On the Relationship between Playing Rationally and Knowing how to Play: A Logical Account

Wojciech Jamroga

Department of Informatics,
Clausthal University of Technology, Germany
wjamroga@in.tu-clausthal.de

Abstract. Modal logics of strategic ability usually focus on capturing what it means for an agent to have a feasible strategy that brings about some property. While there is a general agreement on abilities in scenarios where agents have perfect information, the right semantics for ability under incomplete information is still debated upon. Epistemic Temporal Strategic Logic, an offspring of this debate, can be treated as a logic that captures properties of agents' rational play.
In this paper, we provide a semantics of ETSL that is more compact and comprehensible than the one presented in the original paper by van Otterloo and Jonker. Second, we use ETSL to show that a rational player knows that he will succeed if, and only if, he knows how to play to succeed – while the same is not true for rational coalitions of players.

Keywords: multi-agent systems, theories of agency, game-theoretical foundations, modal logic.

1 Introduction

Modal logics of strategic ability usually focus on capturing what it means for an agent to have a feasible strategy that brings about some property. While there is a general agreement on abilities in scenarios where agents have perfect information, the right semantics for ability under incomplete information is still debated upon. Epistemic Temporal Strategic Logic, proposed by van Otterloo and Jonker [13], is an offspring of this debate, but one that leads in an orthogonal direction to the mainstream solutions. The central operator of ETSL can be read as: "if A play *rationally* to achieve φ (meaning: they never play a dominated strategy), they will achieve φ". Thus, one may treat ETSL as a logic that captures properties of agents' rational play in a sense.

This paper contains two main messages. First, we provide a semantics of ETSL that is more compact and comprehensible than the one presented in [13]. ETSL is underpinned by several exciting concepts. Unfortunately, its semantics is also quite hard to read due to a couple non-standard solutions and a plethora of auxiliary functions, which is probably why the logic never received the attention it deserves. Second, and perhaps more importantly, we use ETSL to show that a rational player knows that he will succeed if, and only if, he knows how to play to succeed – while the same is not true for rational coalitions of players.

2 Reasoning about Abilities of Agents

Modal logics of strategic ability [1, 2] form one of the fields where logic and game theory can successfully meet. The logics have clear possible worlds semantics, are axiomatizable, and have some interesting computational properties. Moreover, they are underpinned by intuitively appealing conceptual machinery for modeling and reasoning about systems that involve multiple autonomous agents.

2.1 ATL: Ability in Perfect Information Games

Alternating-time Temporal Logic (ATL) [1, 2] can be seen as a logic for systems involving multiple agents, that allows one to reason about what agents can achieve in game-like scenarios. Since ATL does not include incomplete information in its scope, it can be seen as a logic for reasoning about agents who always have perfect information about the current state of affairs. Formula $\langle\!\langle A\rangle\!\rangle\varphi$, where A is a coalition of agents, expresses that A have a collective strategy to enforce φ. ATL formulae include temporal operators: "$\bigcirc$" ("in the next state"), $\Box$ ("always from now on") and $\mathcal{U}$ ("until"). Operator $\Diamond$ ("now or sometime in the future") can be defined as $\Diamond\varphi \equiv \top\,\mathcal{U}\,\varphi$. Like in CTL, every occurrence of a temporal operator is preceded by exactly one cooperation modality $\langle\!\langle A\rangle\!\rangle$.[1] Formally, the recursive definition of ATL formulae is:

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

A number of semantics have been defined for ATL, most of them equivalent [3]. In this paper, we use a variant of *concurrent game structures*,

$$M = \langle \mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o\rangle,$$

which includes a nonempty finite set of all agents $\mathbb{A}\mathrm{gt} = \{1, ..., k\}$, a nonempty set of states St, a set of atomic propositions Π, a valuation of propositions $\pi : \Pi \to \mathcal{P}(St)$, and a nonempty set of (atomic) actions Act. Function $d : \mathbb{A}\mathrm{gt}\times St \to \mathcal{P}(Act)$ defines actions available to an agent in a state, and o is a deterministic transition function that assigns an outcome state $q' = o(q, \alpha_1, \ldots, \alpha_k)$ to state q, and a tuple of actions $\langle\alpha_1, \ldots, \alpha_k\rangle$ that can be executed by $\mathbb{A}\mathrm{gt}$ in q. A *strategy* of agent a is a conditional plan that specifies what a is going to do for every possible situation ($s_a : St \to Act$ such that $s_a(q) \in d(a, q)$). A *collective strategy* (called also a *strategy profile*) S_A for a group of agents A is a tuple of strategies S_a, one per agent $a \in A$. A *path* Λ in M is an infinite sequence of states that can be effected by subsequent transitions, and refers to a possible course of action (or a possible computation) that may occur in the system; by $\Lambda[i]$, we denote the ith position on path Λ. Function $out(q, S_A)$ returns the set of all paths that may result from agents A executing strategy S_A from state q onward:

[1] The logic to which such a syntactic restriction applies is sometimes called *"vanilla"* ATL (resp. "vanilla" CTL etc.).

$out(q, S_A) = \{\lambda = q_0 q_1 q_2 ... \mid q_0 = q$ and for every $i = 1, 2, ...$ there exists a tuple of actions $\langle \alpha_1^{i-1}, ..., \alpha_k^{i-1} \rangle$ such that $\alpha_a^{i-1} = S_a(q_{i-1})$ for each $a \in A$, $\alpha_a^{i-1} \in d(a, q_{i-1})$ for each $a \notin A$, and $o(q_{i-1}, \alpha_1^{i-1}, ..., \alpha_k^{i-1}) = q_i\}$.

Now, the semantics of ATL formulae can be given via the following clauses:

$M, q \models p$ iff $q \in \pi(p)$ (where $p \in \Pi$);

$M, q \models \neg\varphi$ iff $M, q \not\models \varphi$;

$M, q \models \varphi \wedge \psi$ iff $M, q \models \varphi$ and $M, q \models \psi$;

$M, q \models \langle\langle A \rangle\rangle \bigcirc \varphi$ iff there is a collective strategy S_A such that, for every $\Lambda \in out(q, S_A)$, we have $M, \Lambda[1] \models \varphi$;

$M, q \models \langle\langle A \rangle\rangle \Box \varphi$ iff there exists S_A such that, for every $\Lambda \in out(q, S_A)$, we have $M, \Lambda[i]$ for every $i \geq 0$;

$M, q \models \langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$ iff there is S_A st. for every $\Lambda \in out(q, S_A)$ there is $i \geq 0$, for which $M, \Lambda[i] \models \psi$, and $M, \Lambda[j] \models \varphi$ for every $0 \leq j < i$.

2.2 Strategic Ability and Incomplete Information

ATL is unrealistic in a sense: real-life agents seldom possess complete information about the current state of the world. *Alternating-time Temporal Epistemic Logic* (ATEL) [12] enriches the picture with an epistemic component, adding to ATL operators for representing agents' knowledge: $K_a\varphi$ reads as "agent a knows that φ". Additional operators $E_A\varphi$, $C_A\varphi$, and $D_A\varphi$ refer to *mutual knowledge* ("everybody knows"), *common knowledge*, and *distributed knowledge* among the agents from A. Models for ATEL extend concurrent game structures with epistemic accessibility relations $\sim_1, ..., \sim_k \subseteq Q \times Q$ (one per agent) for modeling agents' uncertainty; the relations are assumed to be equivalences. We will call such models *concurrent epistemic game structures* (CEGS) in the rest of the paper. Agent a's epistemic relation is meant to encode a's inability to distinguish between the (global) system states: $q \sim_a q'$ means that, while the system is in state q, agent a cannot determine whether it is not in q'. Then:

$M, q \models K_a\varphi$ iff φ holds for every q' such that $q \sim_a q'$.

