e.d.

We proceed with the main proof and define $\hat{\sim}_a$ as the subset of $\sim_a^{T_s}$ which exists between sets St_h and $St_{h'}$ with $h \neq h'$. From Lemma 8 it follows that $\sim_a^{T_s}$ is transitive and that $\hat{\sim}_a$ is symmetric. Moreover, by definition $\hat{\sim}_a \cap (St_h \times St_h) = \emptyset$ for all $h \in St'$.

The fourth condition of Definition 16 is obvious from the definition of the $i_s R$-pando unfolding. It remains to show the fifth condition of Definition 16. Suppose $h_1, h_1' \in St_{\bar{h}_1}$ and $h_1(\sim_{\mathbb{A}\mathrm{gt}})^* h_1'$. Then, also $h_1, h_1' \in \Delta_i$ for some i. From Lemma 6 we obtain that $rel(h_1)(\approx_{\mathbb{A}\mathrm{gt}}^{\mathfrak{M}})^* rel(h_1')$, i.e. that both nodes reside on the same level.

■

The following two lemmata are needed to prove Theorem 3. The first lemma states that the set of epistemic alternatives to any state is the same in the model and in the pando unfolding.

Lemma 9 *Let $\mathfrak{M}$ be an* **iCGS** *and q_0 a state in it. Then, the following property holds: For all $A \subseteq \mathbb{A}\mathrm{gt}$ and all nodes h in $T_s(\mathfrak{M}, q_0)$ we have that $\{q \mid lastr(h) \sim_A^{\mathfrak{M}} q,\ q \in St_{\mathfrak{M}}\} = \{lastr(h') \mid h' \sim_A^{T_s(\mathfrak{M},q_0)} h,\ h' \in St_{T_s(\mathfrak{M},q_0)}\}$.*

Proof. "$\subseteq$": Suppose $h = h'q' \in \Delta_{\mathfrak{M}}^i(q_0)$ and $q' \sim_A^{\mathfrak{M}} q$ and $h' \neq \epsilon$ (the case for $h' = \epsilon$ is clear). Then, there is some $a \in A$ with $q' \sim_a^{\mathfrak{M}} q$ and thus $h'' := \hat{h}\hat{a}q \in \Delta_{\mathfrak{M}}^{i+1}(q_0) \subseteq St_{T_s(\mathfrak{M},q_0)}$ with $\hat{h} \sim_a^{T_s} h$ by definition of the $\Delta_{\mathfrak{M}}^j$'s. By Definition 17 also $h'' \sim_A^{T_s(\mathfrak{M},q_0)} h$. The claim follows as $lastr(h'') = q$.

"$\supseteq$": Suppose $h' \sim_A^{T_s(\mathfrak{M},q_0)} h$ and $h' \in St_{T_s(\mathfrak{M},q_0)}$. The claim is clear if $h, h' \in \Delta_{\mathfrak{M}}^i(q_0)$. According to Definition 17 the remaining case is when $h \in \Delta_{\mathfrak{M}}^i(q)$ and $h' \in \Delta_{\mathfrak{M}}^{i+1}(q)$, or the roles of h and h' switched. Then, $h' = \hat{h}\hat{a}q$, $\hat{h} \sim_a^{T_s} h$ for some $a \in \mathbb{A}\mathrm{gt}$ and $lastr(\hat{h}) \sim_a^{\mathfrak{M}} q$. The claim follows as $lastr(h') = q$ and $lastr(h) \sim_A^{\mathfrak{M}} q$.

■

The next lemma is needed to show that the witnessing strategy which we shall construct in the invariance Theorem 3 is uniform.

Lemma 10 *Let $\mathfrak{M}$ be an* **iCGS**, *q a state in it, $T_s = T_s(\mathfrak{M}, q)$, $h, \hat{h}_1, \hat{h}_2 \in St_{T_s}$, $A \subseteq \mathbb{A}\mathrm{gt}$, and $a \in \mathbb{A}\mathrm{gt}$. If $\hat{h}_1 \sim_A^{T_s} h \sim_A^{T_s} \hat{h}_2$ and $h_1 = \hat{h}_1 h_1^F \sim_a \hat{h}_2 h_2^F = h_2$ with $h_1^F, h_2^F \in \Lambda^{fin}$; then, ($rel(h_1) \approx_a^{\mathfrak{M}} rel(h_2)$, $\rho(\hat{h}_1, h_1) \approx_a^{\mathfrak{M}} \rho(\hat{h}_2, h_2)$ and $|h_1^F| = |h_2^F|$), or $h_1^F = h_2^F = \epsilon$.*

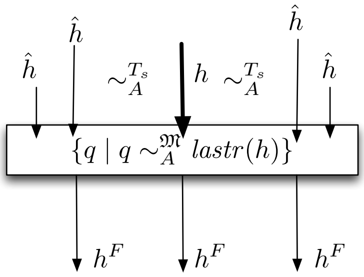

Fig. 16 Setting of the proof of Theorem 3.

