Logics for Reasoning About Strategic Abilities in Multi-player Games

Nils Bulling[1(✉)], Valentin Goranko[2,3], and Wojciech Jamroga[4,5]

[1] Department of Intelligent Systems, TU Delft, Delft, Netherlands
n.bulling@tudelft.nl
[2] Department of Philosophy, Stockholm University, Stockholm, Sweden
valentin.goranko@philosophy.su.se
[3] Department of Mathematics, University of Johannesburg, Johannesburg, South Africa
[4] Institute of Computer Science, Polish Academy of Sciences, Warszawa, Poland
[5] Computer Science and Communication Research Unit and Interdisc. Centre on Security, Reliability and Trust, University of Luxembourg, Luxembourg City, Luxembourg
w.jamroga@ipipan.waw.pl

Abstract. We introduce and discuss basic concepts, ideas, and logical formalisms used for reasoning about strategic abilities in multi-player games. In particular, we present concurrent game models and the alternating time temporal logic ATL^* and its fragment ATL. We discuss variations of the language and semantics of ATL^* that take into account the limitations and complications arising from incomplete information, perfect or imperfect memory of players, reasoning within dynamically changing strategy contexts, or using stronger, constructive concepts of strategy. Finally, we briefly summarize some technical results regarding decision problems for some variants of ATL.

Keywords: Logics · Game theory · Strategic reasoning · Strategic logic · Multi-agent systems · Automated reasoning

1 Introduction: Strategic Reasoning

Strategic reasoning is ubiquitous in the modern world. Our entire lives comprise a complex flux of diverse yet interleaved games that we play in different social contexts with different sets of other players, different rules, objectives and preferences. The outcomes of these games determine not only our sense of success (winning) or failure (losing) in life but also what games we engage to play further, and how. In this process we adopt, consciously or not, and follow, commit, abandon, modify and re-commit again to a stream of local strategies. Thus, we are gradually composing and building a big strategy which, together with all

J. van Benthem et al. (Eds.): Models of Strategic Reasoning, LNCS 8972, pp. 93–136, 2015.
DOI: 10.1007/978-3-662-48540-8_4

that happens in the surrounding environment and the 'butterfly effects' coming from the rest of the world, determines a unique play called Life . . .[1]

After this lyrical-philosophical overture, let us make some more analytic introductory notes on our view of strategic reasoning.

To begin with, one can distinguish two, related yet different, perspectives on strategic reasoning depending on the position of the reasoner[2]:

- Reasoning of the agents (players) *from within the game* on what strategy to adopt in order to best achieve their objectives. This starts with 'zero-order' reasoning from the player's own perspective, but not taking into account the other players' strategic reasoning. Then it evolves into 'first-order' reasoning by only taking into account the other players' zero-order strategic reasoning; then likewise second-, third-, etc. higher-order strategic reasoning, eventually converging to the concept of 'common belief/knowledge of rationality', fundamental in game theory.
- Reasoning of an external observer, *from outside the game*, on what strategies the playing agents can *objectively* adopt in trying to achieve their objectives. This reasoning can again be stratified into conceptual layers by taking into account the players' observational, informational, memory, and reasoning limitations in the game, but also their knowledge or ignorance about the other players' limitations, their objectives, etc. Eventually, a complex hierarchy of levels of 'objective' or 'external' strategic reasoning emerges, that essentially embeds the 'internal' one above.

One can also distinguish different threads of strategic reasoning depending on the rationality assumptions, both for the proponents and the opponents. As we noted, the game-theoretic tradition emphasizes reasoning about *rational players'* strategic behaviour *under the assumption of common belief or knowledge of rationality.* Depending on how this assumption is perceived various solution concepts emerge, describing or prescribing the players' rational strategic behaviour. For the epistemic and doxastic foundations of strategic behaviour of rational agents, focusing on the internal perspective of strategic reasoning, we refer to other chapters in this volume: Bonanno [18], Perea [74] and Pacuit [68]. Another active and promising direction of current research on strategic reasoning, presented in this chapter, does not consider players taking into account any assumptions about the rationality of the other players but analyzes, from an external observer's perspective, the players' *objective abilities* to adopt and to apply strategies that guarantee the achievement of their goals *regardless of the rationality level and strategic behaviour of the opponents.* Thus, when assessing objectively the strategic abilities of individual players or coalitions of players – generically called the 'proponents' – to achieve a specific goal we essentially assume that the remaining

[1] While 'strategy' is commonly defined as a complete conditional plan, we cannot resist noting here John Lennon's famous quote: "Life is what happens while you are busy making other plans".

[2] Roughly corresponding to 'first-person deliberation' vs. 'third-person assessment of strategic action in games' in van Benthem [16].

players – the 'opponents' – play a collective adversary in a strictly competitive game between the proponents and the opponents.

One can also regard the framework presented in this chapter as analyzing the objective strategic abilities of players – possibly impaired by imperfect or incomplete knowledge about the game – to achieve qualitative goals using zero-order reasoning only in concurrent extended multi-player games[3].

2 Concurrent Game Models and Strategic Abilities

Logics of strategic reasoning build upon several fundamental concepts from game theory, the most important being that of a 'strategy'. The notion of strategy adopted in this chapter is classical: a conditional plan that prescribes what action a given agent (or, a coalition of agents) should take in every possible situation that may arise in the system (game) in which they act. This notion will be made mathematically more precise in this chapter, where strategies will be used to provide formal logical semantics.

We start with a technical overview of the basic game-theoretic concepts used later on in this chapter. For more details we refer the reader to e.g. [51,65].

Throughout this chapter we use the terms 'agent' and 'player' as synonyms and consider an arbitrarily fixed nonempty finite[4] set of players/agents $\mathbb{A}\mathrm{gt}$. We also fix a nonempty set of atomic propositions $Prop$ that encode basic properties of game states.

2.1 One-Round Multi-player Strategic Games

The abstract games studied in traditional non-cooperative game theory are usually presented either in *extensive* or in *strategic* form (also known as *normal form*). We first focus on the latter type of games here.

Strategic Games

Definition 1 (Strategic Game Forms and Strategic Games). *A* strategic game form *is a tuple* $(\mathbb{A}\mathrm{gt}, \{\mathsf{Act_a} \mid \mathsf{a} \in \mathbb{A}\mathrm{gt}\}, \mathsf{Out}, \mathsf{out})$ *that consists of a nonempty finite set of players* $\mathbb{A}\mathrm{gt}$*, a nonempty set of* actions *(also known as* moves *or* choices*)* $\mathsf{Act_a}$ *for each player* $\mathsf{a} \in \mathbb{A}\mathrm{gt}$*, a nonempty set of* outcomes Out*, and an* outcome function $\mathsf{out} : \prod_{\mathsf{a}\in\mathbb{A}\mathrm{gt}} \mathsf{Act_a} \to \mathsf{Out}$*, that associates an outcome with every* action profile*; that is, tuple of actions, one for each player*[5]*.*

A strategic game *is a strategic game form endowed with* preference orders $\leq_{\mathsf{a}}$ *on the set of outcomes, one for each player. Often, players' preferences are*

[3] We do, however, discuss briefly in Sect. 5.2 how some concepts of rationality can be expressed in logical languages considered here.

[4] We have no strong reason for this finiteness assumption, other than common sense and technical convenience.

[5] We assume that there is an ordering on $\mathbb{A}\mathrm{gt}$ which is respected in the definition of tuples etc.

expressed by payoff functions $u_{\mathsf{a}} : \mathsf{Out} \to \mathbb{R}$. *Then, the preference relations are implicitly defined as follows:* $o \leq_{\mathsf{a}} o'$ *iff* $u_{\mathsf{a}}(o) \leq u_{\mathsf{a}}(o')$. *Thus, strategic games can be represented either as tuples* $(\mathbb{A}\mathrm{gt}, \{\mathsf{Act}_{\mathsf{a}} \mid \mathsf{a} \in \mathbb{A}\mathrm{gt}\}, \mathsf{Out}, \mathsf{out}, (\leq_{\mathsf{a}})_{\mathsf{a}\in\mathbb{A}\mathrm{gt}})$ *or* $(\mathbb{A}\mathrm{gt}, \{\mathsf{Act}_{\mathsf{a}} \mid \mathsf{a} \in \mathbb{A}\mathrm{gt}\}, \mathsf{Out}, \mathsf{out}, (u_{\mathsf{a}})_{\mathsf{a}\in\mathbb{A}\mathrm{gt}})$.

In traditional game theory outcomes are usually characterized quantitatively by real values called *utilities* or *payoffs*. More generally, outcomes can be abstract objects, ordered by relations $\leq_{\mathsf{a}}$ which represent preferences of players, as in the definition above. Here we abstract from the actual preferences between outcomes and focus on the players' powers to enforce particular properties (sets) of outcome states. Thus, we will use the terms "strategic game" and "strategic game form" interchangeably, assuming that game forms come equipped with some preference orders that have no direct bearing on our discussion.

The intuition behind a strategic game is simple: each player chooses an action from her set of possible actions. All actions are performed *independently and simultaneously*. Thus, all players perform a collective action based on which the outcome function determines the (unique) outcome. Hence, a strategic game typically represents a one-shot interaction.

1\2	*coop*	*defect*
coop	(3, 3)	(0, 5)
defect	(5, 0)	(1, 1)

Fig. 1. Prisoner's Dilemma.

Example 1 (Prisoner's Dilemma as a Strategic Game). We will use a version of the well-known Prisoner's Dilemma game, given in Fig. 1, to illustrate the basic concepts introduced in this section. Each of the two players in the game can choose to cooperate (play action *coop*) or to defect (play action *defect*). Formally, the game is defined as $(\{1, 2\}, \{\mathsf{Act}_1, \mathsf{Act}_2\}, \{o_1, o_2, o_3, o_4\}, \mathsf{out}, (\leq_1, \leq_2))$ with $\mathsf{Act}_1 = \mathsf{Act}_2 = \{\mathit{coop}, \mathit{defect}\}$, $\mathsf{out}(\mathit{coop}, \mathit{coop}) = o_1$, $\mathsf{out}(\mathit{coop}, \mathit{defect}) = o_2$, $\mathsf{out}(\mathit{defect}, \mathit{coop}) = o_3$, and $\mathsf{out}(\mathit{defect}, \mathit{defect}) = o_4$. Moreover, we define $\leq_1$ and $\leq_2$ as the smallest transitive relations with $o_2 \leq_1 o_4 \leq_1 o_1 \leq_1 o_3$ and $o_3 \leq_2 o_4 \leq_2 o_1 \leq_2 o_2$. In the figure we have shown the value of the payoff functions u_1 and u_2 defined as follows: $u_1(o_1) = u_2(o_1) = 3$, $u_1(o_4) = u_2(o_4) = 1$, $u_1(o_2) = u_2(o_3) = 0$, and $u_1(o_3) = u_2(o_2) = 5$.

2.2 Effectivity Functions and Models for Strategic Games

It is important to note that in strategic games none of the players knows in advance the actions chosen by the other players, and therefore has no definitive control on the outcome of the game. So, what power does an individual player or a coalition of players have to influence the outcome in such a game? We will address this fundamental question below in terms of *effectivity functions*, first introduced in cooperative[6] game theory in Moulin and Peleg [64] and in social choice theory in Abdou and Keiding [1], to provide an abstract representation of powers of players and coalitions.

Definition 2 (Effectivity Functions and Models). *Given a set of players* $\mathbb{A}\mathrm{gt}$ *and a set of outcomes* Out*, a* (coalitional) effectivity function (EF) *over* $\mathbb{A}\mathrm{gt}$ *and* Out *is a mapping* $\mathsf{E} : \mathcal{P}(\mathbb{A}\mathrm{gt}) \to \mathcal{P}(\mathcal{P}(\mathsf{Out}))$ *that associates a family of sets of outcomes with each coalition of players.*

A (coalitional) effectivity model *(EM) is a coalitional effectivity function endowed with a labelling* $V : \mathsf{Out} \to \mathcal{P}(Prop)$ *of outcomes with sets of atomic propositions from Prop. The labeling prescribes which atomic propositions are true in a given outcome state.*

Intuitively, for a group of agents $A \subseteq \mathbb{A}\mathrm{gt}$ every element of $\mathsf{E}(A)$ is the set of all possible outcomes that can result from a given joint action of players in A, depending on how the remaining players from $\mathbb{A}\mathrm{gt}$ decide to act. In other words, for every set X in $\mathsf{E}(A)$ the coalition A has a collective action that is guaranteed to yield an outcome in X, regardless of the actions taken by the players in $\overline{A} = \mathbb{A}\mathrm{gt} \setminus A$. Therefore, every element of $\mathsf{E}(A)$ can be regarded as representing a possible joint action of coalition A.

Every strategic game G naturally defines an effectivity function called the α-effectivity function of G and denoted by E^{α}_{G}, which is defined as follows.

Definition 3 (Effectivity in Strategic Games, Pauly [71]). *For a strategic game G, the* α-effectivity function $\mathsf{E}^{\alpha}_{G} : \mathcal{P}(\mathbb{A}\mathrm{gt}) \to \mathcal{P}(\mathcal{P}(\mathsf{Out}))$ *is defined as follows:* $X \in \mathsf{E}^{\alpha}_{G}(A)$ *if and only if there exists a joint action* σ_A *for* A *such that for every joint action* $\sigma_{\overline{A}}$ *of* $\overline{A}$ *we have* $\mathsf{out}(\sigma_A, \sigma_{\overline{A}}) \in X$.