Relations $\sim_A^E$, $\sim_A^C$ and $\sim_A^D$, used to model group epistemics, are derived from the individual relations of agents from A. First, $\sim_A^E$ is the union of relations $\sim_a$, $a \in A$. Next, $\sim_A^C$ is defined as the transitive closure of $\sim_A^E$. Finally, $\sim_A^D$ is the intersection of all the $\sim_a$, $a \in A$. The semantics of group knowledge can be defined as below (for $\mathcal{K} = C, E, D$):

$M, q \models \mathcal{K}_A\varphi$ iff φ holds for every q' such that $q \sim_A^{\mathcal{K}} q'$.

Example 1. **(Gambling Robots)** Two robots (a and b) play a simple card game. The deck consists of Ace, King and Queen (A, K, Q); it is assumed that A beats K, K beats Q, but Q beats A. First, the "environment" agent *env* deals a random card to both robots (face down), so that each player can see his own hand, but he does not know the card of the other player. Then robot a can exchange his card for the one remaining in the deck (action *exch*), or he can

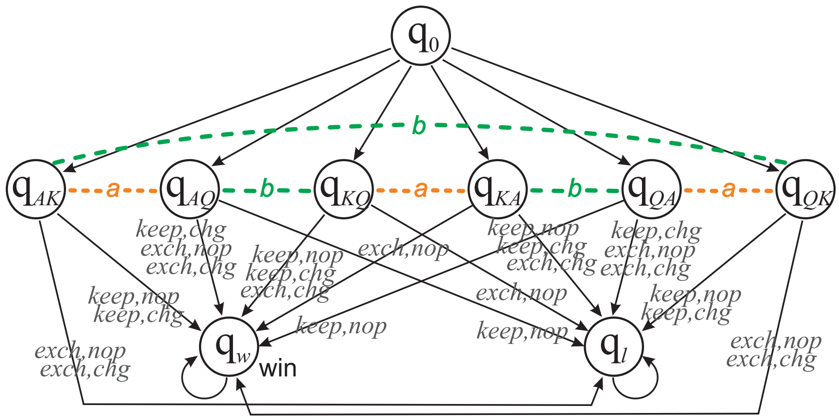

Fig. 1. Gambling Robots game. Arrows represent possible transitions of the system (labeled with tuples of agents' actions); dashed lines connect states that are indiscernible for particular agents.

keep the current one (*keep*). At the same time, robot b can change the priorities of the cards, so that A becomes better than Q (action *chg*) or he can do nothing (*nop*). If a has a better card than b after that, then a win is scored, otherwise the game ends in a "losing" state. A CEGS for the game is shown in Figure 1; we will refer to the model as M_0 throughout the rest of the paper. Note that $M_0, q_0 \models \langle\langle a \rangle\rangle \Diamond \mathsf{win}$ (and even $M_0, q_0 \models K_a \langle\langle a \rangle\rangle \Diamond \mathsf{win}$), although, intuitively, a has no feasible way of ensuring a win. This is a fundamental problem with ATEL, which we discuss briefly below.

It was pointed out in several places that the meaning of ATEL formulae is somewhat counterintuitive [5, 6, 10]. Most importantly, one would expect that an agent's ability to achieve property φ should imply that the agent has enough control and knowledge to *identify* and *execute* a strategy that enforces φ (cf. also [11]). This problem is closely related to the well known distinction between knowledge *de re* and knowledge *de dicto*.

A number of frameworks were proposed to overcome this problem [5, 6, 11, 10, 13, 4], yet none of them seems the ultimate definitive solution. Most of the solutions agree that only *uniform* strategies (i.e., strategies that specify the same choices in indistinguishable states) are really executable. However, in order to identify a successful strategy, the agents must consider not only the courses of action, starting from the current state of the system, but also from states that are indistinguishable from the current one. There are many cases here, especially when group epistemics is concerned: the agents may have common, ordinary or distributed knowledge about a strategy being successful, or they may be hinted the right strategy by a distinguished member (the "boss"), a subgroup ("headquarters committee") or even another group of agents ("consulting company"). Most existing solutions [11, 13, 4] treat only some of the cases (albeit rather in an elegant way), while others [6, 10] offer a more general treatment of the problem at the expense of an overblown logical language (which is by no means elegant).

Recently, a new, non-standard semantics for ability under incomplete information has been proposed in [8, 9], which we believe to be both intuitive, general and elegant. We summarize the proposal in the next section, as we will use it further to capture strategic abilities of agents.

2.3 An Intuitive Semantics for Ability and Knowledge

In [8, 9], a non-standard semantics for the logic of strategic ability and incomplete information has been proposed, which we believe to be finally satisfying. In the semantics, formulae are interpreted over *sets of states* rather than single states. Moreover, we introduce "constructive knowledge" operators $\mathbb{K}_a$, one for each agent a, that yield the set of states, indistinguishable from the current state from a's perspective. Constructive common, mutual, and distributed knowledge is formalized via operators $\mathbb{C}_A, \mathbb{E}_A$, and $\mathbb{D}_A$. The language, which we tentatively call Constructive Strategic Logic (CSL) here, is defined as follows:

$$\varphi ::= p \mid \neg\varphi \mid {\sim}\varphi \mid \varphi \wedge \varphi \mid \langle\langle A\rangle\rangle \bigcirc \varphi \mid \langle\langle A\rangle\rangle \Box \varphi \mid \langle\langle A\rangle\rangle \varphi \,\mathcal{U}\, \varphi \mid C_A\varphi \mid E_A\varphi \mid D_A\varphi \mid \mathbb{C}_A\varphi \mid \mathbb{E}_A\varphi \mid \mathbb{D}_A\varphi.$$

Individual knowledge operators can be derived as: $K_a\varphi \equiv E_{\{a\}}\varphi$ and $\mathbb{K}_a\varphi \equiv \mathbb{E}_{\{a\}}\varphi$. Moreover, we define $\varphi_1 \vee \varphi_2 \equiv \neg(\neg\varphi_1 \wedge \neg\varphi_2)$, and $\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$.

The models are concurrent epistemic game structures again, and we consider only memoryless uniform strategies. Let $\mathrm{img}(q, \mathcal{R})$ be the image of state q with respect to relation $\mathcal{R}$, i.e. the set of all states q' such that $q\mathcal{R}q'$. Moreover, we use $out(Q, S_A)$ as a shorthand for $\cup_{q\in Q}out(q, S_A)$, and $\mathrm{img}(Q, \mathcal{R})$ as a shorthand for $\cup_{q\in Q}\mathrm{img}(q, \mathcal{R})$. The notion of a formula φ being satisfied by a set of states $Q \subseteq St$ in a model M is given through the following clauses.