Proof. Suppose $\hat{h}_1 \sim_b^{T_s} h \sim_c^{T_s} \hat{h}_2$ with $b, c \in A$. From Lemma 3 we obtain $\rho(\hat{h}_1, h_1) \approx_a^{\mathfrak{M}} \rho(\hat{h}_2, h_2)$ by taking $\hat{q}_i = \hat{h}_i$ and $q_i = h_i$ and $q = h$ for $i = 1, 2$.

To prove the lemma, we firstly assume that $h_1, h_2 \in \Delta^i$ for some $i \in \mathbb{N}_0$. Then, by definition $rel(h_1) \approx_a^{\mathfrak{M}} rel(h_2)$ and the claim follows.

Secondly, suppose w.l.o.g. $h_1 \in \Delta^i$ and $h_2 \in \Delta^{i+1}$ for some $i \in \mathbb{N}_0$. In this case $rel(h_2) \in St$ and thus $h_2^F = \epsilon$ and $\rho(\hat{h}_2, h_2) = rel(h_2) \approx_a^{\mathfrak{M}} \rho(\hat{h}_1, h_1)$. The latter implies that $|\rho(\hat{h}_1, h_1)| = 1$ and hence $h_1^F = \epsilon$. ■

Theorem 3 ($\rightsquigarrow$ page 29). *For every* **iCGS** $\mathfrak{M}$*, state* q *in* $\mathfrak{M}$*, and* ***ATL*****-formula* φ*, it holds that*

$$\mathfrak{M}, q \models_{i_sR} \varphi \text{ iff } T_s(\mathfrak{M}, q), q \models_{i_sR} \varphi \text{ iff } T_s(\mathfrak{M}, q), q \models_{i_sr} \varphi.$$

Proof. We show that for every node h in $T_s := T_s(\mathfrak{M}, q)$ it holds that $\mathfrak{M}, lastr(h) \models_{i_sR} \varphi$ iff $T_s(\mathfrak{M}, q), h \models_{i_sr} \varphi$. Then, the claim follows from Propositions 7 and 8 for $h = q$. The proof is done by induction on the structure of φ and is similar to the proof given for Theorem 2.

Base cases:

Propositional case: Straightforward.
Case: $\varphi \equiv \langle\langle A \rangle\rangle \gamma$ where γ contains no nested strategic modalities.

"⇒": Suppose we have $\mathfrak{M}$, $lastr(h) \models_{i_sR} \langle\langle A \rangle\rangle \gamma$ and let s_A be an iR-strategy with

$$(\star)\ \forall \lambda \in out_{\mathfrak{M}}^{i_s}(lastr(h), s_A) : \mathfrak{M}, \lambda \models_{i_sR} \gamma.$$

We construct the *ir*-strategy s'_A in $T_s(\mathfrak{M}, q)$ as follows: for all $\hat{h} \in St_{T_s}$ with $h \sim_A^{T_s} \hat{h}$ and all $h^F \in \Lambda_{\mathfrak{M}}^{fin}(lastr(\hat{h}))$ we set

$$s'_a(\hat{h}(h^F[1, \infty])) := s_a(h^F).$$

We note that we have to exclude the first state in h^F because it is already contained in $\hat{h}$. For all other histories h'' (which do not have the prescribed form) we define $s'_a(h'')$ arbitrarily but in a uniform way. The setting is illustrated in Figure 16.

Clearly, we have that each $\hat{h}(h^F[1, \infty])$ is a valid state in T_s and by Lemma 10, s'_a is well-defined: Suppose, there are $h_1 = \hat{h}_1 h_1^F$ and $h_2 = \hat{h}_2 h_2^F$ with $h_1 = h_2$. Then, also $h_1 \sim_a^{T_s} h_2$ and by Lemma 10 $h_1^F = h_2^F$ which shows that $s'_a(h_1) = s'_a(h_2)$.