Respectively, the β-effectivity function for G *is* $\mathsf{E}^{\beta}_{G} : \mathcal{P}(\mathbb{A}\mathrm{gt}) \to \mathcal{P}(\mathcal{P}(\mathsf{Out}))$*, defined as follows:* $X \in \mathsf{E}^{\beta}_{G}(A)$ *if and only if for every joint action* $\sigma_{\overline{A}}$ *of* $\overline{A}$ *there exists a joint action* σ_A *of* A *(generally, depending on* $\sigma_{\overline{A}}$*) such that* $\mathsf{out}(\sigma_A, \sigma_{\overline{A}}) \in X$.

Intuitively, α-effectivity functions describe the powers of coalitions to *guarantee* outcomes satisfying desired properties while β-effectivity functions describe the abilities of coalitions to *prevent* outcomes satisfying undesired properties.

Since strategic games are determined, α-effectivity and β-effectivity of coalitions are dual to each other in a sense that for every coalition A and $X \subseteq \mathsf{Out}$:

$$X \in \mathsf{E}^{\alpha}_{G}(A) \text{ iff } \overline{X} \notin \mathsf{E}^{\beta}_{G}(\overline{A})$$

where $\overline{X} = \mathsf{Out} \setminus X$. That is, a coalition A can guarantee an outcome with a property X precisely when its complementary coalition $\overline{A}$ cannot prevent it.

[6] Coalitional effectivity can be regarded as a concept of cooperative game theory from the internal perspective of the coalition, but from the external perspective of the other players it becomes a concept of non-cooperative game theory. We will not dwell into this apparent duality here.

Example 2 (Prisoner's Dilemma as Effectivity Model). The Prisoner's Dilemma from Example 1 can also be represented by the following effectivity function over $(\{1,2\},\{o_1,\ldots,o_4\})$: $E(\emptyset)=\{\mathsf{Out}\}$, $E(\{1\})=\{\{o_1,o_2\},\{o_3,o_4\}\}\cup\{X\subseteq\{o_1,\ldots,o_4\}\mid\{o_1,o_2\}\subseteq X$ or $\{o_3,o_4\}\subseteq X\}$, $E(\{2\})=\{\{o_1,o_3\},\{o_2,o_4\}\}\cup\{X\subseteq\{o_1,\ldots,o_4\}\mid\{o_1,o_3\}\subseteq X$ or $\{o_2,o_4\}\subseteq X\}$, and $E(\{1,2\})=\{\{o_1\},\{o_2\},\{o_3\},\{o_4\}\}\cup\{\exists i\in\{1,2,3,4\}$ s.t. $o_i\in X\}=\mathcal{P}(\{o_1,o_2,o_3,o_4\})\backslash\{\emptyset\}$.

Let us adopt atomic propositions representing the payoff values for each agent $\{\mathsf{p}_\mathsf{a}^\mathsf{j}\mid\mathsf{a}\in\mathbb{A}\mathrm{gt},j\in\{0,1,3,5\}\}$ and label the outcomes appropriately. Then, for example, we have $V(o_1)=\{\mathsf{p}_1^3,\mathsf{p}_2^3\}$ and $V(o_2)=\{\mathsf{p}_1^0,\mathsf{u}_2^5\}$.

2.3 Characterization of Effectivity Functions of Strategic Games

Clearly, not every abstract effectivity function defined as above corresponds to strategic games. The following notion captures the properties required for such correspondence.

Definition 4 (True Playability (Pauly [71], Goranko et al. [45]). *An effectivity function* $\mathsf{E}:\mathcal{P}(\mathbb{A}\mathrm{gt})\to\mathcal{P}(\mathcal{P}(\mathsf{Out}))$ *is* truly playable *iff the following conditions hold:*

Outcome monotonicity: $X\in\mathsf{E}(A)$ *and* $X\subseteq Y$ *implies* $Y\in\mathsf{E}(A)$*;*
Liveness: $\emptyset\notin\mathsf{E}(A)$*;*
Safety: $\mathsf{St}\in\mathsf{E}(A)$*;*
Superadditivity: *if* $A_1\cap A_2=\emptyset$, $X\in\mathsf{E}(A_1)$ *and* $Y\in\mathsf{E}(A_2)$, *then* $X\cap Y\in\mathsf{E}(A_1\cup A_2)$*;*
$\mathbb{A}\mathrm{gt}$***-maximality:*** $\overline{X}\notin\mathsf{E}(\emptyset)$ *implies* $X\in\mathsf{E}(\mathbb{A}\mathrm{gt})$*;*
Determinacy: *if* $X\in\mathsf{E}(\mathbb{A}\mathrm{gt})$ *then* $\{x\}\in\mathsf{E}(\mathbb{A}\mathrm{gt})$ *for some* $x\in X$*.*

It is easy to see that every α-effectivity function of a strategic game is truly playable. The converse holds too as stated below.

Representation theorem for effectivity functions.[7] An effectivity function E for $(\mathbb{A}\mathrm{gt},\mathsf{Out})$ is truly playable if and only if there exists a strategic game $G=(\mathbb{A}\mathrm{gt},\{\mathsf{Act}_i\mid i\in\mathbb{A}\mathrm{gt}\},\mathsf{Out},\mathsf{out})$ such that $\mathsf{E}_G^\alpha=\mathsf{E}$, see [45,71].

Actual α-effectivity in Strategic Games. The notion of effectivity in game G can be refined to the "actual" α-effectivity function of G that collects *precisely* the sets of outcomes of collective actions available to the coalition without closing the sets under outcome monotonicity. Formally, given a strategic game $G=(\mathbb{A}\mathrm{gt},\{\mathsf{Act}_i\mid i\in\mathbb{A}\mathrm{gt}\},\mathsf{Out},\mathsf{out})$, a coalition A and a joint action σ_A we define $\mathsf{outcome\_states}(\sigma_A)$ as the set of all possible outcomes that can result from σ_A:

$$\mathsf{outcome\_states}(\sigma_A)=\{\mathsf{out}(\sigma_A,\sigma_{\overline{A}})\mid\sigma_{\overline{A}}\text{ is a joint action for }\overline{A}\}.$$

[7] This representation theorem was first proved in Pauly [71] for so called "playable" effectivity functions, without the *Determinacy* requirement. It has been recently shown in [45] that, for games with infinite outcome spaces, "playability" is not sufficient. The *Determinacy* condition was identified and added to define "truly playable" effectivity functions and prove a correct version of the representation theorem.

We define the *actual* α*-effectivity function* $\widehat{\mathsf{E}}_G : \mathcal{P}(\mathbb{A}\text{gt}) \to \mathcal{P}(\mathcal{P}(\mathsf{Out}))$ as the family of all outcome sets effected by possible joint actions of A:

$$\widehat{\mathsf{E}}_G(A) = \{\mathsf{outcome\_states}(\sigma_A) \mid \sigma_A \text{ is a joint action for } A\}.$$

Clearly, the standard α-effectivity function for G can now be obtained by closure under outcome-monotonicity:

$$\mathsf{E}_G(A) = \{Y \mid X \subseteq Y \text{ for some } X \in \widehat{\mathsf{E}}_G(A)\}.$$

Conversely, obtaining $\widehat{\mathsf{E}}_G$ from E_G for games with infinite outcome states is not so straightforward because $\widehat{\mathsf{E}}_G$ may not be uniquely determined by E_G, so the notion of actual effectivity is at least as interesting and perhaps more important than "standard", outcome-monotone effectivity. We refer the interested reader for further discussion and details to [44].

2.4 Strategic Abilities in Concurrent Game Models

Extensive game forms allow to model turn-based games, where at every non-terminal position only one player is allowed to make a move. In this section we discuss more general (as we explain further) "concurrent" games, where at every position all players make their moves simultaneously.

Extensive Games meet Repeated Games: Concurrent Game Structures. Strategic games are usually interpreted as one-step games. Especially in evolutionary game theory, they are often considered in a repeated setting: game G is played a number of times, and the payoffs from all rounds are aggregated. Concurrent game structures from [8], which are essentially equivalent to multi-player game frames from [71] (see Goranko [42]), generalize the setting of repeated games by allowing *different* strategic games to be played at different stages. This way we obtain multi-step games that are defined on some state space, in which every state is associated with a strategic game with outcomes being states again. The resulting game consists of successive rounds of playing one-step strategic games where the outcome of every round determines the successor state, and therefore the strategic game to be played at the next round. Alternatively, one can see concurrent game structures as a generalization of extensive game forms where simultaneous moves of different players are allowed, as well as loops to previously visited states.

Definition 5 (Concurrent Game Structures and Models). *A* concurrent game structure *(CGS) is a tuple*

$$\mathcal{S} = (\mathbb{A}\text{gt}, \mathsf{St}, \mathsf{Act}, \mathsf{act}, \mathsf{out})$$

which consists of a non-empty finite set of players $\mathbb{A}\text{gt} = \{1, \ldots, \mathsf{k}\}$*, a non-empty set of states*[8] St*, a non-empty set of atomic actions* Act*, a function* act :

[8] The set of states is assumed finite in [8] but that restriction is not necessary for our purposes. In Sect. 6.3 we even rely on the fact that the set of states can be infinite.

$\mathbb{A}\text{gt} \times \mathsf{St} \to \mathcal{P}(\mathsf{Act}) \setminus \{\emptyset\}$ *that defines the set of actions available to each player at each state, and a (deterministic) transition function* out *that assigns a unique* successor (outcome) state $\mathsf{out}(q, \alpha_1, \dots, \alpha_k)$ *to each state* q *and each tuple of actions* $\langle \alpha_1, \dots, \alpha_k \rangle$ *such that* $\alpha_a \in \mathsf{act}(a, q)$ *for each* $a \in \mathbb{A}\text{gt}$ *(i.e., each* α_a *that can be executed by player* a *in state* q*).*

A concurrent game model *(CGM) over a set of atomic propositions Prop is a CGS endowed with a labelling* $V : \mathsf{St} \to \mathcal{P}(Prop)$ *of game states with subsets of Prop, thus prescribing which atomic propositions are true at a given state.*

Thus, all players in a CGS execute their actions synchronously and the combination of these actions together with the current state determines the transition to a successor state in the CGS.

Note that turn-based extensive form games can be readily represented as concurrent game structures by assigning at each non-leaf state the respective set of actions to the player whose turn it is to move from that state, while allowing a single action 'pass' to all other players at that state. At leaf states all players are only allowed to 'pass' and the result of such collective pass action is the same state, thus looping there forever.

Example 3 (Prisoner's Escape). A CGM $\mathcal{M}_{esc}$ is shown in Fig. 2 modeling the following scenario. A prison has two exits: the rear exit guarded by the guard Alex and the front exit guarded by the guard Bob. The prison is using the following procedure for exiting (e.g., for the personnel): every person authorized to exit the prison is given secret passwords, one for every guard. When exiting

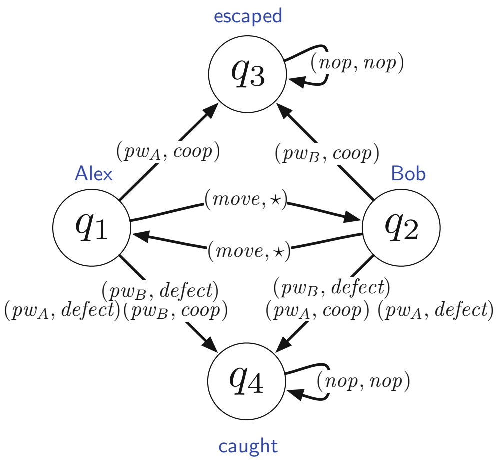

Fig. 2. Prisoner's escape modelled as CGM $\mathcal{M}_{esc}$. An action tuple (a_1, a_2) consists of an action of Frank (a_1) and Charlie (a_2). $\star$ is a placeholder for any action available at the very state; e.g., the tuple $(move, \star)$ leading from state q_1 to q_2 is a shortcut for the tuples $(move, defect)$ and $(move, coop)$. Loops are added to the "final states" q_3 and q_4 where action *nop* is the only available action for both players. We leave the formal definition to the reader.

the prison the guard must be given the password associated with him/her. If a person gives a wrong password to any guard, he is caught and arrested. Now, Frank is a prisoner who wants to escape, of course. Somehow Frank has got a key for his cell and has learned the passwords for each of the guards. Charlie is an internal guard in the prison and can always see when Frank is going to any of the exits. Frank has bribed Charlie to keep quiet and not to warn the other guards. Charlie can cooperate (actions *coop*), by keeping quiet, or can defect (action *defect*), by alerting the guards. Thus, the successful escape of Frank depends on Charlie's cooperation.

Global Coalition Effectivity Functions and Models. Every CGS $\mathcal{S}$ can be associated with a *global effectivity function* $\mathsf{E} : \mathsf{St} \times \mathcal{P}(\mathbb{A}\text{gt}) \to \mathcal{P}(\mathcal{P}(\mathsf{St}))$ that assigns a (local) α-effectivity function $\mathsf{E}_q = \mathsf{E}(q, \cdot)$ to every state $q \in \mathsf{St}$, generated by the strategic game associated with q in $\mathcal{S}$. These can be accordingly extended to *global effectivity models* by adding valuation of the atomic propositions.

Global effectivity functions and models have been introduced abstractly in [71–73] (called there 'effectivity frames and models'). The global effectivity functions generated by concurrent game structures are characterized in [45,71] by the true playability conditions listed in Sect. 2.3, applied to every E_q.

The idea of effectivity functions has also been extended to *path effectivity functions* in [44]. They will not be discussed here; the reader is referred to that paper for more details.