$M, Q \models p$ iff $q \in \pi(p)$ for every $q \in Q$;
$M, Q \models \neg\varphi$ iff $M, Q \not\models \varphi$;
$M, Q \models {\sim}\varphi$ iff $M, q \not\models \varphi$ for every $q \in Q$;
$M, Q \models \varphi \wedge \psi$ iff $M, Q \models \varphi$ and $M, Q \models \psi$;
$M, Q \models \langle\langle A\rangle\rangle \bigcirc \varphi$ iff there exists S_A such that, for every $\Lambda \in out(Q, S_A)$, we have that $M, \{\Lambda[1]\} \models \varphi$;
$M, Q \models \langle\langle A\rangle\rangle \Box \varphi$ iff there exists S_A such that, for every $\Lambda \in out(Q, S_A)$ and $i \geq 0$, we have $M, \{\Lambda[i]\} \models \varphi$;
$M, Q \models \langle\langle A\rangle\rangle \varphi \,\mathcal{U}\, \psi$ iff there exists S_A such that, for every $\Lambda \in out(Q, S_A)$, there is $i \geq 0$ for which $M, \{\Lambda[i]\} \models \psi$ and $M, \{\Lambda[j]\} \models \varphi$ for every $0 \leq j < i$;
$M, Q \models \mathcal{K}_A\varphi$ iff $M, q \models \varphi$ for every $q \in \mathrm{img}(Q, \sim_A^{\mathcal{K}})$ (where $\mathcal{K} = C, E, D$);
$M, Q \models \hat{\mathcal{K}}_A\varphi$ iff $M, \mathrm{img}(Q, \sim_A^{\mathcal{K}}) \models \varphi$ (where $\hat{\mathcal{K}} = \mathbb{C}, \mathbb{E}, \mathbb{D}$ and $\mathcal{K} = C, E, D$, respectively).

We will also write $M, q \models \varphi$ as a shorthand for $M, \{q\} \models \varphi$, and this is the notion of satisfaction (in single states) that we are ultimately interested in – but that notion is defined in terms of the satisfaction in sets of states.

Now, $\mathbb{K}_a\langle\langle a\rangle\rangle\varphi$ expresses the fact that a has a single strategy that enforces φ from *all* states indiscernible from the current state, instead of stating that φ can be achieved from *every* such state *separately* (what $K_a\langle\langle a\rangle\rangle\varphi$ says, which is very much in the spirit of standard epistemic logic). More generally, the first kind of formulae refer to *having a strategy "de re"* (i.e. having a successful strategy and knowing the strategy), while the latter refer to *having a strategy "de dicto"* (i.e. only knowing that *some* successful strategy is available; cf. [6]). Note also that the property of having a winning strategy in the current state (but not necessarily even knowing *about* it) is simply expressed with $\langle\langle a\rangle\rangle\varphi$. Capturing different ability levels of coalitions is analogous, with various "epistemic modes" of collective recognizing the right strategy.

Example 2. Robot a has no winning strategy in the starting state of the game: $M_0, q_0 \models \neg\langle\langle a\rangle\rangle\Diamond\mathsf{win}$, which implies that it has neither a strategy "de re" nor "de dicto" ($M_0, q_0 \models \neg\mathbb{K}_a\langle\langle a\rangle\rangle\Diamond\mathsf{win} \wedge \neg K_a\langle\langle a\rangle\rangle\Diamond\mathsf{win}$). On the other hand, he has a successful strategy in q_{AK} (just play *keep*) and he knows he has one (because another action, $exch$, is bound to win in q_{AQ}); still, the knowledge is not constructive, since a does not know which strategy is the right one in the current situation: $M_0, q_{AK} \models \langle\langle a\rangle\rangle\bigcirc\mathsf{win} \wedge K_a\langle\langle a\rangle\rangle\bigcirc\mathsf{win} \wedge \neg\mathbb{K}_a\langle\langle a\rangle\rangle\bigcirc\mathsf{win}$. Also, b's playing chg enforces a transition to q_w for both q_{AQ}, q_{KQ}, so $M_0, q_{AQ} \models \mathbb{K}_b\langle\langle b\rangle\rangle\bigcirc\mathsf{win}$ (robot b has a strategy "de re" to enforce a win from q_{AQ}).

Finally, $q_{QK} \models \langle\langle a,b\rangle\rangle\Diamond\mathsf{win} \wedge E_{\{a,b\}}\langle\langle a,b\rangle\rangle\Diamond\mathsf{win} \wedge C_{\{a,b\}}\langle\langle a,b\rangle\rangle\Diamond\mathsf{win} \wedge \neg\mathbb{E}_{\{a,b\}}\langle\langle a,b\rangle\rangle\Diamond\mathsf{win} \wedge \mathbb{D}_{\{a,b\}}\langle\langle a,b\rangle\rangle\Diamond\mathsf{win}$: in q_{QK}, the robots have a collective strategy to enforce a win, and they all know it (they even have common knowledge about it); on the other hand, they cannot identify the right strategy as a team – they can only see one if they share knowledge at the beginning (i.e., in q_{QK}).

3 Epistemic Temporal Strategic Logic

A very interesting variation on the theme of combining strategic, epistemic and temporal aspects of a multi-agent system was proposed in [13]. Epistemic Temporal Strategic Logic (ETSL) digs deeper in the repository of game theory, and focuses on the concept of *undominated strategies*. Thus, its variant of cooperation modalities has a different flavor than the ones from ATL, ATEL, CSL etc. In a way, formula $\langle\langle A\rangle\rangle\varphi$ in ETSL can be summarized as:

"If A play *rationally* to achieve φ (meaning: they never play a dominated strategy), they will achieve φ".

ETSL can be treated as a logic that describes the outcome of *rational play* under incomplete information,[2] in the same way as CSL can be seen as a logic that captures agents' strategic abilities (regardless of whether the agents play rationally

[2] We emphasize that this is a specific notion of rationality (i.e., agents are assumed to *play only undominated strategies*). Game theory proposes several other rationality criteria as well, based e.g. on Nash equilibrium, dominant strategies, or Pareto efficiency. In fact, it is easy to imagine ETSL-like logics based on these notions instead.

or not). The main claim we propose in this paper is that a rational player knows that he will succeed if, and only if, he has a strategy "de re" to succeed – while the same is not true for rational coalitions of players. However, before we present and discuss the claim formally in Section 4, we must re-write the semantics of ETSL in several respects.

First, the original semantics of ETSL is defined only for finite turn-based acyclic game models with epistemic accessibility relations, and we will generalize the semantics to concurrent epistemic game structures. Next, the semantics comes with a plethora of auxiliary functions and definitions (and a couple of omissions), which makes it rather hard to read. In fact, this is probably the reason why the logic never received the attention it deserves, and it is definitely worth trying to make the semantics more compact. Finally, the authors of [13] propose that a model should include also a "grand strategy profile" $S_{\mathbb{A}\mathrm{gt}}$, defining the actual strategies of all agents (or at least constraining them in some way, since non-deterministic strategies are allowed in ETSL). While the idea seems interesting in itself (a similar idea was later exploited e.g. in [7] to allow for explicit analysis of strategies and reasoning about strategy revision), we will show that it does not introduce a finer-grained analysis of "vanilla" ETSL formulas: if a formula holds in M, q for one strategy profile, it holds in M, q for all the other strategy profiles, too. Moreover, it can be proved that the semantics of cooperation modalities $\langle\langle A\rangle\rangle$ is the same regardless of whether we consider non-deterministic strategies or not. In consequence, we will be able to show a "vanilla" ETSL semantics expressed entirely in terms of concurrent epistemic game structures and their states.