In the following we show that s'_a *is uniform*. Let h_1 and h_2 be two histories with $h_1 \sim_a^{T_s} h_2$.

1. Assume that both nodes have the form from above, i.e. $h_1 = \hat{h}_1 h_1^F$, $h_1 = \hat{h}_1 h_1^F$ and $h_1 \sim_A^{T_s} h \sim_A^{T_s} h_2$. Then, uniformity follows from Lemma 10.
2. Choices for two histories h_1 and h_2 where at least one does not have the required form can be defined in a uniform way by definition because $\sim_a^{T_s}$ is an equivalence relation.

Finally, we show that the sets of outcome paths are isomorphic wrt. both strategies. By Lemma 9 we have that for all states q_0 in $\mathfrak{M}$ the following holds $q_0 \sim_A^{\mathfrak{M}} lastr(h)$ iff there is a history h' with $h' \sim_A^{T_s} h$ and $lastr(h') = q_0$. We denote one of these histories h' by $h(q_0)$. Then, by construction of s'_A we have that

$$q_0 q_1 q_2 \dots \in out^{i_s}_{\mathfrak{M}}(last(h), s_A) \text{ iff } (h')(h'q_1)(h'q_1q_2)\dots \in out^{i_s}_{T_s(\mathfrak{M},q)}(h, s'_A) \text{ for some } h' = h(q_0).$$

Since the valuation of propositions does only depend on the final state of a history and by $(\star)$ we have $T_s(\mathfrak{M}, q), h \models_{i_s r} \langle\langle A\rangle\rangle\gamma$.

$\Leftarrow$: For the other direction, suppose we have $T_s(\mathfrak{M}, q), h \models_{i_s r} \langle\langle A\rangle\rangle\gamma$. So, there is an *ir*-strategy s_A such that

$$(\star)\ \forall \lambda \in out^{i_s}_{T_s(\mathfrak{M},q)}(h, s_A) : T_s(\mathfrak{M}, q), \lambda \models_{i_s r} \gamma.$$

We construct a witnessing *iR*-strategy s'_A in $\mathfrak{M}$ as follows: $s'_a(h^F) = s_a(\hat{h}\hat{a}h^F)$ for every $a \in A$, $\hat{h} \sim_A^{T_s} h$ for $h^F \in \Lambda^{fin}(q')$ with $q' \sim_A^{\mathfrak{M}} lastr(h)$ and arbitrary but in a uniform way for all other histories. It is easy to verify that each strategy s'_a is uniform and well-defined. Moreover, s'_A yields an equivalent (apart from the notational differences) set of outcome paths as above. We have $\mathfrak{M}, lastr(h) \models_{i_s R} \langle\langle A\rangle\rangle\gamma$.

Induction step:

Case: $\varphi \equiv \psi_1 \wedge \psi_2$. Straightforward.

Case: $\varphi \equiv \neg\psi$. $\mathfrak{M}, lastr(h) \models_{i_s R} \neg\psi$ iff not $\mathfrak{M}, lastr(h) \models_{i_s R} \psi$ iff (by induction hypothesis) not $T_s(\mathfrak{M}, q), h \models_{i_s R} \psi$ iff $T_s(\mathfrak{M}, q), h \models_{i_s R} \neg\psi$.

Case: $\varphi \equiv \langle\langle A\rangle\rangle\gamma$. By induction hypothesis we have for each history h in $T_s(\mathfrak{M}, q)$ and each strict state-subformula φ' of γ that $\mathfrak{M}, lastr(h) \models_{i_s R} \varphi'$ iff $T_s(\mathfrak{M}, q), h \models_{i_s r} \varphi'$. For any maximal strict state-subformula φ' we label all states h in $T_s(\mathfrak{M}, q)$ and states $lastr(h)$ in $\mathfrak{M}$ with a new proposition $\mathsf{p}_{\varphi'}$ iff φ' holds in this very state. Then, we replace each φ' in φ with proposition $\mathsf{p}_{\varphi'}$ and the claim follows by induction.