2.5 Strategies and Strategic Ability

Strategies in Concurrent Game Models. A *path* in a CGS/CGM is an infinite sequence of states that can result from subsequent transitions in the structure/model. A *strategy* of a player a in a CGS/CGM $\mathcal{M}$ is a conditional plan that specifies what a should do in each possible situation. Depending on the type of memory that we assume for the players, a strategy can be *memoryless* (alias *positional*), formally represented with a function $\mathbf{s_a} : \mathsf{St} \to \mathsf{Act}$, such that $\mathbf{s_a}(q) \in \mathsf{act_a}(q)$, or *memory-based* (alias *perfect recall*), represented by a function $\mathbf{s_a} : \mathsf{St}^+ \to \mathsf{Act}$ such that $\mathbf{s_a}(\langle \ldots, q\rangle) \in \mathsf{act_a}(q)$, where St^+ is the set of *histories*, i.e., finite sequences of states in $\mathcal{M}$. The latter corresponds to players with perfect recall of the past states; the former corresponds to players whose memory, if any, is entirely encoded in the current state of the system. Intermediate options, where agents have bounded memory, have been studied by Ågothes and Walther [5], but will not be discussed here.

A *joint strategy* of a group of players $A = \{\mathsf{a}_1, ..., \mathsf{a}_r\}$ is simply a tuple of strategies $\mathbf{s}_A = \langle \mathbf{s}_{\mathsf{a}_1}, ..., \mathbf{s}_{\mathsf{a}_r}\rangle$, one for each player from A. We denote player a's component of the joint strategy $\mathbf{s}_A$ by $\mathbf{s}_A[\mathsf{a}]$. Then, in the case of positional joint strategy $\mathbf{s}_A$, the action that $\mathbf{s}_A[\mathsf{a}]$ prescribes to player a at state q is $\mathbf{s}_A[\mathsf{a}](q)$; respectively, $\mathbf{s}_A[\mathsf{a}](\pi)$ is the action that a memory-based joint strategy $\mathbf{s}_A$ prescribes to a from the finite path (i.e., history) π. By a slight abuse of notation, we will use $\mathbf{s}_A(q)$ and $\mathbf{s}_A(\pi)$ to denote the joint actions of A in state q and history π, respectively.

Outcomes of Strategies and Strategic Abilities. The outcome set function $\mathsf{outcome\_states}$ can be naturally extended from joint actions to all strategy profiles applied at a given state (respectively, history) in a given CGS (or CGM). Then $\mathsf{outcome\_states}(q, \mathbf{s}_A)$ (respectively, $\mathsf{outcome\_states}(\pi, \mathbf{s}_A)$) returns the set of all possible successor states that can result from applying a given positional (respectively, memory-based) joint strategy $\mathbf{s}_A$ of the coalition A at state q (respectively, at history π). Formally,

$$\mathsf{outcome\_states}(q, \mathbf{s}_A) = \{\mathsf{out}(q, \mathbf{s}_A(q), \mathbf{s}_{\overline{A}}(q)) \mid \mathbf{s}_{\overline{A}} \text{ is a joint strategy of } \overline{A}\}.$$

The local actual effectivity function $\widehat{\mathsf{E}}_{\mathcal{S}}$, defining the coalitional powers at every state q in $\mathcal{S}$, is defined explicitly as

$$\widehat{\mathsf{E}}_{\mathcal{S}}(q, A) = \{\mathsf{outcome\_states}(q, \mathbf{s}_A) \mid \mathbf{s}_A \text{ is a memoryless joint strategy of } A\}.$$

As before, the standard α-effectivity functions for $\mathcal{S}$ can be obtained by closure under outcome-monotonicity:

$$\mathsf{E}_{\mathcal{S}}(q, A) = \{Y \mid X \subseteq Y \text{ for some } X \in \widehat{\mathsf{E}}_{\mathcal{S}}(q, A)\}.$$

Likewise for $\mathsf{outcome\_states}(\pi, \mathbf{s}_A)$, $\widehat{\mathsf{E}}_{\mathcal{S}}(\pi, A)$, and $\mathsf{E}_{\mathcal{S}}(\pi, A)$ which we will not further discuss here.

Example 4 (Prisoner's Escape Continued). $\widehat{\mathsf{E}}_{\mathcal{S}}(q_1, \{\text{Frank}\}) = \{\{q_2\}, \{q_4\}, \{q_3, q_4\}\}$, where $\mathcal{S}$ denotes the underlying CGS of $\mathcal{M}_{esc}$.

We extend the function $\mathsf{outcome\_states}$ to a function $\mathsf{outcome\_plays}$ that returns the set of all *plays*, i.e., all paths $\lambda \in \mathsf{St}^\omega$ that can be realised when the players in A follow strategy $\mathbf{s}_A$ from a given state q (respectively, history π) onward. Formally, for memoryless strategies this is defined as:

$$\mathsf{outcome\_plays}(q, \mathbf{s}_A) = \{\lambda = q_0, q_1, q_2 \ldots \mid q_0 = q \text{ and for each } j \in \mathbb{N} \text{ there exists an action profile for all players } \langle \alpha_1^j, \ldots, \alpha_{\mathsf{k}}^j \rangle \text{ such that } \alpha_{\mathsf{a}}^j \in \mathsf{act}_{\mathsf{a}}(q_j) \text{ for every } \mathsf{a} \in \mathbb{A}\mathrm{gt},\ \alpha_{\mathsf{a}}^j = \mathbf{s}_A[\mathsf{a}](q_j) \text{ for every } \mathsf{a} \in A, \text{ and } q_{j+1} = \mathsf{out}(q_j, \alpha_{\mathsf{a}_1}^j, \ldots, \alpha_{\mathsf{a}_k}^j)\}.$$

The definition for memory-based strategies is analogous: $\mathsf{outcome\_plays}(q, \mathbf{s}_A)$ consists of all plays of the game that start in q and can be realised as a result of each player in A following its individual memory-based strategy in $\mathbf{s}_A$, while the remaining players act in any way that is admissible by the game structure.

Example 5 (Prisoner's Escape Continued). Suppose the guard Charlie, who is a friend with the guard Bob and does not want to cause him trouble, adopts the *memoryless strategy* to cooperate with Frank if he goes to the rear exit (i.e., at state q_1) by not warning Alex, but to defect and warn Bob if Frank decides to go to the front exit, i.e. at state q_2. Naturally, Frank does not know that. The set of possible outcome plays enabled by this strategy and starting from state q_1 is:

$$\{(q_1 q_2)^\omega, q_1 (q_2 q_1)^n q_3^\omega, (q_1 q_2)^n q_4^\omega \mid n \in \mathbb{N}\}.$$

Suppose now being at q_1 Frank decides, for his own reasons, to try to escape through the front exit. Frank's strategy is to move at q_1 and to give the password to the guard at q_2. The resulting play from that strategy profile is the play $q_1 q_2 q_4^\omega$.

Memory-based strategies are more flexible. For instance, a memory-based strategy for Charlie could be one where he defects *the first time* Frank appears at any given exit, but thereafter cooperates at the rear exit (say, because he was then given more money by Frank) and defects at the front exit. That strategy enables the following set of plays from q_1:

$$\{(q_1 q_2)^\omega, q_1 q_4^\omega, (q_1 q_2)^n q_4^\omega, q_1 (q_2 q_1)^{n+1} q_3^\omega \mid n \in \mathbb{N}\}.$$

So, if Frank tries to escape as soon as possible with Charlie's support, he will fail; however, if he decides to first move to the front exit (and to pay Charlie extra money) and tries to escape the second time he appears at the front exit, he may succeed.

Note that there is no *memoryless strategy* that would allow Frank to escape if Charlie adopts the strategy specified above. This is because Frank's memoryless strategy must specify *the same action* in each state every time he is at that state, regardless of the history of the game up to that point. Thus, either Frank tries to escape through one of the exits right away, or he executes *move* forever, or gets caught.

A fundamental question regarding a concurrent game model is: *what can a given player or coalition achieve in that game*? So far the objectives of players and coalitions are not formally specified, but a typical objective would be to reach a state satisfying a given property, e.g. a winning state. Generally, an objective is a property of plays, for instance one can talk about winning or losing plays for the given player or coalition. More precisely, if the current state of the game is q we say that a coalition of players A can (is sure to) achieve an objective O from that state if there is a joint strategy $\mathbf{s}_A$ for A such that every play from $\mathsf{outcome\_plays}(q, \mathbf{s}_A)$ satisfies the objective O. The central problem that we discuss in the rest of this chapter is *how to use logic to formally specify strategic objectives of players and coalitions and how to formally determine their abilities to achieve such objectives.*

3 Logics for Strategic Reasoning and Coalitional Abilities

Logic and game theory have a long and rich history of interaction which we will not discuss here and refer the reader to e.g. [15]. Here, we will focus on the role of logic in formalizing and structuring reasoning about strategic abilities in multi-player games.

3.1 Expressing Local Coalitional Powers: Coalition Logic

The concept of α-effectivity in strategic games (Definition 3) has the distinct flavour of a *non-normal modal operator with neighbourhood semantics*, see [33], and this observation was utilized by Pauly who introduced in [71,73] a multi-modal logic capturing coalitional effectivity in strategic games, called *Coalition Logic (* CL *).* CL extends classical propositional logic with a family of modal operators $[A]$ parameterized with coalitions, i.e. subsets of the set of agents $\mathbb{A}\mathrm{gt}$. Intuitively, the formula $[A]\varphi$ says that *coalition A has, at the given game state, the power to guarantee an outcome satisfying* φ. Formally, operator $[A]$ is interpreted in global effectivity models $\mathcal{M} = (\mathsf{E}, V)$ as follows:

$$\mathcal{M}, q \models [A]\varphi \;\text{ iff }\; \|\varphi\|_{\mathcal{M}} \in \mathsf{E}_q(A),$$

where $\|\varphi\|_{\mathcal{M}} := \{s \in \mathsf{St} \mid \mathcal{M}, s \models \varphi\}$.

This implicitly defines the semantics of CL in every concurrent game model $\mathcal{M}$, in terms of the generated global α-effectivity function $\mathsf{E}_{\mathcal{M}}$.

Coalition logic is a very natural language to express *local* strategic abilities of players and coalitions; that is, their powers to guarantee desired properties in the *successor states*.

Example 6 In the following we state some properties expressed in CL.

1. *"If Player 1 has an action to guarantee a winning successor state, then Player 2 cannot prevent reaching a winning successor state."*

$$[1]\,\mathit{Win1} \rightarrow \neg[2]\neg\,\mathit{Win1}.$$

2. *"Player 1 has an action to guarantee a successor state where she is rich, and has an action to guarantee a successor state where she is happy, but has no action to guarantee a successor state where she is both rich and happy."*

$$[1]\mathit{Rich} \wedge [1]\mathit{Happy} \wedge \;\neg[1](\mathit{Rich} \wedge \mathit{Happy}).$$

3. *"None of players 1 and 2 has an action ensuring an outcome state satisfying Goal, but they have a collective action ensuring such an outcome state."*

$$\neg[1]\,\mathit{Goal} \wedge \neg[2]\,\mathit{Goal} \wedge [1,2]\,\mathit{Goal}.$$

Example 7 (Prisoner's Escape: Example 3 Continued). Let us denote hereafter Frank by f and Charlie by c. Then we have $\mathcal{M}_{esc}, q_1 \models \neg[f]\mathsf{escaped}$ and $\mathcal{M}_{esc}, q_1 \models [f,c]\mathsf{escaped}$.

3.2 Expressing Long-Term Strategic Abilities in the Logic ATL*

While CL is suitable for expressing local properties and immediate abilities, it cannot capture *long-term* strategic abilities of players and coalitions. For these, we need to extend the language of CL with more expressive temporal operators. That was done in [71,72] where Pauly introduced the *Extended Coalition Logic* ECL, interpreted essentially (up to notational difference) on concurrent game models. Independently, a more expressive logical system called *Alternating-Time Temporal Logic*, ATL^* (and its syntactic fragment ATL) was introduced and studied by Alur, Henzinger and Kupferman in a series of papers, see [6–8] as a logic for reasoning about open systems. The main syntactic construct of ATL^* is a formula of type $\langle\!\langle A\rangle\!\rangle\gamma$, intuitively meaning:

"The coalition A has a collective strategy to guarantee the satisfaction of the objective γ on every play enabled by that strategy"[9].

As shown in [42,43] Pauly's ECL is directly embeddable into ATL, so we will not discuss ECL further, but will focus on ATL and ATL* interpreted over concurrent game models.

ATL* and its Fragment ATL : Syntax and Semantics. Formally, the alternating-time temporal logic ATL* is a multimodal logic extending the linear time temporal logic LTL– comprising the temporal operators $\mathcal{X}$ ("at the next state"), $\mathcal{G}$ ("always from now on") and $\mathcal{U}$ ("until") – with *strategic path quantifiers* $\langle\langle A \rangle\rangle$ indexed with coalitions A of players. There are two types of formulae of ATL*: *state formulae* that constitute the logic, and which are evaluated at game states, and *path formulae*, which are evaluated on game plays. These are respectively defined by the following grammars, where $A \subseteq \mathbb{A}\mathrm{gt}, \mathsf{p} \in Prop$:

State formulae: $\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\gamma$,
Path formulae: $\gamma ::= \varphi \mid \neg\gamma \mid \gamma \wedge \gamma \mid \mathcal{X}\gamma \mid \mathcal{G}\gamma \mid \gamma\,\mathcal{U}\,\gamma$.