3.1 The Semantics Made Easier to Read

Formulae of ETSL come with no restriction wrt grouping of temporal operators:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid \langle\langle A\rangle\rangle\varphi \mid \bigcirc\varphi \mid \Box\varphi \mid \varphi\,\mathcal{U}\,\psi \mid K_a\varphi.$$

After some re-writing (and having it generalized to general game structures, not only turn-based trees), the semantics can be given as follows. Strategies are allowed to be non-deterministic, i.e. $S_a : St \to \mathcal{P}(Act)$.[3] We require strategies to be uniform, although [13] does not do it explicitly (we take it as a simple omission, because otherwise many claims in that paper seem to be false). A collective strategy (strategy profile) S_A is a tuple of strategies, one per agent from A. S_a^0 is the "neutral strategy" with no restriction on a's actions ($S_a^0(q) = Act$ for each $q \in St$), and strategy profile S_A^0 assigns neutral strategies to agents from A. Moreover, we generalize function $out(q, S_A)$ to handle nondeterministic strategies too; in $out'(q, S_A)$, "$\alpha_a^{i-1} = S_a(q_{i-1})$" is replaced with $\alpha_a^{i-1} \in S_a(q_{i-1})$.
Now, the semantics can be given through the following clauses (the semantics for p, $\neg\varphi$ and $\varphi \wedge \psi$ is analogous to the one presented in Section 2.1):

[3] To preserve seriality ("time flows forever"), we assume that $S_a(q) \neq \varnothing$ for all $q \in St$.

$M, S_{\mathbb{A}\mathrm{gt}}, q \models \langle\!\langle A\rangle\!\rangle\varphi$ iff for all strategies T_A, undominated wrt q, φ, we have $M, (T_A, S^0_{\mathbb{A}\mathrm{gt}\setminus A}), q \models \varphi$;

$M, S_{\mathbb{A}\mathrm{gt}}, q \models \bigcirc\varphi$ iff for every $\Lambda \in out'(q, S_{\mathbb{A}\mathrm{gt}})$ we have $M, S_{\mathbb{A}\mathrm{gt}}, \Lambda[1] \models \varphi$;

$M, S_{\mathbb{A}\mathrm{gt}}, q \models \Box\varphi$ iff for every $\Lambda \in out'(q, S_{\mathbb{A}\mathrm{gt}})$ and $i \geq 0$ we have $M, S_{\mathbb{A}\mathrm{gt}}, \Lambda[i] \models \varphi$;

$M, S_{\mathbb{A}\mathrm{gt}}, q \models \varphi\,\mathcal{U}\,\psi$ iff for every $\Lambda \in out'(q, S_{\mathbb{A}\mathrm{gt}})$ there is $i \geq 0$ such that $M, S_{\mathbb{A}\mathrm{gt}}, \Lambda[i] \models \psi$ and for all j such that $0 \leq j < i$ we have $M, S_{\mathbb{A}\mathrm{gt}}, \Lambda[j] \models \varphi$;

$M, S_{\mathbb{A}\mathrm{gt}}, q \models K_a\varphi$ iff for all $q \sim_a q'$ we have $M, (S_{\mathbb{A}\mathrm{gt}}(a), S^0_{\mathbb{A}\mathrm{gt}\setminus\{a\}}), q' \models \varphi$.

Definition 1. *Strategy S_A* dominates *T_A with respect to formula φ, model M, and state q, if S_A achieves φ better then T_A, i.e. iff:*

1. *for every q' such that $q \sim_A q'$: if $M, (T_A, S^0_{\mathbb{A}\mathrm{gt}\setminus A}), q' \models \varphi$ then also $M, (S_A, S^0_{\mathbb{A}\mathrm{gt}\setminus A}), q' \models \varphi$, and*
2. *there exists q' such that $q \sim_A q'$, and $M, (S_A, S^0_{\mathbb{A}\mathrm{gt}\setminus A}), q' \models \varphi$, and $M, (T_A, S^0_{\mathbb{A}\mathrm{gt}\setminus A}), q \not\models \varphi$.*

Remark 1. Definition 1 uses epistemic relation $\sim_A$. However, epistemic accessibility relations are defined only for individual agents in [13], which is perhaps another omission. In this study, we take the liberty to fix $\sim_A$ as $\sim^E_A$.

We also point out that ETSL can be extended with collective epistemic operators E_A, C_A, D_A in a straightforward manner.

Example 3. Consider the gambling robots again. Robot a has two undominated strategies wrt $\bigcirc\mathsf{win}, M, q_{AK}$: namely, to play *exch* in both q_{AK}, q_{AQ}, or to play *keep* in both (other choices do not matter). Since playing *exch* fails in q_{AK}, so: $M_0, q_{AK} \not\models \langle\!\langle a\rangle\!\rangle\bigcirc\mathsf{win}$. Furthermore, playing *keep* is the only undominated strategy in q_{KQ} and q_{KA} (and it succeeds only in q_{KQ}). Thus, $M_0, q_{KQ} \models \langle\!\langle a\rangle\!\rangle\bigcirc\mathsf{win}$, and $M_0, q_{KA} \not\models \langle\!\langle a\rangle\!\rangle\bigcirc\mathsf{win}$. Hence, $M_0, q_{KQ} \not\models K_a\langle\!\langle a\rangle\!\rangle\bigcirc\mathsf{win}$.

3.2 A Few Properties

In this section, we present several properties of ETSL formulae that will allow us to give an even simpler semantic definition of "vanilla" ETSL.

Proposition 1. *For every "vanilla" ETSL formula φ, concurrent epistemic game structure M, and state q in M: $M, S_{\mathbb{A}\mathrm{gt}}, q \models \varphi$ iff $M, S'_{\mathbb{A}\mathrm{gt}}, q \models \varphi$ for any pair of "grand" strategy profiles $S_{\mathbb{A}\mathrm{gt}}, S'_{\mathbb{A}\mathrm{gt}}$.*

Proof. By induction on the structure of φ. Note that it is sufficient to prove the implication one way, as the choice of $S_{\mathbb{A}\mathrm{gt}}, S'_{\mathbb{A}\mathrm{gt}}$ is completely arbitrary.

Case $\varphi \equiv p$: $M, S_{\mathbb{A}\mathrm{gt}}, q \models p$, so $q \in \pi(q)$, so $M, S'_{\mathbb{A}\mathrm{gt}}, q \models p$.

Case $\varphi \equiv \neg\psi$: $M, S_{\mathbb{A}\mathrm{gt}}, q \models \neg\psi$, so $M, S_{\mathbb{A}\mathrm{gt}}, q \not\models \psi$, so (by induction hypothesis) $M, S'_{\mathbb{A}\mathrm{gt}}, q \not\models \psi$, so $M, S'_{\mathbb{A}\mathrm{gt}}, q \models \neg\psi$. (As the choice of $S_{\mathbb{A}\mathrm{gt}}, S'_{\mathbb{A}\mathrm{gt}}$ was completely arbitrary, the implication holds the other way too.)

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

Case $\varphi \equiv \langle\langle A\rangle\rangle \bigcirc \psi$: $M, S_{\mathbb{A}gt}, q \models \langle\langle A\rangle\rangle \bigcirc \psi$ iff $M, (T_A, S^0_{\mathbb{A}gt\setminus A}), \Lambda[1] \models \varphi$ for all undominated T_A and $\Lambda \in out'(q, (T_A, S^0_{\mathbb{A}gt\setminus A}))$. Note that the latter condition does not refer to $S_{\mathbb{A}gt}$, so $M, S'_{\mathbb{A}gt}, q \models \langle\langle A\rangle\rangle \bigcirc \psi$ too.

Cases $\varphi \equiv \langle\langle A\rangle\rangle \Box \psi$ and $\varphi \equiv \langle\langle A\rangle\rangle \psi_1 \mathcal{U} \psi_2$: analogous.