■

References

1. T. Ågotnes. A note on syntactic characterization of incomplete information in atel. In *Procedings of the Workshop on Knowledge and Games, Liverpool*, pages 34–42, 2004.
2. T. Ågotnes. Action and knowledge in alternating-time temporal logic. *Synthese*, 149(2):377–409, 2006. Section on Knowledge, Rationality and Action.
3. 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.
4. T. Ågotnes and D. Walther. A logic of strategic ability under bounded memory. *Journal of Logic, Language and Information*, 18(1):55–77, 2009.
5. N. Alechina, B. Logan, N.H. Nga, and A. Rakib. A logic for coalitions with bounded resources. In *Proceedings of IJCAI*, pages 659–664, 2009.
6. N. Alechina, B. Logan, N.H. Nga, and A. Rakib. Resource-bounded alternating-time temporal logic. In *Proceedings of AAMAS*, pages 481–488, 2010.
7. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 100–109. IEEE Computer Society Press, 1997.
8. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Lecture Notes in Computer Science*, 1536:23–60, 1998.
9. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
10. R. Alur, T.A. Henzinger, O. Kupferman, and M.Y. Vardi. Alternating refinement relations. In *Proceedings of CONCUR*, volume 1466 of *LNCS*, pages 163–178, 1998.

11. T. Brihaye, A. Da Costa Lopes, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. In *Proceedings of LFCS*, volume 5407 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2009.
12. N. Bulling and J. Dix. Modelling and verifying coalitions using argumentation and ATL. *Inteligencia Artificial, Revista Iberoamericana de Inteligencia Artificial*, 14(46):45–73, 2010.
13. N. Bulling, J. Dix, and W. Jamroga. Model checking logics of strategic ability: Complexity. In M. Dastani, K. Hindriks, and J.-J. Meyer, editors, *Specification and Verification of Multi-Agent Systems*, pages 125–159. Springer, 2010.
14. N. Bulling and B. Farwer. Expressing properties of resource-bounded systems: The logics RTL* and RTL. In *Proceedings of CLIMA*, pages 22–45, 2009.
15. N. Bulling and B. Farwer. On the (un-)decidability of model checking resource-bounded agents. In *Proceedings of ECAI*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pages 567–572. IOS Press, 2010.
16. N. Bulling and W. Jamroga. What agents can probably enforce. *Fundamenta Informaticae*, 93(1-3):81–96, 2009.
17. N. Bulling and W. Jamroga. Verifying agents with memory is harder than it seemed. *AI Communications*, 23:380–403, 2010.
18. N. Bulling and W. Jamroga. Alternating epistemic mu-calculus. In *Proceedings of IJCAI-11*, pages 109–114, 2011.
19. N. Bulling, W. Jamroga, and J. Dix. Reasoning about temporal properties of rational play. *Annals of Mathematics and Artificial Intelligence*, 53(1-4):51–114, 2008.
20. K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. In *Proceedings of CONCUR*, pages 59–73, 2007.
21. E.M. Clarke and E.A. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proceedings of Logics of Programs Workshop*, volume 131 of *Lecture Notes in Computer Science*, pages 52–71, 1981.
22. E. Davis. Knowledge preconditions for plans. *Journal of Logic and Computation*, 4(5):721–766, 1994.
23. C. Dima, C. Enea, and D.P. Guelev. Model-checking an alternating-time temporal logic with knowledge, imperfect information, perfect recall and communicating coalitions. In *Proceedings of GANDALF*, pages 103–117, 2010.
24. C. Dima and F.L. Tiplea. Model-checking atl under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
25. H. Ghaderi, H. Levesque, and Y. Lespérance. A logical theory of coordination and joint ability. In *Proceedings of AAAI'07*, pages 421–426. AAAI Press, 2007.
26. V. Goranko. Coalition games and alternating temporal logics. In J. van Benthem, editor, *Proceedings of TARK VIII*, pages 259–272. Morgan Kaufmann, 2001.
27. V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004.
28. V. Goranko and W. Jamroga. State and path effectivity models for logics of multi-player games. In *Proceedings of AAMAS*, 2012. To appear.
29. V. Goranko, W. Jamroga, and P. Turrini. Strategic games and truly playable effectivity functions. In *Proceedings of AAMAS2011*, pages 727–734, 2011.
30. V. Goranko and D. Shkatov. Tableau-based decision procedures for logics of strategic ability in multiagent systems. *ACM Trans. Comput. Log.*, 11(1), 2009.
31. V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, 353(1):93–117, 2006.
32. D.P. Guelev, C. Dima, and C. Enea. An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. *Journal of Applied Non-Classical Logics*, 21(1):93–131, 2011.
33. 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.
34. P. Hawke. Coordination, almost perfect information and strategic ability. In *Proceedings of LAMAS*, 2010.
35. W. Jamroga. Some remarks on alternating temporal epistemic logic. In B. Dunin-Keplicz and R. Verbrugge, editors, *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003)*, pages 133–140, 2003.