The formal semantics of ATL* was initially based on *alternating transition systems* in [6,7], and subsequently reworked for concurrent game models, as follows[10]. Let $\mathcal{M}$ be a CGM, q a state in $\mathcal{M}$, and $\lambda = q_0 q_1 \ldots$ be a path in $\mathcal{M}$. For every $i \in \mathbb{N}$ we define $\lambda[i] = q_i$, and denote by $\lambda[0..i]$ the prefix $q_0 q_1 \ldots q_i$, and by $\lambda[i..\infty]$ the suffix $q_i q_{i+1} \ldots$ of λ. The semantics of ATL* is given as follows (cf. [8]). For state formulae:

$\mathcal{M}, q \models \mathsf{p}$ iff $q \in V(\mathsf{p})$, for $\mathsf{p} \in Prop$;
$\mathcal{M}, q \models \neg\varphi$ iff $\mathcal{M}, q \not\models \varphi$;
$\mathcal{M}, q \models \varphi_1 \wedge \varphi_2$ iff $\mathcal{M}, q \models \varphi_1$ and $\mathcal{M}, q \models \varphi_2$;
$\mathcal{M}, q \models \langle\langle A \rangle\rangle\gamma$ iff there is a joint strategy $\mathbf{s}_A$ for A such that $\mathcal{M}, \lambda \models \gamma$ for every play $\lambda \in \mathsf{outcome\_plays}(q, \mathbf{s}_A)$;

and for path formulae:
$\mathcal{M}, \lambda \models \varphi$ iff $\mathcal{M}, \lambda[0] \models \varphi$ for any state formula φ;
$\mathcal{M}, \lambda \models \neg\gamma$ iff $\mathcal{M}, \lambda \not\models \gamma$;
$\mathcal{M}, \lambda \models \gamma_1 \wedge \gamma_2$ iff $\mathcal{M}, \lambda \models \gamma_1$ and $\mathcal{M}, \lambda \models \gamma_2$;
$\mathcal{M}, \lambda \models \mathcal{X}\gamma$ iff $\mathcal{M}, \lambda[1, \infty] \models \gamma$;
$\mathcal{M}, \lambda \models \mathcal{G}\gamma$ iff $\mathcal{M}, \lambda[i, \infty] \models \gamma$ for every $i \geq 0$; and
$\mathcal{M}, \lambda \models \gamma_1 \,\mathcal{U}\, \gamma_2$ iff there is i such that $\mathcal{M}, \lambda[i, \infty] \models \gamma_2$ and $\mathcal{M}, \lambda[j, \infty] \models \gamma_1$ for all $0 \leq j < i$.

The other Boolean connectives and constants $\top$ and $\bot$ are defined as usual. The operator $\mathcal{F}$ ("sometime in the future") is defined as $\mathcal{F}\varphi \equiv \top\,\mathcal{U}\,\varphi$.[11]

[9] We use the terms *objective* and *goal* of a coalition A as synonyms, to indicate the subformula γ of the formula $\langle\langle A \rangle\rangle\gamma$. In doing so, we ignore the issue of whether agents may have (common) goals, how these goals arise, etc.

[10] As proved in [42,43], under natural assumptions the two semantics are equivalent.

[11] Of course, $\mathcal{G}$ is definable as $\neg\mathcal{F}\neg$, but keeping it as a primitive operator in the language is convenient when defining the sublanguage ATL.

The logic ATL^* is very expressive, often more than necessary. This expressiveness comes at a high computational price which can be avoided if we settle for a reasonably designed fragment which is still sufficient in many cases. The key idea is to restrict *the combination of temporal operators* in the language. That can be achieved by imposing a syntactic restriction on the construction of formulae: occurrences of temporal operators must be immediately preceded by strategic path quantifiers. The result is the logic ATL defined by the following grammar, for $A \subseteq \mathbb{A}\text{gt}, \mathsf{p} \in \mathit{Prop}$:

$$\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A\rangle\rangle \mathcal{X}\varphi \mid \langle\langle A\rangle\rangle \mathcal{G}\varphi \mid \langle\langle A\rangle\rangle (\varphi \,\mathcal{U}\, \varphi).$$

For example, $\langle\langle A\rangle\rangle \mathcal{G}\,\mathcal{F}\,\mathsf{p}$ is an ATL^* formula but not an ATL formula whereas $\langle\langle A\rangle\rangle \mathcal{G}\,\langle\langle B\rangle\rangle \mathcal{F}\,\mathsf{p}$ is also an ATL formula. Thus, the coalitional objectives in ATL formulae are quite simple. As a consequence, it turns out that for the formulae of ATL the two notions of strategy, memoryless and memory-based, yield the same semantics [8,47].

Note that CL can be seen as the fragment of ATL involving only Booleans and operators $\langle\langle A\rangle\rangle \mathcal{X}$, whereas ECL also involves the operator $\langle\langle A\rangle\rangle \mathcal{G}$ (denoted in [71,72] by $[A^*]$). Both logics inherit the semantics of ATL on concurrent game models.

Example 8 (Prisoner's Escape Continued). We express some properties of the escape scenario from Example 3 in ATL^*. (We recall that we denote Frank by f and Charlie by c.) We remark that all but the last formula belong to ATL.

1. $\mathcal{M}_{esc}, q_1 \models \neg\langle\langle f\rangle\rangle \mathcal{F}\,\mathsf{escaped}$: Frank cannot guarantee to escape on his own (from q_1).
2. $\mathcal{M}_{esc}, q_1 \models \langle\langle f, c\rangle\rangle \mathcal{F}\,\mathsf{escaped}$: if Frank and Charlie cooperate then they can ensure that Frank eventually escapes.
3. $\mathcal{M}_{esc}, q_1 \models \langle\langle c\rangle\rangle \mathcal{G}\,\neg\mathsf{escape}$: Charlie can guarantee that Frank never escapes.
4. $\mathcal{M}_{esc}, q_1 \models \neg\langle\langle c\rangle\rangle \mathcal{F}\,\mathsf{caught}$: Charlie cannot guarantee that Frank is caught.
5. $\mathcal{M}_{esc}, q_1 \models \langle\langle f\rangle\rangle \mathcal{X}\,(\mathsf{Bob} \wedge \langle\langle f, c\rangle\rangle \mathcal{X}\,\mathsf{escaped})$: Frank has a strategy to reach the front exit guarded by Bob in the next step and then escape with the help of Charlie.
6. $\mathcal{M}_{esc}, q_1 \models \langle\langle f\rangle\rangle \mathcal{G}\,\mathcal{F}\,\mathsf{Alex}$: Frank can guarantee to reach the rear exit guarded by Alex infinitely many times.

3.3 From Branching-Time Temporal Logics to ATL^*

We have introduced ATL^* from a game-theoretic perspective as a logic for reasoning about players' strategic abilities. An alternative approach, in fact the one adopted by its inventors in [8], is to introduce $\mathsf{ATL}/\mathsf{ATL}^*$ as a generalization of the branching-time temporal logic $\mathsf{CTL}/\mathsf{CTL}^*$ to enable reasoning about open systems. Indeed, $\mathsf{CTL}/\mathsf{CTL}^*$ can be regarded as a 1-player version of $\mathsf{ATL}/\mathsf{ATL}^*$ where – assuming the singleton set of agents is $\{\mathsf{i}\}$ – the existential path quantifier E is identified with $\langle\langle \mathsf{i}\rangle\rangle$ and the universal path quantifier A is identified with $\langle\langle\emptyset\rangle\rangle$. Indeed, we leave it to the reader to check that the semantics of $\langle\langle \mathsf{i}\rangle\rangle\varphi$

and $\langle\langle\emptyset\rangle\rangle\varphi$ in any single-agent CGM $\mathcal{M}$ coincide with the semantics of $\mathsf{E}\,\varphi$ and $\mathsf{A}\,\varphi$ in $\mathcal{M}$ regarded as a transition system with transitions determined by the possible actions of the agent i, respectively.

ATL/ATL* can be regarded – at least formally, but see a discussion further – as a multi-agent extension of CTL/CTL* resulting into a more refined quantification scheme over the paths, respectively computations, enabled by some collective strategy of the given coalition.

4 Variations in Reasoning About Strategies

In this section, we discuss two interesting and important directions of extending the basic pattern of reasoning about agents' strategies and abilities. First, we investigate limitations and inadequacies stemming from the compositionality of the semantics of CL and ATL* that seem to be in conflict with the concept of strategy commitment. We discuss variant notions of strategic ability that attempt to resolve these problems. Then, we briefly summarize some attempts at reasoning about outcomes of *particular strategies*, rather than the mere existence of suitable plans.

4.1 Persistence of Strategic Committments

Strategic Commitment and Persistence in the Semantics of ATL*. Agents in actual multi-agent systems commit to strategies and relinquish their commitments in pursuit of their individual and collective goals in a dynamic, pragmatic, and often quite subtle way. While the semantics of ATL* is based on the standard notion of strategy, it appears that it does not capture adequately all aspects of strategic behaviour. For instance, the meaning of the ATL* formula $\langle\langle A\rangle\rangle\gamma$ is that the coalition A has a collective strategy, say s_A, to bring about the truth of γ *if the agents in A follow that strategy*. However, according to the formal semantics of ATL*, as introduced in [8], the evaluation of γ in the possible plays of the system enabled by s_A *does not take that strategy into account anymore*. That is, if γ contains a subformula $\langle\langle B\rangle\rangle\psi$, then in the evaluation of $\langle\langle B\rangle\rangle\psi$ the agents in $A \cap B$ are free to choose *any (other)* strategy as part of the collective strategy of B claimed to exist to justify the truth of ψ. Thus, the semantics of ATL* *does not commit the agents in A to the strategies they adopt in order to bring about the truth of the formula $\langle\langle A\rangle\rangle\gamma$*. This is in agreement with the semantics of path quantifiers in CTL*, where it is natural to express claims like $\mathsf{E}\mathcal{G}\,\mathsf{E}\varphi$ read as "there is a path, such that from any state of that path the system can deviate to another path which satisfies φ". One may argue that this feature disagrees with the game-theoretic view of a strategy as a full conditional plan that *completely specifies the agent's future behavior*. To see the problem more explicitly, consider the ATL formula $\langle\langle \mathsf{i}\rangle\rangle\mathcal{G}\,(\gamma \wedge \langle\langle \mathsf{i}\rangle\rangle\mathcal{X}\,\neg\gamma)$. Depending on how orthodoxly or liberally one adopts the concept of strategic commitment, the requirement expressed – that agent i has a strategy to ensure both that γ holds forever and that it can always alter that strategy to reach a non-γ state – may be considered satisfiable or not. This issue has been independently addressed in different ways in [3,4,19,75,81], where various proposals have been made in order to incorporate strategic commitment and persistent strategies in the syntax and semantics of ATL^*.

Paradoxes of Non-persistence. We continue with two more similar examples to argue that non-persistent strategies can lead to apparently counterintuitive descriptions of strategic ability.

Example 9 (Non-renewable Resource). Consider a system with a shared resource, where we are interested in reasoning about whether agent a has access to the resource. Let p denote the fact that agent a controls the resource. The ATL formula $\langle\langle a\rangle\rangle\mathcal{X}\, p$ expresses the claim that a is able to obtain control of the resource in the next moment, if it chooses to. Now imagine that agent a does not need to access the resource all the time, but it would like to be able to control the resource any time it needs it. Intuitively, this is expressed in ATL by the formula $\langle\langle a\rangle\rangle\mathcal{G}\,\langle\langle a\rangle\rangle\mathcal{X}\, p$, saying that a has a strategy which guarantees that, in any future state of the system, a can always force the next state to be one where a controls the resource.

Now, consider the single-agent system $\mathcal{M}_0$ from Fig. 3. We have that $\mathcal{M}_0, q_1 \models \langle\langle a\rangle\rangle\mathcal{X}\, p$: a can choose action α_2, which guarantees that p is true next. But we also have that $\mathcal{M}_0, q_1 \models \langle\langle a\rangle\rangle\mathcal{G}\,\langle\langle a\rangle\rangle\mathcal{X}\, p$: a's strategy in this case is to always choose α_1, which guarantees that the system will stay in q_1 forever and, as we have seen, $\mathcal{M}_0, q_1 \models \langle\langle a\rangle\rangle\mathcal{X}\, p$. However, this system does not have exactly the property we had in mind because by following that strategy, agent a dooms itself to *never access the resource* – in which case it is maybe counter-intuitive that $\langle\langle a\rangle\rangle\mathcal{X}\, p$ should be true. In other words, a can ensure that it is forever *able* to access the resource, but only by never *actually* accessing it.[12] Indeed, while a can force the *possibility* of achieving p to be true forever, the actual achievement of p destroys that possibility.

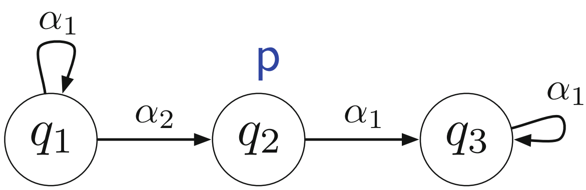

Fig. 3. Having the cake *or* eating it: model $\mathcal{M}_0$ with a single agent a. The transitions between states are labeled by the actions chosen by agent a.

Example 10 (Nested Strategic Operators). Non-persistence of strategic commitments in nested strategic formulas (like in $\langle\langle a\rangle\rangle\mathcal{G}\,\langle\langle a\rangle\rangle\mathcal{X}\, p$) also contradicts the

[12] This is the famous "have the cake or eat it" dilemma. One can keep *being able to eat the cake*, but only by never eating the cake.

observation that a player's choice constrains the outcomes that can be achieved by other players. Consider the ATL^* formula $\langle\langle A\rangle\rangle\langle\langle B\rangle\rangle\gamma$. It is easy to see that, according to the semantics of ATL^*, the formula is equivalent to $\langle\langle B\rangle\rangle\gamma$ for any pair A, B of coalitions (intersecting or not). Thus, none of A's strategy can influence the outcome of B's play, which is opposite to what we typically assume in strategic reasoning.