Case $\varphi \equiv K_a\psi$: $M, S_{\mathbb{A}gt}, q \models K_a\psi$, so $M, (S_{\mathbb{A}gt}(a), S^0_{\mathbb{A}gt\setminus\{a\}}), q' \models \psi$ for all $q \sim_a q'$. By induction hypothesis, also $M, (S'_{\mathbb{A}gt}(a), S^0_{\mathbb{A}gt\setminus\{a\}}), q' \models \psi$ for all $q \sim_a q'$, so $M, S'_{\mathbb{A}gt}, q \models K_a\psi$.

Remark 2. We point out that restricting the scope of Proposition 1 to "vanilla" ETSL formulae is important. In particular, the epistemic opertor K_a has a non-standard interpretation when the full language of ETSL is considered.

Proposition 2. *Let $\Phi \equiv \bigcirc\psi, \Box\psi$, or $\psi_1 \mathcal{U} \psi_2$ where ψ, ψ_1, ψ_2 are "vanilla" ETSL formulae. Moreover, let $|\Phi|$ denote the set of paths for which Φ holds; formally, $|\bigcirc\psi| = \{\Lambda \mid M, \Lambda[1] \models \psi\}$, $|\Box\psi| = \{\Lambda \mid \forall_i M, \Lambda[i] \models \psi\}$, and $|\psi_1 \mathcal{U} \psi_2| = \{\Lambda \mid \exists_i (M, \Lambda[i] \models \psi_2 \wedge \forall_{0\le j<i} M, \Lambda[j] \models \psi_1\}$.*

Then, S_A dominates T_A wrt Φ, M, and q iff:

1. *for every q', $q \sim^E_A q'$: if $out(q', T_A) \subseteq |\Phi|$ then also $out(q', S_A) \subseteq |\Phi|$, and*
2. *there exists q', $q \sim^E_A q'$, such that $out(q', S_A) \subseteq |\Phi|$ and $out(q', T_A) \not\subseteq |\Phi|$.*

Proof. Straightforward from the definition.

Remark 3. Note that dominance can be characterized in an even more compact way. Let $succ_{q,\Phi}(S_A) = \{q \in \mathrm{img}(q, \sim^E_A) \mid out(q, S_A) \subseteq |\Phi|\}$ be the set of states from $\mathrm{img}(q, \sim^E_A)$, for which s_a succeeds to enforce Φ. Now, S_A dominates T_A wrt Φ, M, q iff $succ_{q,\Phi}(T_A) \subsetneq succ_{q,\Phi}(S_A)$.

Proposition 3. *Let $\Phi \equiv \bigcirc\psi, \Box\psi$, or $\psi_1 \mathcal{U} \psi_2$ where ψ, ψ_1, ψ_2 are "vanilla" ETSL formulae. Strategy T_A is dominated wrt Φ, M, q by a strategy S_A iff it is dominated wrt Φ, M, q by a* deterministic *strategy S'_A.*

Proof. $\Rightarrow$: Let T_A be dominated by S_A (wrt φ, M, q). We construct the deterministic strategy S'_A by fixing arbitrary (uniform) choices out of S_A. Formally, for every agent $a \in A$ and abstraction class $\mathrm{img}(q', \sim_a) \subseteq St$ such that $S_a(q') = \{\alpha, \alpha', ...\}$, we fix $S'_a(q'') = \alpha$ for all $q'' \in \mathrm{img}(q', \sim_a)$. (By uniformity of S_A, we have $\alpha \in S_a(q'')$ for all $q'' \in \mathrm{img}(q', \sim_a)$, so S'_A is a valid strategy.) First, this enforces uniformity of S'_A. Second, $out(\bar{q}, S'_A) \subseteq out(\bar{q}, S_A)$ for all $\bar{q} \in St$ (by definition of out). Thus, we can use Proposition 2 to show that S'_A dominates T_A, which concludes the proof.

$\Leftarrow$: Straightforward.

Proposition 4. *Let Φ be as above. Then, $M, S_{\mathbb{A}gt}, q \models \langle\langle A\rangle\rangle\Phi$ iff for all* deterministic *strategies T_A, undominated wrt Φ, we have $M, (T_A, S^0_{\mathbb{A}gt\setminus A}), q \models \Phi$.*

Proof. $\Rightarrow$: Straightforward.

$\Leftarrow$: Assume that $M, (T_A, S^0_{\mathbb{A}\mathrm{gt}\setminus A}), q \models \Phi$ for all deterministic strategies T_A, undominated wrt Φ, and suppose that there is a nondeterministic undominated S_A such that $M, (S_A, S^0_{\mathbb{A}\mathrm{gt}\setminus A}), q \not\models \Phi$. Let us fix a deterministic uniform strategy S'_A out of S_A in a similar way as in Proposition 3. Now, $out(\bar{q}, S'_A) \subseteq out(\bar{q}, S_A)$ for all $\bar{q} \in St$, so $out(q', S_A) \subseteq |\Phi|$ implies $out(q', S'_A) \subseteq |\Phi|$ (S'_A is never worse than S_A wrt Φ). Moreover, $out(q, S'_A) \subseteq |\Phi|$ and $out(q, S_A) \not\subseteq |\Phi|$. By Proposition 2, S'_A dominates S_A, so S_A is dominated – a contradiction.

3.3 ETSL in Terms of Concurrent Epistemic Game Structures

We have shown that, for "vanilla" ETSL, strategies do not have to be referred explicitly in the interpretation of formulae (Propositions 1 and 2). Moreover, we can restrict the set of considered strategies to deterministic strategies (Propositions 3 and 4). In consequence, we can express the semantics of "vanilla" ETSL equivalently in ATL-like fashion:

$M, q \models \langle\langle A\rangle\rangle \bigcirc\varphi$ iff for every strategy S_A, undominated wrt $q, \bigcirc\varphi$, and every $\Lambda \in out(q, S_A)$, we have that $M, \Lambda[1] \models \varphi$;

$M, q \models \langle\langle A\rangle\rangle \Box\varphi$ iff for every strategy S_A, undominated wrt $q, \Box\varphi$, and every $\Lambda \in out(q, S_A)$ and $i \geq 0$ we have $M, \Lambda[i] \models \varphi$;

$M, q \models \langle\langle A\rangle\rangle \varphi\,\mathcal{U}\,\psi$ iff for every strategy S_A, undominated wrt $q, \varphi\,\mathcal{U}\,\psi$, and every $\Lambda \in out(q, S_A)$, there is $i \geq 0$ such that $M, \Lambda[i] \models \psi$ and for all j such that $0 \leq j < i$ we have $M, \Lambda[j] \models \varphi$.

Only uniform deterministic strategies are taken into account. The semantics of p, $\neg\varphi$, $\varphi \wedge \psi$, and the epistemic operators is the same as for ATL and ATEL.

4 Playing Rationally vs. Knowing how to Play

We can finally present the main result of this paper, namely, that a rational player knows that he will succeed if, and only if, he has a strategy "de re" to succeed. The result holds under the assumption that the model is finite,[4] or more generally, that it includes at least one undominated strategy.

Moreover, we show that having common knowledge how to succeed is, in general, a stronger property than knowing that one will succeed for rational coalitions of players. That is, if rational agents have common knowledge about a winning strategy, then they have common knowledge that they will succeed – but the converse is not true any more. Surprisingly enough, it turns out that the relationship is strictly reverse for distributed knowledge: if a rational coalition has distributed knowledge that it will succeed, then it has distributed knowledge about a winning strategy – but not necessarily the other way around. For mutual knowledge, the relationship holds neither way.