36. W. Jamroga. On the relationship between playing rationally and knowing how to play: A logical account. In C. Freksa, M. Kohlhase, and K. Schill, editors, *Proceedings of KI 2006*, volume 4314 of *Lecture Notes in Artificial Intelligence*, pages 419–433. Springer Verlag, 2006.
37. W. Jamroga. A temporal logic for stochastic multi-agent systems. In *Proceedings of PRIMA'08*, volume 5357 of *LNCS*, pages 239–250, 2008.
38. W. Jamroga and T. Ågotnes. Modular interpreted systems: A preliminary report. Technical Report IfI-06-15, Clausthal University of Technology, 2006.
39. 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.
40. W. Jamroga and N. Bulling. Comparing variants of strategic ability. In *Proceedings of EUMAS2010*, 2010.
41. W. Jamroga and N. Bulling. Comparing variants of strategic ability. In *Proceedings of IJCAI-11*, pages 252–257, 2011.
42. W. Jamroga and J. Dix. Model checking abilities of agents: A closer look. *Theory of Computing Systems*, 42(3):366–410, 2008.
43. W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
44. W. Jamroga, W. van der Hoek, and M. Wooldridge. Intentions and strategies in game-like scenarios. In Carlos Bento, Amílcar Cardoso, and Gaël Dias, editors, *Progress in Artificial Intelligence: Proceedings of EPIA 2005*, volume 3808 of *Lecture Notes in Artificial Intelligence*, pages 512–523. Springer Verlag, 2005.
45. G. Jonker. Feasible strategies in Alternating-time Temporal Epistemic Logic. Master thesis, University of Utrecht, 2003.
46. M. Kacprzak and W. Penczek. Unbounded model checking for Alternating-time Temporal Logic. In *Proceedings of AAMAS-04*, 2004.
47. M. Köster and P. Lohmann. Abstraction for model checking modular interpreted systems over ATL. In *Proceedings of AAMAS*, pages 1129–1130, 2011.
48. F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. *Logical Methods in Computer Science*, 4:7, 2008.
49. Y. Lesperance, H.J. Levesque, F. Lin, and R.B. Scherl. Ability and knowing how in the situation calculus. *Studia Logica*, 66(1):165–186, 2000.
50. F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. What makes ATL* decidable? A decidable fragment of strategy logic. In *Proceedings of CONCUR*, pages 193–208, 2012.
51. F. Mogavero, A. Murano, and M.Y. Vardi. Reasoning about strategies. In *Proceedings of FSTTCS*, pages 133–144, 2010.
52. M. Pauly. *Logic for Social Software*. PhD thesis, University of Amsterdam, 2001.
53. M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
54. 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.
55. H. Schnoor. Deciding epistemic and strategic properties of cryptographic protocols. Technical Report tr_1012, Kiel University, 2010.
56. Henning Schnoor. Strategic planning for probabilistic games with incomplete information. In *Proceedings of AAMAS'10*, pages 1057–1064, 2010.
57. P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.
58. W. van der Hoek, W. Jamroga, and M. Wooldridge. A logic for strategic reasoning. In *Proceedings of AAMAS'05*, pages 157–164, 2005.
59. 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.
60. W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In C. Castelfranchi and W.L. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-02)*, pages 1167–1174. ACM Press, New York, 2002.
61. W. van der Hoek and M. Wooldridge. Cooperation, knowledge and time: Alternating-time Temporal Epistemic Logic and its applications. *Studia Logica*, 75(1):125–157, 2003.
62. S. van Otterloo and G. Jonker. On Epistemic Temporal Strategic Logic. *Electronic Notes in Theoretical Computer Science*, XX:35–45, 2004. Proceedings of LCMAS'04.

63. S. van Otterloo, W. van der Hoek, and M. Wooldridge. Preferences in game logics. In *Proceedings of AAMAS-04*, pages 152–159, 2004.
64. 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.
65. D. Walther, W. van der Hoek, and M. Wooldridge. Alternating-time temporal logic with explicit strategies. In D. Samet, editor, *Proceedings TARK XI*, pages 269–278. Presses Universitaires de Louvain, 2007.
66. N. Yadav and S. Sardiña. Reasoning about agent programs using ATL-like logics. In *Proceedings of JELIA*, pages 437–449, 2012.