Alternative Semantics of Strategic Play. What are the alternatives? Let us analyze them using the example formula $\langle\langle 1,2\rangle\rangle\mathcal{G}\,\langle\langle 2,3\rangle\rangle\mathcal{X}\,p$.

1. **Irrevocable Strategies.** At the point of evaluation of $\langle\langle 1,2\rangle\rangle\mathcal{G}\,\langle\langle 2,3\rangle\rangle\mathcal{X}\,p$ the strategies of agents 1 and 2 are selected and fixed. When evaluating the subformula $\langle\langle 2,3\rangle\rangle\mathcal{X}\,p$ only the strategy of agent 3 can vary. A natural, straightforward way of obtaining this semantics with minimal change to the standard semantics of ATL^* is to *update* the model when agents choose a strategy, so that their future choices must be consistent with that strategy, but otherwise keeping semantics (definition of strategies, etc.) as is. We call these *irrevocable strategies* (see [3]), since a commitment to a strategy can never be revoked in this semantics, and denote by IATL the version of ATL adopting (memoryless) irrevocable strategies in its semantics.
2. **Strategy Contexts.** At the point of evaluation of $\langle\langle 1,2\rangle\rangle\mathcal{G}\,\langle\langle 2,3\rangle\rangle\mathcal{X}\,p$ the strategies of agents 1 and 2 are selected and fixed, but when evaluating the subformula $\langle\langle 2,3\rangle\rangle\mathcal{X}\,p$ agent 2 is granted the freedom to *change its strategy in order to achieve the current goal*, i.e. $\mathcal{X}\,p$. Thus, both agents 2 and 3 can choose new strategies, and moreover they can do that under the assumption that agent 1 remains committed to his strategy selected at the point of evaluation of $\langle\langle 1,2\rangle\rangle\mathcal{G}\,\langle\langle 2,3\rangle\rangle\mathcal{X}\,p$. This is a simple case of what we will later call *strategy contexts.*

ATL with Irrevocable Strategies. A strategy in game theory is usually understood as a plan that completely prescribes the player's behaviour, in all conceivable situations and for all future moments. An alternative semantics for strategic quantifiers takes this into account by adopting *irrevocable* strategies, implemented through the mechanism of *model update.*

Definition 6 (Model Update). *Let $\mathcal{M}$ be a CGM, A a coalition, and s_A a strategy for A. The* update of $\mathcal{M}$ by s_A, *denoted $M \dagger s_A$, is the model $\mathcal{M}$ where the choice of each agent $i \in A$ is fixed by the strategy $s_A[i]$; that is, $d_i(q) = \{s_i(q)\}$ for each state q.*

The semantics of ATL^* *with irrevocable strategies* (IATL^*) is now defined as follows, where q is a state in a CGS $\mathcal{M}$:

$[\mathcal{M}, q \models \langle\langle A\rangle\rangle\gamma$ iff there is a joint strategy s_A such that for every path $\lambda \in \mathsf{outcome\_plays}_{\mathcal{M}}(q, s_A)$ we have $\mathcal{M} \dagger s_A, \lambda \models \gamma$.

Depending on whether memory-based strategies, or only memoryless strategies, are allowed two different versions of ATL with irrevocable strategies emerge: MATL and IATL. For further details on these, we refer the reader to [3,4].

ATL with Strategy Contexts. A somewhat different and more flexible approach has been proposed by Brihaye et al. [19]. Instead of a "hard" model update that transforms the CGM according to the chosen strategy, the model is kept intact and the strategy is only added to the *strategy context*. The context collects strategies being currently executed, and hence influences the outcome paths that can occur. On the other hand, since the model itself does not change, each strategy can be revoked – in particular when an agent chooses another strategy in a nested cooperation modality. Formally, let s_A be a joint strategy of agents A (the current strategy context), and let t_B be a new joint strategy of agents B. We define the *context update* $s_A \circ t_B$ as the joint strategy f for agents in $A \cup B$ such that $f[i] = t_B[i]$ for $i \in B$ and $f[i] = s_A[i]$ for $i \in A \setminus B$. That is, the new strategies from t_B are added to the context, possibly replacing some of the previous ones. The semantic rule for strategic modalities becomes:

$\mathcal{M}, q, f \models \langle\langle A \rangle\rangle \gamma$ iff there is a joint strategy s_A for the agents in A such that for every path $\lambda \in \mathsf{outcome\_plays}(q, f \circ s_A)$ we have that $\mathcal{M}, \lambda, f \circ s_A \models \gamma$.

Additionally, $\mathcal{M}, q \models \varphi$ iff $\mathcal{M}, q, f_\emptyset \models \varphi$ where $f_\emptyset$ is the only joint strategy of the empty coalition (i.e., the empty tuple).

For more details and a thorough analysis of the model checking problem for ATL with strategy contexts, we refer the reader to [19]. A proof of the undecidability of the satisfiability problem for ATL with strategy contexts can be found in [79].

4.2 Making Strategies Explicit

In this section, we discuss several proposed variations of ATL with explicit references to strategies in the logical language.

Counterfactual ATL (CATL), proposed by van der Hoek et al. [52], extends ATL with operators of "counterfactual commitment" $\mathcal{C}_i(\sigma, \varphi)$ where i is an agent, σ is a term symbol standing for a strategy, and φ is a formula. The informal reading of $\mathcal{C}_i(\sigma, \varphi)$ is: "*if it were the case that agent i committed to strategy σ, then φ would hold*". The semantics is based on model updates, like the IATL semantics presented in Sect. 4.1:

$$\mathcal{M}, q \models \mathcal{C}_i(\sigma, \varphi) \text{ iff } \mathcal{M} \dagger [\![\sigma]\!]_i, q \models \varphi$$

where $[\![\sigma]\!]_i$ is the strategy of agent i denoted by the strategy term σ.

ATL *with intentions* (ATLI), proposed by Jamroga et al. [59], is similar to CATL, but its counterfactual operators have a different flavour: $(\mathbf{str}_i\sigma)\varphi$ reads as "*suppose that agent i* intends *to play strategy σ, then φ holds*". An intention is a kind of commitment – it persists – but it can be revoked by switching to another intention. Semantically, this is done by an additional "marking" of the intended actions in the concurrent game model. Moreover, strategies can be nondeterministic, which provides semantic tools for e.g. partial strategies as well as explicit release of commitments. Thus, Jamroga et al. [59] provide in fact the semantics of ATL based on strategy contexts (here called intentions).

However, ATLI does not allow to quantify over intentions, and hence allows only for limited context change. ATLI and its richer variant called ATLP ("ATL with plausibility" see [29]) have been used to e.g. characterize game-theoretic solution concepts and outcomes that can be obtained by rational agents. We discuss this and show some examples in Sect. 5.

Alternating-time temporal logic with explicit strategies (ATLES), see [81], is a revised version of CATL which dispenses with the counterfactual operators. Instead, strategic modalities are subscripted by *commitment functions* which are partial functions of the form $\rho = \{a_1 \mapsto \sigma_1, \ldots, a_l \mapsto \sigma_l\}$ where each a_j is an agent and σ_j is a strategy term. The meaning of a formula such as $\langle\langle A \rangle\rangle_\rho \mathcal{G}\, \varphi$ is that there exists a strategy for $A \cup \{a_1, \ldots, a_l\}$ where each a_j is required to play $[\![\sigma_j]\!]$ such that φ will hold. Note, that ATLES formulae also involve strategy commitment. Consider, for instance, formula $\langle\langle A \rangle\rangle_\rho \mathcal{G}\, \langle\langle A \rangle\rangle_\rho \mathcal{F}\, \varphi$. If A is a subset of the domain of ρ then in the evaluation of the subformula $\langle\langle A \rangle\rangle_\rho \mathcal{F}\, \varphi$, A is bound to play the same joint strategy it selected for the outer modality $\langle\langle A \rangle\rangle_\rho \mathcal{G}$.

Alternating-time temporal epistemic logic with actions (ATEL − A), proposed by Ågotnes [2], enables reasoning about the interplay between explicit strategies of bounded length and agents' knowledge.

Strategy Logic, introduced by Chatterjee et al. [31,32], treats strategies in two-player turn-based games as explicit first-order objects and enables specifying important properties of non-zero-sum games in a simple and natural way. In particular, the one-alternation fragment of strategy logic subsumes ATL* and is strong enough to express the existence of Nash equilibria and secure equilibria.

The idea of treating strategies explicitly in the language and quantifying over them is subsequently followed up in a series of papers, e.g. in [61–63] where strategy logic is extended and generalized to concurrent games, and a decidable fragment (as complex as ATL*) of it is identified and studied.

5 Reasoning About Games

ATL* and its variations are closely related to basic concepts in game theory. Firstly, their models are derived from those used in game theory. Secondly, their semantics are based on the notions of *strategies* and their *outcomes*, central in a game-theoretic analysis. In this section we give a brief overview of how to relate game theory and strategic logics. We begin with the relation between games (as viewed and analyzed in game theory) and concurrent game models. Then, we present logics which can be used to *characterize* solution concepts and logics which can *use* such solution concepts to reason about the outcome of games and the ability of rational players. For a more substantial treatment on solution concepts we refer the reader to the chapters by Bonanno [18], Pacuit [68], and Perea [74] in this book.

5.1 Representing Games as Concurrent Games Models

Standard models of modal logics correspond to strategic games, as shown in [9,52]. Moreover, concurrent game models have a close relationship to strategic and

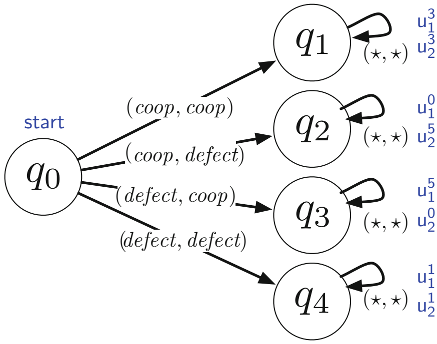

Fig. 4. Prisoner's Dilemma modelled as CGM $\mathcal{M}_{pris}$.

extensive form games, cf. [59]. We illustrate the correspondence with two examples of how strategic and extensive game frames compare to concurrent game models. The major difference is that CGMs lack the notion of payoff/outcome. However, we recall after [9,59] that CGMs can embed strategic games (cf. Example 11) and extensive games with perfect information (cf. Example 12) in a natural way. This can be done, e.g., by adding auxiliary propositions to the leaf nodes of tree-like CGMs that describe the payoffs of agents. Under this perspective, concurrent game structures can be seen as a strict generalisation of extensive form games.

In formal terms, consider first any strategic game and let U be the set of all possible utility values in it. For each value $v \in U$ and agent $a \in \mathbb{A}\mathrm{gt}$, we introduce a proposition $\mathsf{u}_\mathsf{a}^\mathsf{v}$ and put $\mathsf{u}_\mathsf{a}^\mathsf{v} \in V(q)$ iff a gets a payoff v in state q.

Example 11 (Prisoner's Dilemma as CGM). The Prisoner's Dilemma (Example 1) can also be represented by the following CGM:

$$(\{1,2\}, \{q_0,\ldots,q_4\}, \{\mathit{defect}, \mathit{coop}\}, \mathsf{Act}, \mathsf{out}, V)$$

with $\mathsf{Act}(\mathsf{a}, q) = \{\mathit{defect}, \mathit{coop}\}$ for all players a and states q, $\mathsf{out}(q_0, \mathit{coop}, \mathit{coop}) = q_1$, $\mathsf{out}(q_0, \mathit{coop}, \mathit{defect}) = q_2$, $\mathsf{out}(q_0, \mathit{defect}, \mathit{coop}) = q_3$, $\mathsf{out}(q_0, \mathit{defect}, \mathit{defect}) = q_4$, and $\mathsf{out}(q_i, a_1, a_2) = q_i$ for $i = 1,\ldots,4$ and $a_1, a_2 \in \{\mathit{defect}, \mathit{coop}\}$. The CGM is shown in Fig. 4 where the labeling function V is defined over $\mathit{Prop} = \{\mathsf{start}\} \cup \{\mathsf{u}_\mathsf{a}^\mathsf{v} \mid \mathsf{a} \in \mathbb{A}\mathrm{gt}, \mathrm{v} \in \{0,1,3,5\}\}$ as shown in the figure; e.g., we have $V(q_2) = \{\mathsf{u}_1^0, \mathsf{u}_2^5\}$ representing that players 1 and 2 receive utility of 0 and 5, respectively, if strategy profile $(\mathit{coop}, \mathit{defect})$ is played.

Example 12 (Bargaining). This example shows that CGMs are also rich enough to model (possibly infinite) extensive form games. Consider bargaining with time discount (cf. [65,76]). Two players, a_1 and a_2, bargain over how to split goods worth initially $w_0 = 1$ euro. After each round without agreement, the subjective worth of the goods reduces by *discount rates* δ_1 (for player a_1) and δ_2 (for player a_2). So, after t rounds the goods are worth $\langle \delta_1^t, \delta_2^t \rangle$, respectively. Subsequently, a_1 (if t is even) or a_2 (if t is odd) makes an offer to split the goods in proportions

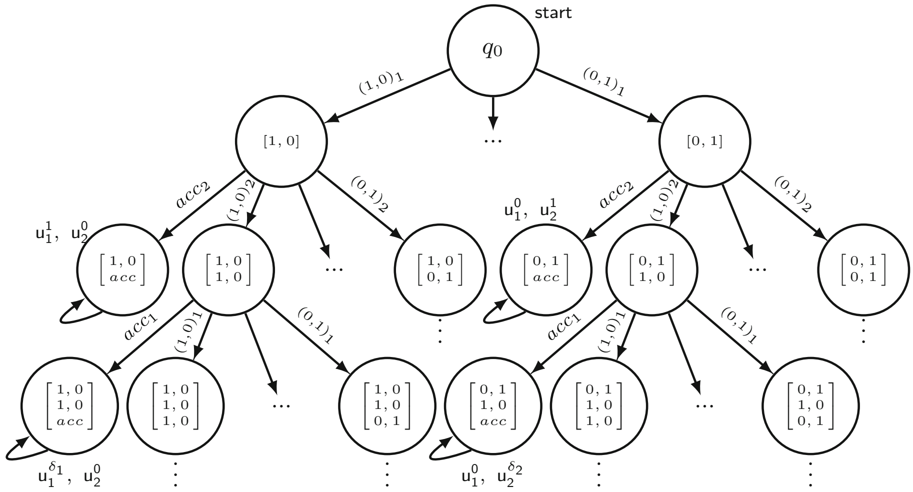

Fig. 5. CGM $\mathcal{M}_{barg}$ modeling the bargaining game.