In what follows, we use $\models_{\mathrm{ETSL}}$ and $\models_{\mathrm{CSL}}$ to denote the ETSL and CSL satisfaction relation, respectively.

[4] We use the term "finite model" to denote a CEGS with a *finite set of states St*.

4.1 Rational Play of Individual Agents

We begin with two important lemmas.

Lemma 1. *Given a finite model M, state q in M, formula Φ and agent a, there is a strategy s_a which is undominated wrt M, q, Φ.*

Proof. First, we consider the simpler case when the set of actions Act is finite. In such a case, the set of strategies is also finite, and the dominance relation is transitive and antireflexive. Suppose that every strategy is dominated; then, there must be a strategy which is dominated by itself – a contradiction.

We sketch the proof for infinite Act as follows. We partition the infinite set of strategies into equivalence classes, such that strategies in the same class have the same outcome paths for every state q (i.e., $s_a \approx t_a$ iff $\forall_q out(q, s_a) = out(q, t_a)$). Obviously, if s_a dominates t_a, then all strategies $s'_a \approx s_a$ dominate t_a too. Now, at every state q (and therefore at every point on a path from $out(q', s_a)$) there is a finite number of possible sets of successor states (the actual set being determined by the choice $s_a(q)$). Moreover, the same choice must be taken at every further occurrence of the same state q on a path, since s_a is a memoryless strategy. In consequence, there is only a finite number of different sets of outcome paths, and hence a finite number of the equivalence classes. Again, dominance is transitive and antireflexive, so an undominated strategy must exist.

Remark 4. Note that the result in Lemma 1 does not extend to CEGS with infinite state spaces. Consider the game of "Fuzzy Blackjack" (called so all the more because our robots play it usually after having consumed too much machine oil). Only a single player is necessary, and we use positive real numbers as states and actions (i.e., $St = Act = \mathbb{R}_+$). When the player chooses a number in state q, the number is added to the state: $o(q, \alpha) = q + \alpha$. The values below 1 are the winning ones, i.e. $\pi(\mathsf{win}) = (0, 1)$ (it should be 21, but this would make the game too complicated for a drunken robot). Moreover, the robot cannot distinguish between the states below 1: $q \sim_a q'$ for all $q, q' \in (0, 1)$. Now, there is no undominated strategy wrt $0.5, \bigcirc\mathsf{win}$.

To prove this, suppose that a strategy s_a is undominated. The strategy is uniform, so $s_a(q) = \alpha$ for some $\alpha \in \mathbb{R}_+$ and all $q \in (0, 1)$. Obviously, $\alpha \in (0, 1)$, because else s_a never succeeds. Now, the set of states in which s_a is successful is: $succ_{0.5,\bigcirc\mathsf{win}}(s_a) = (0, 1 - \alpha)$. Let $t_a(q) = q + \alpha/2$. Now, $succ_{0.5,\bigcirc\mathsf{win}}(t_a) = (0, 1 - \alpha/2) \supsetneq succ_{0.5,\Phi}(s_a)$ – a contradiction. Note also that:

- If we replace $\mathbb{R}_+$ with the set of positive rational numbers, the result is the same. So, there may be no undominated strategies even when we restrict St and Act to countable sets.
- In order to show the same for countable St and *finite* Act, it is sufficient to modify the example so that $Act = \{0, 1, call\}$, and the initial state and every subsequent action $\alpha = 0, 1$ are simply stored in the resulting state. Now $o(q, call)$ takes the initial state q_0 and the string of 0s and 1s $\alpha_1, ..., \alpha_n$ stored in q, and returns $q' = q_0 + (0.\alpha_1...\alpha_n 1)_2$. For such a game, there is no undominated strategy wrt $0.5, \Diamond\mathsf{win}$.

Lemma 2. *Given M, q, Φ, a, if there is an undominated strategy wrt M, q, Φ, then there is also an undominated strategy wrt M, q', Φ for every $q' \in \mathrm{img}(q, \sim_a)$.*

Proof. Take any s_a undominated wrt M, q, Φ (*). Suppose now that s_a is dominated by some strategy t_a wrt another state $q' \in \mathrm{img}(q, \sim_a)$ (**).

1. By (*) and Prop. 2: $\forall_{q'' \in \mathrm{img}(q, \sim_a)} \; (out(q'', t_a) \subseteq |\Phi| \Rightarrow out(q'', s_a) \subseteq |\Phi|)$.
2. By (**) and Prop. 2: $\exists_{q'' \in \mathrm{img}(q', \sim_a)} \; (out(q'', t_a) \subseteq |\Phi| \wedge out(q'', s_a) \not\subseteq |\Phi|)$.

Moreover, $\mathrm{img}(q, \sim_a) = \mathrm{img}(q', \sim_a)$ because is $\sim_a$ is an equivalence relation – which gives a contradiction between (1) and (2).

Remark 5. We note that Lemma 2 may hold even for indistinguishability relations that are not equivalences. In fact, it is sufficient to require that $\sim_a$ is transitive. In that case, $q' \in \mathrm{img}(q, \sim_a)$ and $q'' \in \mathrm{img}(q', \sim_a)$ implies that $q'' \in \mathrm{img}(q, \sim_a)$, and we also get the contradiction.

We are ready to prove the main claim of this paper now.

Theorem 1. *Let us consider only finite models, and formulae $\Phi \equiv \bigcirc\psi, \Box\psi$, or $\psi_1 \,\mathcal{U}\, \psi_2$ where ψ, ψ_1, ψ_2 are "vanilla" ETSL formulae. An agent has a strategy "de re" to enforce Φ if, and only if, he knows that his rational play will bring about Φ. Formally, for every finite M and state q in M:*

$$M, q \models_{\mathrm{ETSL}} K_a\langle\langle a\rangle\rangle\Phi \quad \textit{iff} \quad M, q \models_{\mathrm{CSL}} \mathbb{K}_a\langle\langle a\rangle\rangle\Phi.$$

Proof. Induction on the structure of Φ. We prove the theorem for the case $\Phi \equiv \Box\psi$. Other cases are analogous.

$\Rightarrow$: Let $M, q \models_{\mathrm{ETSL}} K_a\langle\langle a\rangle\rangle\Box\psi$. Then, $\forall_{q' \in \mathrm{img}(q, \sim_a)} M, q' \models_{\mathrm{ETSL}} \langle\langle a\rangle\rangle\Box\psi$, and hence $M, q \models_{\mathrm{ETSL}} \langle\langle a\rangle\rangle\Box\psi$ in particular. By Lemmas 1 and 2, there is a strategy s_a, undominated wrt $M, q', \Box\psi$ for every $q' \in \mathrm{img}(q, \sim_a)$.
Then: $\forall_{q' \in \mathrm{img}(q, \sim_a)} \forall_{\Lambda \in out(q', s_a)} \forall_i M, \Lambda[i] \models_{\mathrm{ETSL}} \Box\psi$. By the induction hypothesis, also $\forall_{q' \in \mathrm{img}(q, \sim_a)} \forall_{\Lambda \in out(q', s_a)} \forall_i M, \Lambda[i] \models_{\mathrm{CSL}} \psi$. Thus, $\forall_{\Lambda \in out(\mathrm{img}(q, \sim_a), s_a)} \forall_i$ $M, \Lambda[i] \models_{\mathrm{CSL}} \psi$ and so $M, \mathrm{img}(q, \sim_a) \models_{\mathrm{CSL}} \langle\langle a\rangle\rangle\Box\psi$, and finally $M, q \models_{\mathrm{CSL}} \mathbb{K}_a\langle\langle a\rangle\rangle\Box\psi$.