$\langle x, 1-x \rangle$, and the other player accepts or rejects it. If the offer is accepted, then a_1 takes $x\delta_1^t$, and a_2 gets $(1-x)\delta_2^t$; otherwise the game continues.

The CGM corresponding to this extensive form game is shown in Fig. 5. Note that the model has a tree-like structure with infinite depth and an infinite branching factor. Nodes represent various states of the negotiation process, and arcs show how agents' moves change the state of the game. A node label refers to the history of the game for better readability. For instance, $\begin{bmatrix} 0,1 \\ 1,0 \\ acc \end{bmatrix}$ has the meaning that in the first round 1 offered $\langle 0, 1 \rangle$ which was rejected by 2. In the next round 2's offer $\langle 1, 0 \rangle$ has been accepted by 1 and the game has ended.

5.2 Characterization of Solution Concepts and Abilities

Rationality can be approached in different ways. Research within game theory understandably favours work on the *characterization* of various types of rationality (and defining most appropriate solution concepts). Applications of game theory, also understandably, tend toward *using* the solution concepts in order to predict the outcome in a given game; in other words, to "solve" the game. In this section we discuss logics which address both aspects. A natural question is *why* we need logics for describing and using solution concepts. In our opinion there are at least three good reasons: (i) Logical descriptions of solution concepts help for better understanding of their inner structures; e.g. interrelations can be proven by means of logical reasoning. (ii) Model checking provides an automatic way to verify properties of games and strategy profiles; e.g. whether a given profile is a Nash equilibrium in a given game or whether there is a Nash equilibrium at all. (iii) Often, the logical characterization of solution concepts is a necessary first step for using them to reason about rational agents in a flexible way, i.e. for allowing a flexible description of rational behavior rather than having a static pre-defined notion, hard-coded in the semantics of a logic.

We first give an overview of logics able to characterize solution concepts before we consider logics to reason about rational agents using characterizations of solution concepts.

Characterizing Solution Concepts in Strategic Games. In [50], a modal logic for characterizing solution concepts was presented. The main construct of the logic is $[\beta]\varphi$ where β ranges over preference relations, and complete and partial strategy profiles. The three kinds of operators have the following meaning, where i, $pref_i$, and σ, represent a player, her preference relation, and a complete strategy profile, respectively:

$[pref_i]\varphi$: φ holds in all states at least as preferred to i as the current one.
$[\sigma]\varphi$: φ will hold in the final state if all players follow σ.
$[\sigma_{-i}]\varphi$: φ will hold in all final states if all players, apart from i, follow σ.

These basic operators can be used to describe solution concepts. For instance, the formula $BR_i(\sigma) \equiv (\neg[\sigma_{-i}]\neg[pref_i]\varphi) \rightarrow [\sigma]\varphi$ expresses that σ_i is a *best response* to σ_{-i} with respect to φ: if there is a strategy for i (note that σ_{-i} does not fix a strategy for i) such that the reachable state satisfies φ and is among the most preferred ones for player i; then, the strategy σ_i (which is included in σ) does also bring about φ. Then, the property that σ is a Nash equilibrium can be captured with the formula $NE(\sigma) \equiv \bigwedge_{i\in\mathbb{A}\mathrm{gt}} BR_i(\sigma)$.

The above characterization of Nash equilibrium illustrates that, in order to assign properties to specific strategies, the strategies (or better: associated syntactic symbols) must be explicit in the object language. In Sect. 4.2 we have discussed some ATL-like logics of this kind that allow to reason about the outcome of specific strategies.

In ATLI, proposed by Jamroga et al. [59] (cf. Sect. 4.2) for example, best response strategies can be characterized as follows (where U is assumed to be a *finite* set of utility values, and $\mathsf{u}_\mathsf{a}^{\geq \mathrm{v}} \equiv \bigvee_{v\in U} \mathsf{u}_\mathsf{a}^\mathrm{v}$ expresses that agent a gets a utility value of at least v):

$$BR_a(\sigma) \equiv (\mathbf{str}_{\mathbb{A}\mathrm{gt}\setminus\{a\}}\sigma[\mathbb{A}\mathrm{gt}\setminus\{a\}]) \bigwedge_{v\in U} \Big((\langle\langle a\rangle\rangle\mathcal{F}\,\mathsf{u}_\mathsf{a}^\mathrm{v}) \rightarrow ((\mathbf{str}_a\sigma[a])\langle\langle\emptyset\rangle\rangle\mathcal{F}\,\mathsf{u}_\mathsf{a}^{\geq\mathrm{v}})\Big).$$

$BR_a(\sigma)$ refers to $\sigma[a]$ being a best response strategy for a against $\sigma[\mathbb{A}\mathrm{gt}\setminus\{a\}]$. The first counterfactual operator occurring in $BR_a(\sigma)$ fixes the strategies for all players except a. Then, each conjunct corresponds to a utility value v and expresses that if player a has a strategy to eventually achieve v (given the fixed strategies of the other players); then, a's strategy $\sigma[a]$ does eventually guarantee at least v. That is, $\sigma[a]$ is at least as good as any other strategy of a against the other players' strategies $\sigma[\mathbb{A}\mathrm{gt}\setminus\{a\}]$. The best response strategy allows to characterize Nash equilibria and subgame perfect Nash equilibria:

$$NE(\sigma) \equiv \bigwedge_{a\in\mathbb{A}\mathrm{gt}} BR_a(\sigma) \qquad \text{and} \qquad SPN(\sigma) \equiv \langle\langle\emptyset\rangle\rangle\mathcal{G}\,NE(\sigma).$$

Example 13 (Prisoner's Dilemma Continued). We continue Example 11. Suppose the strategy term σ represents the strategy profile in which both players execute *coop* in q_0 and an arbitrary action in the "final" states q_1, q_2, q_3, q_4. Let us now justify $\mathcal{M}_{dil}, q_0 \models BR_1(\sigma)$. The first operator $(\mathbf{str}_{\mathbb{A}\mathrm{gt}\setminus\{1\}}\sigma[\mathbb{A}\mathrm{gt} \setminus \{1\}])$ fixes the strategy of player 2, i.e. to cooperate. Then, player 1 has a strategy to obtain payoff u_1^3 and u_1^0 (expressed by $\langle\!\langle 1\rangle\!\rangle\mathcal{F}\,\mathsf{u}_1^{\mathsf{v}}$) by playing *coop* and *defect* in q_0, respectively. Hence, player 1's strategy contained in σ also guarantees a payoff of u_1^3 (expressed by $(\mathbf{str}_1\sigma[1])\langle\!\langle\emptyset\rangle\!\rangle\mathcal{F}\,\mathsf{u}_1^{\geq \mathsf{v}}$). This shows that player 1's strategy contained in σ is indeed a best response to player 2's strategy contained in σ.

We have another look at the given characterization of a best response strategy; in particular, at the temporal operator $\mathcal{F}$ used in the characterization. The antecedent $\langle\!\langle a\rangle\!\rangle\mathcal{F}\,\mathsf{u}_{\mathsf{a}}^{\mathsf{v}}$ requires that player a achieves v somewhere along every resulting path; it is true for the *greatest* value v along the path. In contrast, if we replace $\mathcal{F}$ by $\mathcal{G}$ the antecedent is only satisfied for the smallest value v; for, it has to be true in *every* state along a path. In general, we can use any of the *unary* temporal operators $\mathcal{X}, \mathcal{G}, \mathcal{F}, \_\mathcal{U}\psi, \psi\mathcal{U}\_$ and define variants BR_a^T, NE^T, SPN^T where T stands for any of these temporal operators and replaces $\mathcal{F}$ everywhere in the characterizations above. We refer to them as T-best response etc., each corresponding to a *different temporal pattern* of utilities. For example, we may assume that *agent a gets v* if a utility of at least v is guaranteed for every time moment ($T = \mathcal{G}$), or if it is eventually achieved ($T = \mathcal{F}$), and so on. In [59] it is shown that the $\mathcal{F}$-Nash equilibrium corresponds to its game-theoretic counterpart. This is obvious from the way games were encoded into CGMs: utility values were added to terminal states.

In Bulling et al. [29], these concepts are further generalized to *general solution concepts*. They evaluate strategies with respect to path formulae: the utility of a strategy depends on the truth of specific path formula. Furthermore, ATL *with plausibility* is introduced which extends ATL with intentions in several respects.

Further approaches for characterizing solution concepts, which we cannot discuss in detail due to lack of space, are proposed in [9,14,52].

Reasoning about the Outcome of Rational Play. The logics discussed in the previous paragraph allow to characterize game-theoretic solution concepts. It is also interesting to *use* game-theoretic solution concepts to reason about rational players. Although players have limited ability to predict the future often some lines of action seem more sensible or realistic than others. If a rationality criterion is available, we obtain means to focus on a proper subset of possible plays and to reason about the abilities of players.

Game logic with preferences (GLP), proposed by van Otterloo et al. [53], was designed to address the *outcome* of rational play in extensive form games with perfect information. The central idea of GLP is facilitated by the *preference operator* $[A : \varphi]$, interpreted as follows: If the truth of φ can be enforced by group A, then we remove from the model all the actions of A that do *not* enforce it and evaluate ψ in the resulting model. Thus, the evaluation of GLP formulae is underpinned by the assumption that *rational agents satisfy their preferences* *whenever they can*. This is a way of *using* solution concepts to reason about rational outcome.

The ideas behind ATLI and GLP were combined and extended in ATL *with plausibility (* ATLP *)*, proposed by Bulling et al. [29]. The logic allows to reason about various rationality assumptions of agents in a flexible way. For this purpose, sets of rational strategy profiles can be specified in the object language in order to analyze agents' play if only these strategy profiles were allowed. For example, if we again consider the Prisoner's Dilemma CGM from Example 13, a typical formula has the following form: $\mathbf{set\text{-}pl}\,\sigma.NE^{\mathcal{F}}(\sigma)\mathbf{Pl}\,\langle\!\langle\{1\}\rangle\!\rangle\mathcal{X}\,((\mathsf{u}_1^3\wedge\mathsf{u}_2^3)\vee(\mathsf{u}_1^1\wedge\mathsf{u}_2^1))$. The formula expresses that if it is supposed to be rational to follow $\mathcal{F}$-Nash equilibrium strategy profiles; then, player 1 can guarantee that the players will both get a payoff of 1 or both get a payoff of 3. Similar to ATLI, $NE^{\mathcal{F}}(\sigma)$ describes all Nash equilibrium strategies. The term $\sigma.NE^{\mathcal{F}}(\sigma)$ collects all these strategy profiles and the operator $\mathbf{set\text{-}pl}\cdot$ assumes that they describe rational behavior. Finally, operator **Pl** assumes that all agents play indeed rationally, and restricts their choices to rational ones; that is, Nash equilibria in this example. The restriction to rational behavior rules out all other alternatives. The logic also allows to characterize generalized versions of classical solution concepts through the characterization of patterns of payoffs by temporal formulae and quantification over strategy profiles. For further details, we refer to [24,29]. In [26] ATLP was enriched with an epistemic dimension, more precisely combined with the logic CSL discussed in Sect. 6.4, to reason about rational players under incomplete information.

Game Logic (GL) from Parikh [69] is another logic to reason about games, more precisely about determined two-player games. It builds upon propositional dynamic logic (PDL) and extends it with new operators. The work in [53,66,67] commits to a particular view of rationality (Nash equilibria, undominated strategies etc.). Finally, we would also like to mention the related work in [14] on rational dynamics and in [17] on modal logic and game theory.

6 Strategic Reasoning Under Incomplete Information

6.1 Incomplete Information Models and Uniform Strategies

The decision making capabilities and abilities of strategically reasoning players are influenced by the knowledge they possess about the world, other players, past actions, etc. So far we have considered structures of *complete and (almost) perfect information* in the sense that players are completely aware of the rules and structure of the game system and of the current state of the play. The only information they lack is the choice of actions of the other players at the current state. However, in reality this is rarely the case: usually players have only partial information about the structure and the rules of the game, as well as the precise history and the current state of the game. It is important to note that strategic ability crucially depends on the players' knowledge. In the following we are concerned with the following question: *What can players achieve in a game if they are only partially informed about its structure and the current state?*

Following the tradition of epistemic logic, we model the players' incomplete information[13] by *indistinguishability relations* $\sim \subseteq \mathsf{St} \times \mathsf{St}$ on the state space. We write $q \sim_{\mathsf{a}} q'$ to describe player a's inability to discern between states q and q'. Both states appear identical from a's perspective. The indistinguishability relations are traditionally assumed to be equivalence relations. The knowledge of a player is then determined as follows: a player *knows* a property O in a state q if O is the case in all states indistinguishable from q for that player.

How does the incomplete information, modelled by an indistinguishability relation, affect the abilities of players? In the case of ATL^* with complete information, abilities to achieve goals were modeled by strategies defined on states or (play) histories, i.e. memoryless and memory-based strategies. A basic assumption in the case of incomplete information is that *a player has the same choices in indistinguishable states*, otherwise he/she would have a way to discern between these states.

Formally, a concurrent game with incomplete information is modelled by a *concurrent epistemic game structure* (CEGS), which is a tuple

$$\mathcal{S} = (\mathbb{A}\mathrm{gt}, \mathsf{St}, \{\sim_{\mathsf{a}} \mid \mathsf{a} \in \mathbb{A}\mathrm{gt}\}, \mathsf{Act}, \mathsf{act}, \mathsf{out})$$

where $(\mathbb{A}\mathrm{gt}, \mathsf{St}, \mathsf{Act}, \mathsf{act}, \mathsf{out})$ is a CGS (cf. Definition 5) and $\sim_{\mathsf{a}}$ is the indistinguishability relation of player a over St, one per agent in $\mathbb{A}\mathrm{gt}$, such that if $q \sim_{\mathsf{a}} q'$ then $\mathsf{act}_{\mathsf{a}}(q) = \mathsf{act}_{\mathsf{a}}(q')$.

Just like a CGM, a *concurrent epistemic game model* (CEGM) is defined by adding to a CEGS a labeling of game states with sets of atomic propositions. Note that models of perfect information (CGMs) can be seen as a special case of CEGMs where each $\sim_{\mathsf{a}}$ is the smallest reflexive relation (i.e., such that $q \sim_{\mathsf{a}} q'$ iff $q = q'$).

Example 14 (Prisoner's Escape with Incomplete Information). We now explore the consequences of incomplete information in the Prisoner's escape scenario from Example 3. Recall that Frank knows the two passwords but suppose now that he does not know which one is for which guard. Equivalently, we can assume that he does not know how the guards look and which guard is at which exit. Hence, Frank does not know which password to use where. Surely, Charlie knows the guards and who is at which exit. In this setting, Frank is still able to escape with Charlie's active help. That is, Frank asks Charlie about the guards, which we now model explicitly with the action *ask*. When asked, Charlie replies by telling the truth. But at states q_1 and q_2 Charlie still has the choice of cooperating by keeping quiet or defecting by warning the guards when he sees Frank going to the respective exit. A CEGM $\mathcal{M}'_{esc}$ modelling this scenario is shown in Fig. 6.

[13] Traditionally in game theory two different terms are used to indicate lack of information: "incomplete" and "imperfect". Usually, the former refers to uncertainties about the game structure and rules, while the latter refers to uncertainties about the history, current state, etc. of the specific *play of the game*. Here we will use the latter term in about the same sense, whereas we will use "incomplete information" more loosely, to indicate any possible relevant lack of information.

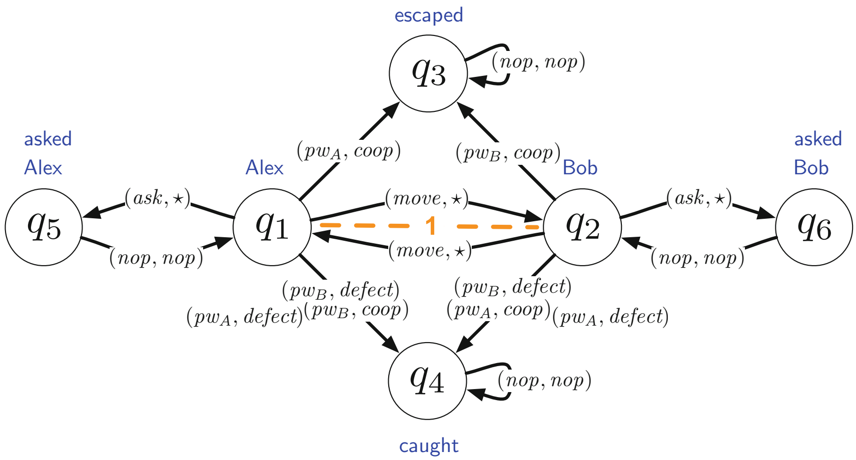

Fig. 6. Prisoner's escape with incomplete information.

The relation $q_1 \sim_1 q_2$ represents that player 1 (Frank), does not know which guard is at which entrance.

How do these perceptual limitations affect the agents' abilities? Under the assumption of complete information, Frank and Charlie can guarantee that Frank will eventually escape from q_1 and q_2 by a simple memoryless strategy for both: $s_{\{1,2\}}(q_1) = (pw_A, coop)$, $s_{\{1,2\}}(q_2) = (pw_B, coop)$, and arbitrarily defined for the other states in $\mathcal{M}_{esc}$. This strategy, however, is *not* feasible if Frank's incomplete information is taken into account because it prescribes to him different actions in the indistinguishable states q_1 and q_2. Actually, it is easy to see that in that sense there is no feasible *memoryless* strategy which achieves the property from q_1 or q_2. This is so because a strategy must be successful from all epistemic alternatives for the player (Frank). For example, his action prescribed by the strategy at q_2 must also be successful from q_1 and vice versa. This claim will become precise later, when we present the formal semantics.

However, Frank has a feasible *memory-based* strategy which guarantees that he can eventually escape, again in cooperation with Charlie. Firstly, from state q_1 or q_2 Frank asks Charlie about the guards, thus *learns* about the environment, and then goes back to use the correct password and to escape if Charlie cooperates. Formally, the reason for a successful memory-based strategy is that the *histories* $q_1q_5q_1$ and $q_2q_6q_2$ can be distinguished by Frank.

We will analyze the interaction between memory and information more formally in Sect. 6.3.

The above example indicates that the notion of strategy must be refined in order to be consistent with the incomplete information setting. An executable strategy must assign the same choices to indistinguishable situations. Such strategies are called *uniform*, e.g. see [58].

Definition 7 (Uniform Strategies). *Let $\mathcal{M}$ be a CEGM over sets of states St. A memoryless strategy s_{a} (over $\mathcal{M}$) is* uniform *if the following condition is satisfied:*

$$for\ all\ states\ q, q' \in \mathsf{St},\ if\ q \sim_{\mathsf{a}} q'\ then\ s_{\mathsf{a}}(q) = s_{\mathsf{a}}(q').$$

For memory-based *strategies we lift the indistinguishability between states to indistinguishability between (play) histories. Two* histories $\pi = q_0 q_1 \ldots q_n$ *and* $\pi' = q'_0 q'_1 \ldots q'_{n'}$ *are said to be* indistinguishable *for agent* a*, denoted by* $\pi \approx_{\mathsf{a}} \pi'$*, if and only if,* $n = n'$ *and* $q_i \sim_{\mathsf{a}} q'_i$ *for* $i = 0, \ldots, n$.[14]
A memory-based strategy s_{a} *is* uniform *if the following condition holds:*

$$for\ all\ histories\ \pi, \pi' \in \mathsf{St}^+,\ if\ \pi \approx_{\mathsf{a}} \pi'\ then\ s_{\mathsf{a}}(\pi) = s_{\mathsf{a}}(\pi').$$

Analogously to perfect information, a uniform joint strategy for a group A is a tuple of individual uniform strategies, one per member of A.[15]

6.2 Expressing Strategic Ability Under Uncertainty in ATL*

Agents' incomplete information and use of memory can be incorporated into ATL* in different ways, see e.g. [8, 56–58, 78]. In [78] a natural taxonomy of four strategy types was proposed: I (respectively i) stands for *complete* (respectively *incomplete*) *information*, and R (respectively r) refers to *perfect recall* (respectively *no recall*). The approach of Schobbens et al. [78] was syntactic in the sense that cooperation modalities were extended with subscripts: $\langle\langle A \rangle\rangle_{xy}$ where x indicates the use of memory in the strategies (memory-based if $x = R$ / memoryless if $x = r$) and y indicates the information setting (complete information if $y = I$ and incomplete information if $y = i$).

Here, we take a semantic approach. We assume that the object language of ATL/ATL* stays the same, but the semantics is parameterized with the strategy type – yielding four different semantic variants of the logic, labeled accordingly ($\mathsf{ATL_{IR}}$, $\mathsf{ATL_{Ir}}$, $\mathsf{ATL_{iR}}$, $\mathsf{ATL_{ir}}$). As a consequence, we obtain the following semantic relations:

$\models_{IR}$: complete information and memory-based strategies;
$\models_{Ir}$: complete information and memoryless strategies;
$\models_{iR}$: incomplete information and memory-based strategies;
$\models_{ir}$: incomplete information and memoryless strategies.

Given a CEGM $\mathcal{M}$, the two complete information semantic variants are obtained by updating the main semantic clause from Sect. 3.2 as follows:

[14] This corresponds to the notion of synchronous perfect recall according to [41].

[15] Note that uniformity of a joint strategy is based on individual epistemic relations, rather than any collective epistemic relation (representing, e.g., A's common, mutual, or distributed knowledge, cf. Sect. 6.4). This is because executability of agent a's choices within strategy s_A should only depend on what a can observe and deduce.

Alternative semantics where uniformity of joint strategies is defined in terms of knowledge of the group as a whole have been discussed in [36, 48].

$\mathcal{M}, q \models_{IR} \langle\langle A \rangle\rangle \gamma$ iff there is a *memory-based* joint strategy $\mathbf{s}_A$ for A such that $\mathcal{M}, \lambda \models \gamma$ for every play $\lambda \in \mathsf{outcome\_plays}(q, \mathbf{s}_A)$;
$\mathcal{M}, q \models_{Ir} \langle\langle A \rangle\rangle \gamma$ iff there is a *memoryless* joint strategy $\mathbf{s}_A$ for A such that $\mathcal{M}, \lambda \models \gamma$ for every play $\lambda \in \mathsf{outcome\_plays}(q, \mathbf{s}_A)$.

For the imperfect information variants we have:

$\mathcal{M}, q \models_{iR} \langle\langle A \rangle\rangle \gamma$ iff there is a *uniform memory-based* joint strategy $\mathbf{s}_A$ for A such that $\mathcal{M}, \lambda \models \gamma$ for every play $\lambda \in \bigcup_{q' \in \mathsf{St} \text{ s.t. } q \sim_A q'} \mathsf{outcome\_plays}(q', \mathbf{s}_A)$;
$\mathcal{M}, q \models_{ir} \langle\langle A \rangle\rangle \gamma$ iff there is a *uniform memoryless* joint strategy $\mathbf{s}_A$ for A such that $\mathcal{M}, \lambda \models \gamma$ for every play $\lambda \in \bigcup_{q' \in \mathsf{St} \text{ s.t. } q \sim_A q'} \mathsf{outcome\_plays}(q', \mathbf{s}_A)$;

where $\sim_A := \bigcup_{\mathsf{a} \in A} \sim_{\mathsf{a}}$ is used to capture the *collective knowledge of coalition* A. It is clear from the definition that this particular notion of collective knowledge refers to what everybody in A knows, e.g. [41]. For a discussion of other possibilities, we refer the reader to Sect. 6.4.

Example 15 (Prisoner's Escape: Example 14 Continued). We formalize some properties from Example 14.

1. For all $q \in \mathsf{St} \backslash \{q_4\}$ we have $\mathcal{M}'_{esc}, q \models_{Ir} \langle\langle f, c \rangle\rangle \mathcal{F}\, \mathsf{escaped}$: under the assumption of complete information coalition $\{f, c\}$ can guarantee that Frank can escape by using a memoryless strategy.
2. $\mathcal{M}'_{esc}, q_1 \not\models_{ir} \langle\langle f, c \rangle\rangle (\neg \mathsf{asked})\, \mathcal{U}\, \mathsf{escaped}$: under the assumption of incomplete information Frank and Charlie cannot guarantee that Frank will eventually escape without asking Charlie about the identity of the guards. This is true even in the case of memory-based strategies:
$\mathcal{M}'_{esc}, q_1 \not\models_{iR} \langle\langle f, c \rangle\rangle (\neg \mathsf{asked})\, \mathcal{U}\, \mathsf{escaped}$.
3. $\mathcal{M}'_{esc}, q_1 \models_{iR} \langle\langle f, c \rangle\rangle \mathcal{F}\, \mathsf{escaped}$: under the assumption of incomplete information Frank and Charlie can guarantee that Frank will eventually escape by using a uniform memory-based strategy.

In [28] incomplete information has been additionally classified according to *objective* and *subjective* ability. Here, we only consider *subjective* ability; that is, $\langle\langle A \rangle\rangle \gamma$ means that A is not only able to execute the right strategy but A can also identify the strategy. The mere existence of a winning strategy (without A being able to find it) is not sufficient under this interpretation. This is why, when evaluating $\langle\langle A \rangle\rangle \gamma$ in state q, all epistemic alternatives of q with respect to $\sim_A$ are taken into account. Again, we will discuss some other possibilities in Sect. 6.4.

Finally, we would like to add a note on the treatment of nested strategic modalities in ATL. When a nested strategic modality is interpreted, the new strategy does not take into account the previous sequence of events: Agents are effectively forgetting what they have observed before. This can lead to counter-intuitive behaviors in the presence of perfect recall and incomplete information. To overcome this, just recently a "no-forgetting semantics" for ATL has been proposed in Bulling et al. [30].

6.3 Comparing Semantics of Strategic Ability

Semantic variants of ATL are 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? Different answers to these questions induce different semantics of strategic ability, and they clearly give rise to different analysis 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 is whether the semantic variants generate different sets of valid (and dually satisfiable) sentences. In this section, we show a comparison of the validity sets for ATL^* with respect to the four semantic variants presented in the previous section. A detailed analysis and technical results can be found in [28].