$\Leftarrow$: Let $M, q \models_{\mathrm{CSL}} \mathbb{K}_a\langle\langle a\rangle\rangle\Box\psi$, i.e. $M, \mathrm{img}(q, \sim_a) \models_{\mathrm{CSL}} \langle\langle a\rangle\rangle\Box\psi$. Consider $q' \in \mathrm{img}(q, \sim_a)$. By transitivity of $\sim_a$, we have $\mathrm{img}(q', \sim_a) \subseteq \mathrm{img}(q, \sim_a)$, so also $\forall_{q' \in \mathrm{img}(q, \sim_a)} M, \mathrm{img}(q', \sim_a) \models_{\mathrm{CSL}} \langle\langle a\rangle\rangle\Box\psi$. Then, for every $q' \in \mathrm{img}(q, \sim_a)$, there must be s_a such that $\forall_{q'' \in \mathrm{img}(q', \sim_a)} \forall_{\Lambda \in out(q'', s_a)} \forall_i M, \Lambda[i] \models_{\mathrm{CSL}} \psi$, and hence (by induction) $\forall_{q'' \in \mathrm{img}(q', \sim_a)} \forall_{\Lambda \in out(q'', s_a)} \forall_i M, \Lambda[i] \models_{\mathrm{ETSL}} \psi$. So, $succ_{q', \Box\psi}(s_a) = \mathrm{img}(q', \sim_a)$, and therefore $succ_{q', \Box\psi}(t_a) = \mathrm{img}(q', \sim_a)$ for every other undominated strategy t_a (otherwise t_a would be dominated by s_a). Thus, $M, q' \models_{\mathrm{ETSL}} \langle\langle a\rangle\rangle\Box\psi$ for every $q' \in \mathrm{img}(q, \sim_a)$, and finally $M, q \models_{\mathrm{ETSL}} K_a\langle\langle a\rangle\rangle\Box\psi$.

Theorem 2. *More generally, for every Φ as above, and M, q such that there exists an undominated strategy wrt M, q, Φ:* $\;M, q \models_{\mathrm{ETSL}} K_a\langle\langle a\rangle\rangle\Phi$ *iff* $M, q \models_{\mathrm{CSL}} \mathbb{K}_a\langle\langle a\rangle\rangle\Phi$.

4.2 Rational Coalitions Are at Disadvantage

Beside some philosophical insight into the nature of knowledge and rational play, Theorems 1 and 2 provide us with an alternative way of decomposing strategic abilities under incomplete information into a strategic and epistemic part. The definition of the strategic dimension is more sophisticated and less straightforward than usually; on the other hand, we do not pay the price of a non-standard satisfaction relation. Unfortunately, such decomposition is not valid any more when abilities of collective agents are concerned. Now, the relationship is much more limited: if a coalition has *common* knowledge how to play, then it has also common knowledge that rational play will be successful; the same does *not* hold for other types of collective knowledge. Moreover, the converse relationship is guaranteed for distributed knowledge, but *not* for common nor mutual knowledge.

Theorem 3. *Let $\Phi \equiv \bigcirc\psi, \Box\psi$, or $\psi_1 \mathcal{U} \psi_2$ where ψ, ψ_1, ψ_2 are "vanilla" ETSL formulae. Then, if a coalition has common knowledge how to play, then it has common knowledge that rational play will be successful:*

$$\text{if} \quad M, q \models_{\text{CSL}} \mathbb{C}_A\langle\!\langle A\rangle\!\rangle\Phi \quad \text{then} \quad M, q \models_{\text{ETSL}} C_A\langle\!\langle A\rangle\!\rangle\Phi.$$

The same holds for neither mutual nor distributed knowledge.

Proof. **Common knowledge:** Let $M, q \models_{\text{CSL}} \mathbb{K}_A\langle\!\langle A\rangle\!\rangle\Box\psi$, i.e. $M, \text{img}(q, \sim_A^C) \models_{\text{CSL}} \langle\!\langle A\rangle\!\rangle\Box\psi$. Consider $q' \in \text{img}(q, \sim_A^C)$. We have $\text{img}(q', \sim_A^E) \subseteq \text{img}(q', \sim_A^C) \subseteq \text{img}(q, \sim_A^C)$, so also $\forall_{q' \in \text{img}(q, \sim_A^C)} M, \text{img}(q', \sim_A^E) \models_{\text{CSL}} \langle\!\langle A\rangle\!\rangle\Box\psi$. Then, for every $q' \in \text{img}(q, \sim_A^C)$, there must be S_A such that $\forall_{q'' \in \text{img}(q', \sim_A^E)} \forall_{\Lambda \in out(q'', S_A)} \forall_i M, \Lambda[i] \models_{\text{CSL}} \psi$, and hence (by induction) $\forall_{q'' \in \text{img}(q', \sim_A^E)} \forall_{\Lambda \in out(q'', S_A)} \forall_i M, \Lambda[i] \models_{\text{ETSL}} \psi$. So, $succ_{q', \Box\psi}(S_A) = \text{img}(q', \sim_A^E)$, and therefore $succ_{q', \Box\psi}(T_A) = \text{img}(q', \sim_A^E)$ for every other undominated strategy T_A (otherwise T_A would be dominated by S_A). Thus, $M, q' \models_{\text{ETSL}} \langle\!\langle A\rangle\!\rangle\Box\psi$ for every $q' \in \text{img}(q, \sim_A^C)$, and finally $M, q \models_{\text{ETSL}} C_A\langle\!\langle A\rangle\!\rangle\Box\psi$.

Mutual knowledge: for a counterexample, consider a modification of the game from Figure 1, in which a third robot c is introduced. The robot can only execute *nop*, and its epistemic relation $\sim_c = \{(q, q) \mid q \in St\} \cup \{(q_{KQ}, q_{KA}), (q_{KA}, q_{KQ})\}$, i.e. c can distinguish all states except q_{KQ}, q_{KA}. Moreover, the transition function is slightly changed: now, $o(q_{KA}, keep, nop) = q_w$. For the resulting system M_1, we have that $M_1, q_{AQ} \models_{\text{CSL}} \mathbb{E}_{\{b,c\}}\langle\!\langle b, c\rangle\!\rangle\bigcirc\mathsf{win}$, but at the same time $M_1, q_{AQ} \not\models_{\text{ETSL}} E_{\{a,c\}}\langle\!\langle a, c\rangle\!\rangle\bigcirc\mathsf{win}$ because $M_1, q_{KQ} \not\models_{\text{ETSL}} \langle\!\langle a, c\rangle\!\rangle\bigcirc\mathsf{win}$.

Distributed knowledge: analogously, $M_1, q_{KQ} \models_{\text{CSL}} \mathbb{D}_{\{b,c\}}\langle\!\langle b, c\rangle\!\rangle\bigcirc\mathsf{win}$, yet $M_1, q_{KQ} \not\models_{\text{ETSL}} D_{\{a,c\}}\langle\!\langle a, c\rangle\!\rangle\bigcirc\mathsf{win}$ because $M_1, q_{KQ} \not\models_{\text{ETSL}} \langle\!\langle a, c\rangle\!\rangle\bigcirc\mathsf{win}$.

Theorem 4. *Let $\Phi \equiv \bigcirc\psi, \Box\psi$, or $\psi_1 \mathcal{U} \psi_2$ where ψ, ψ_1, ψ_2 are "vanilla" ETSL formulae, and let M be a finite CEGS.*[5] *Then, if A have distributed knowledge*

[5] Alternatively, we can request that A have at least one undominated strategy for every relevant state.

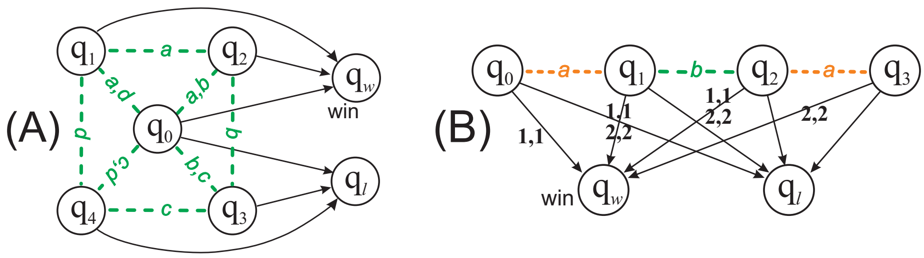

Fig. 2. **(A)** Model M_2: four agents a, b, c, d, epistemic relations shown with the dashed lines, $Act = \{1, 2, 3, 4\}$. Transitions: $o(q_i, j, j, j, j) = q_w$ for $j \neq i$, otherwise the system proceeds to the "losing" state q_l; **(B)** Model M_3: two agents a, b, two actions $1, 2$. The tuples of actions that are absent in the graph lead to q_l.

that rational play will bring about Φ, then they have distributed knowledge how to play to bring about Φ. Formally:

$$\textit{if} \quad M, q \models_{\text{ETSL}} D_A\langle\langle A\rangle\rangle\Phi \quad \textit{then} \quad M, q \models_{\text{CSL}} \mathbb{D}_A\langle\langle A\rangle\rangle\Phi.$$

The same holds for neither mutual nor common knowledge.

Proof. (sketch) **Distributed knowledge:** the proof is analogous to the proofs of Lemma 2 and Theorem 1 (part $\Rightarrow$), as we can exploit the fact that $\sim_A^D$ is transitive, and $\text{img}(q, \sim_A^D) \subseteq \text{img}(q, \sim_A^E)$.

Mutual knowledge: for a counterexample, consider model M_2 from Figure 2A. Let $\overline{q}$ denote the state "opposite" to q, i.e. $\overline{q_1} = q_3$, $\overline{q_2} = q_4$ etc. Furthermore, let $S^i_{\mathbb{A}\text{gt}}$ denote the strategy of playing $\langle i, i, i, i\rangle$ in all states. Now, $S^i_{\mathbb{A}\text{gt}}$ is the only undominated strategy wrt $\overline{q_i}, \bigcirc\mathsf{win}$ for $i = 1, ..., 4$, and $S^1_{\mathbb{A}\text{gt}}, ..., S^4_{\mathbb{A}\text{gt}}$ are exactly the strategies undominated wrt $q_0, \bigcirc\mathsf{win}$. So, $M_2, q_i \models_{\text{ETSL}} \langle\langle\mathbb{A}\text{gt}\rangle\rangle\bigcirc\mathsf{win}$ for every $i = 0, 1, ..., 4$, and therefore $M_2, q_0 \models_{\text{ETSL}} E_{\mathbb{A}\text{gt}}\langle\langle\mathbb{A}\text{gt}\rangle\rangle\bigcirc\mathsf{win}$. On the other hand, there is no single strategy that succeeds for all $q_0, q_1, ..., q_4$.

Common knowledge: consider model M_3 from Figure 2B. Let $S_{\{a,b\}}$ be the strategy "play $\langle 1, 1\rangle$ everywhere", and $T_{\{a,b\}}$ be "play $\langle 2, 2\rangle$ everywhere". Note that $S_{\{a,b\}}$ is the only undominated strategy wrt $q, \bigcirc\mathsf{win}$ for $q = q_0, q_1$, and $T_{\{a,b\}}$ is the only undominated strategy wrt $q, \bigcirc\mathsf{win}$ for $q = q_2, q_3$. Thus, for every $q = q_0, ..., q_3$: $M_3, q \models_{\text{ETSL}} \langle\langle a, b\rangle\rangle\bigcirc\mathsf{win}$, and hence $M_3, q_1 \models_{\text{ETSL}} C_{\{a,b\}}\langle\langle a, b\rangle\rangle\bigcirc\mathsf{win}$. On the other hand, $M_3, q_1 \not\models_{\text{CSL}} \mathbb{C}_{\{a,b\}}\langle\langle a, b\rangle\rangle\bigcirc\mathsf{win}$.

5 Conclusions

In this paper, the relationship between rational play and knowing how to play is investigated in a formal way. To this end, we dust off Epistemic Temporal Strategic Logic by van Otterloo and Jonker [13], and propose a simpler semantics expressed entirely in terms of concurrent epistemic game structures and their states; we prove that the new semantics is equivalent to the original one for "vanilla" ETSL formulae. ETSL serves as a device for talking about the outcome of rational play (in the sense that agents are assumed to play only undominated strategies). To capture properties of the other kind ("knowing how to play"), we use the recent proposal of Constructive Strategic Logic [8, 9].

The main result of this paper states that, for finite models, *a rational player knows that he will succeed if, and only if, he knows how to succeed.* We also show that the relationship is much more limited for rational coalitions. That is, if rational agents have common knowledge about a winning strategy, then they have common knowledge that they will succeed – but the converse is not guaranteed any more. Moreover, it turns out that the relationship is *strictly reverse* for distributed knowledge: if a rational coalition has distributed knowledge that it will succeed, then it has distributed knowledge about a winning strategy – but not necessarily the other way around. Finally, for mutual knowledge, the relationship does not hold either way in general. This is a curious result, and one that may lead to interesting philosophical conclusions.

References

1. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. In *Proceedings of FOCS*, pages 100–109. IEEE Computer Society Press, 1997.
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
3. V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004.
4. A. Herzig and N. Troquard. Knowing how to play: Uniform choices in logics of agency. In *Proceedings of AAMAS'06*, 2006. To appear.
5. W. Jamroga. Some remarks on alternating temporal epistemic logic. In *Proceedings of FAMAS 2003*, pages 133–140, 2003.
6. W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
7. W. Jamroga, W. van der Hoek, and M. Wooldridge. Intentions and strategies in game-like scenarios. In *Proceedings of EPIA 2005*, volume 3808 of *Lecture Notes in Artificial Intelligence*, pages 512–523. Springer Verlag, 2005.
8. W. Jamroga and Thomas Ågotnes. Constructive knowledge: What agents can achieve under incomplete information. Technical Report IfI-05-10, Clausthal University of Technology, 2005.
9. Wojciech Jamroga and Thomas Ågotnes. What agents can achieve under incomplete information. In *Proceedings of AAMAS'06*, 2006.
10. G. Jonker. Feasible strategies in Alternating-time Temporal Epistemic Logic. Master thesis, University of Utrecht, 2003.
11. P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2), 2004.
12. W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of AAMAS-02*, pages 1167–1174, 2002.
13. S. van Otterloo and G. Jonker. On Epistemic Temporal Strategic Logic. *ENTCS*, 126:77–92, 2005. Proceedings of LCMAS'04.