The comparison of the validity sets is important for at least two reasons. Firstly, many logicians identify a logic with the set of sentences that are valid in the logic. Thus, by comparing validity sets we compare the respective logics in the traditional sense. Perhaps more importantly, 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 (CEGS). The difference between games "induced" by different semantics lies in available strategies and the winning conditions for them.

We recall that we use superscripts (e.g., '*') to denote the syntactic variant of ATL, and subscripts to denote the semantic variant being used. For example, $\mathsf{ATL}^*_{\mathsf{ir}}$ denotes the language of ATL^* interpreted with the semantic relation $\models_{ir}$, that is, the one which assumes incomplete information and memoryless strategies. Moreover, we will use $\mathit{Valid}(L)$ to denote the set of validities of logic L, and $\mathit{Sat}(L)$ to denote the set of satisfiable formulas in L.

Perfect vs. Incomplete Information. We begin by comparing properties of games with limited information to those where players can always recognize the current state of the world. Firstly, we recall that complete information can be seen as a special case of incomplete information: each CGM can be seen as a CEGM in which each indistinguishability relation is taken as the smallest reflexive relation. Hence, every valid formula of $\mathsf{ATL}^*_{\mathsf{ir}}$ is also a validity of $\mathsf{ATL}^*_{\mathsf{Ir}}$: if there were a CEGM $\mathcal{M}$ with $\mathcal{M} \not\models_{Ir} \varphi$ then also $\mathcal{M} \not\models_{ir} \varphi$ would be the case. On the other hand, the formula $\langle\!\langle A\rangle\!\rangle\mathcal{F}\,\varphi \leftrightarrow \varphi \vee \langle\!\langle A\rangle\!\rangle\mathcal{X}\,\langle\!\langle A\rangle\!\rangle\mathcal{F}\,\varphi$ is a validity of $\mathsf{ATL}_{\mathsf{Ir}}$ but not of $\mathsf{ATL}_{\mathsf{ir}}$, which shows that the containment is strict even in the limited syntactic fragment of ATL.[16]

The argument for $\mathsf{ATL}_{\mathsf{iR}}$ vs. $\mathsf{ATL}_{\mathsf{IR}}$ is analogous. Thus, we get that $\mathit{Valid}(\mathsf{ATL}_{\mathsf{ir}}) \subsetneq \mathit{Valid}(\mathsf{ATL}_{\mathsf{Ir}})$ and $\mathit{Valid}(\mathsf{ATL}_{\mathsf{ir}}) \subsetneq \mathit{Valid}(\mathsf{ATL}_{\mathsf{Ir}})$, and the same for the broader language of ATL^*.

[16] The equivalence between $\langle\!\langle A\rangle\!\rangle\mathcal{F}\,\varphi$ and $\varphi \vee \langle\!\langle A\rangle\!\rangle\mathcal{X}\,\langle\!\langle A\rangle\!\rangle\mathcal{F}\,\varphi$ is extremely important since it provides a fixpoint characterization of $\langle\!\langle A\rangle\!\rangle\mathcal{F}\,\varphi$. The fact that $\langle\!\langle A\rangle\!\rangle\mathcal{F}\,\varphi \leftrightarrow \varphi \vee \langle\!\langle A\rangle\!\rangle\mathcal{X}\,\langle\!\langle A\rangle\!\rangle\mathcal{F}\,\varphi$ is not valid under incomplete information is one of the main reasons why constructing verification and satisfiability checking algorithms is so difficult for incomplete information strategies.

Memory-based vs. Memoryless Strategies. The comparison of memory-based and memoryless strategies is technically more involved. Firstly, we observe that for any *tree-like* CGM $\mathcal{M}$ the sets of memory-based and memoryless strategies coincide. Secondly, one can show that every CGM $\mathcal{M}$ and state q in $\mathcal{M}$ can be unfolded into an equivalent (more precisely, bisimilar) tree-like CGM $T(\mathcal{M}, q)$ as in [3]. These two observations imply that $\mathsf{ATL}^*_{\mathsf{Ir}} \subseteq \mathsf{ATL}^*_{\mathsf{IR}}$; for, if $\mathcal{M}, q \not\models_{IR} \varphi$ then $T(\mathcal{M}, q), q \not\models_{IR} \varphi$ (by the latter observation) and $T(\mathcal{M}, q), q \not\models_{Ir} \varphi$ (by the first observation). Moreover, the formula $\varphi \equiv \langle\!\langle A \rangle\!\rangle(\mathcal{F}\,\varphi_1 \wedge \mathcal{F}\,\varphi_2) \leftrightarrow \langle\!\langle A \rangle\!\rangle\mathcal{F}\,((\varphi_1 \wedge \langle\!\langle A \rangle\!\rangle\mathcal{F}\,\varphi_2) \vee (\varphi_2 \wedge \langle\!\langle A \rangle\!\rangle\mathcal{F}\,\varphi_1))$ is a validity of $\mathsf{ATL}^*_{\mathsf{IR}}$ but not of $\mathsf{ATL}^*_{\mathsf{Ir}}$, which shows that the inclusion is strict.[17] Note, however, that φ is *not* a formula of ATL. Indeed, it is well known that the semantics given by $\models_{IR}$ and $\models_{Ir}$ coincide in ATL, cf. [8,78]. As a consequence, we obtain that $\mathit{Valid}(\mathsf{ATL}^*_{\mathsf{Ir}}) \subsetneq \mathit{Valid}(\mathsf{ATL}^*_{\mathsf{IR}})$ and $\mathit{Valid}(\mathsf{ATL}_{\mathsf{Ir}}) = \mathit{Valid}(\mathsf{ATL}_{\mathsf{IR}})$.

We observe that strict subsumption holds already for the language of ATL^+ which allows cooperation modalities to be followed by a Boolean combination of simple path formulae.

Finally, we consider the effect of memory in the incomplete information setting. The idea is the same as for perfect information, but the unfolding of a CEGM into an equivalent tree-like CEGM is technically more complex, as one has to take into account the indistinguishability relations (see [28] for details). To show that the inclusion is strict, we use $\langle\!\langle A \rangle\!\rangle\mathcal{X}\,\langle\!\langle A \rangle\!\rangle\mathcal{F}\,\varphi \rightarrow \langle\!\langle A \rangle\!\rangle\mathcal{F}\,\varphi$ which is valid in $\mathsf{ATL}_{\mathsf{iR}}$ but not in $\mathsf{ATL}_{\mathsf{ir}}$.[18] Thus, we get that $\mathit{Valid}(\mathsf{ATL}_{\mathsf{ir}}) \subsetneq \mathit{Valid}(\mathsf{ATL}_{\mathsf{iR}})$, and analogously for the broader language of ATL^*.

Summary. We have obtained above the following hierarchy of logics:

$$\mathit{Valid}(\mathsf{ATL}^*_{\mathsf{ir}}) \subsetneq \mathit{Valid}(\mathsf{ATL}^*_{\mathsf{iR}}) \subsetneq \mathit{Valid}(\mathsf{ATL}^*_{\mathsf{Ir}}) \subsetneq \mathit{Valid}(\mathsf{ATL}^*_{\mathsf{IR}}),$$
$$\text{and } \mathit{Valid}(\mathsf{ATL}_{\mathsf{ir}}) \subsetneq \mathit{Valid}(\mathsf{ATL}_{\mathsf{iR}}) \subsetneq \mathit{Valid}(\mathsf{ATL}_{\mathsf{Ir}}) = \mathit{Valid}(\mathsf{ATL}_{\mathsf{IR}}).$$

Equivalently, we can observe the following pattern in the sets of *satisfiable* sentences:

$$\mathit{Sat}(\mathsf{ATL}^*_{\mathsf{IR}}) \subsetneq \mathit{Sat}(\mathsf{ATL}^*_{\mathsf{Ir}}) \subsetneq \mathit{Sat}(\mathsf{ATL}^*_{\mathsf{iR}}) \subsetneq \mathit{Sat}(\mathsf{ATL}^*_{\mathsf{ir}}),$$
$$\text{and } \mathit{Sat}(\mathsf{ATL}_{\mathsf{IR}}) = \mathit{Sat}(\mathsf{ATL}_{\mathsf{Ir}}) \subsetneq \mathit{Sat}(\mathsf{ATL}_{\mathsf{iR}}) \subsetneq \mathit{Sat}(\mathsf{ATL}_{\mathsf{ir}}).$$

The first, and most important, conclusion is that all four semantic variants of ability are different with respect to the properties of games they induce. Moreover, the results capture formally the usual intuition: complete information is a particular case of incomplete information, memory-based games are special

[17] The formula expresses decomposability of conjuctive goals: being able to achieve $\varphi_1 \wedge \varphi_2$ must be equivalent to having a strategy that achieves first φ_1 and φ_2, or vice versa. It is easy to see that the requirement holds for agents with perfect memory, but not for ones bound to use memoryless strategies (and hence to play the same action whenever the game comes back to a previously visited state).

[18] The formula states that, if A has an opening move and a follow-up strategy to achieve eventually φ, then both strategies can be combined into a single strategy enforcing eventually φ already from the initial state.

cases of memoryless games, and information is a more distinguishing factor than memory.

On a more general level, the results 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, the study reveals that some natural properties – usually taken for granted when reasoning about actions – 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), decomposability of conjunctive goals, and the duality between necessary and obtainable outcomes in a game (cf. [28] for an example). The first kind of property is especially important for practical purposes, since fixpoint equivalences provide the basis for most model checking and satisfiability checking algorithms. Last but not least, the results show that the language of ATL^* is sufficiently expressive to distinguish the main notions of ability.

6.4 Epistemic Extensions of ATL

Reasoning about Knowledge. In this section we consider how the language of ATL can be combined with that of *epistemic logic*, in order to reason about the interplay of knowledge and ability more explicitly. The basic epistemic logic involves modalities for individual agent's knowledge K_i, with $K_i\varphi$ interpreted as "agent i knows that φ". Additionally, one can consider modalities for collective knowledge of groups of agents: *mutual knowledge* ($E_A\varphi$: "everybody in group A knows that φ"), *common knowledge* ($C_A\varphi$: "all the agents in A know that φ, and they know that they know it etc."), and *distributed knowledge* ($D_A\varphi$: "if the agents could share their individual information, they would be able to recognize that φ").

The formal semantics of epistemic operators is defined in terms of indistinguishability relations $\sim_1, ..., \sim_k$, given for instance in a concurrent epistemic game model:

$$\mathcal{M}, q \models K_i\varphi \text{ iff } \mathcal{M}, q' \models \varphi \text{ for all } q' \text{ such that } q \sim_i q'.$$

The accessibility relation corresponding to E_A is defined as $\sim_A^E = \bigcup_{i\in A} \sim_i$, and the semantics of E_A becomes

$$\mathcal{M}, q \models E_A\varphi \text{ iff } \mathcal{M}, q' \models \varphi \text{ for all } q' \text{ such that } q \sim_A^E q'.$$

Likewise, common knowledge C_A is given semantics in terms of the relation $\sim_A^C$ defined as the transitive closure of $\sim_A^E$:

$$\mathcal{M}, q \models C_A\varphi \text{ iff } \mathcal{M}, q' \models \varphi \text{ for all } q' \text{ such that } q \sim_A^C q'.$$

Finally, distributed knowledge D_A is based on the relation $\sim_A^D = \bigcap_{i\in A} \sim_i$, with the semantic clause defined analogously. For a more extensive exposition of epistemic logic, we refer the reader to [41,49,54].

Bringing Strategies and Knowledge Together: ATEL**.** The *alternating-time temporal epistemic logic* ATEL was introduced in [55,56] as a straightforward combination of the multi-agent epistemic logic and ATL in order to formalize reasoning about the interaction of knowledge and abilities of agents and coalitions. ATEL enables specification of various modes and nuances of interaction between knowledge and strategic abilities, e.g.: $\langle\langle A\rangle\rangle\varphi \to \mathrm{E}_A\langle\langle A\rangle\rangle\varphi$ (if group A can bring about φ then everybody in A knows that they can), $\mathrm{E}_A\langle\langle A\rangle\rangle\varphi \land \neg\mathrm{C}_A\langle\langle A\rangle\rangle\varphi$ (the agents in A have mutual knowledge but not common knowledge that they can enforce φ); $\langle\langle i\rangle\rangle\varphi \to \mathrm{K}_i\neg\langle\langle \mathbb{A}\mathrm{gt}\setminus\{i\}\rangle\rangle\neg\varphi$ (if i can bring about φ then she knows that the rest of agents cannot prevent it), etc.

Models of ATEL are *concurrent epistemic game models* (CEGM): $\mathcal{M} = (\mathbb{A}\mathrm{gt}, \mathsf{St}, \mathsf{Act}, \mathsf{act}, \mathsf{out}, V, \sim_1, ..., \sim_k)$ combining the CGM-based models for ATL and the multi-agent epistemic models. That is, the same models are used for ATEL and the Schobbens' $\mathsf{ATL}_{\mathsf{xy}}$ variants of ATL as those presented in Sect. 6.2. The semantics of ATEL simply combines the semantic clauses from ATL and those from epistemic logic.

While ATEL extends both ATL and epistemic logic, it also raises a number of conceptual problems. 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. Unfortunately, neither of these can be expressed in ATEL.[19] A number of approaches have been proposed to overcome this problem. Most of the solutions agree that only *uniform* strategies (i.e., strategies that specify the same choices in indistinguishable states) are really executable, cf. our exposition of ATL variants for incomplete information in Sect. 6.2. 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 "leader"), a subgroup ("headquarters committee") or even another group of agents ("consulting company").

Epistemic Levels of Strategic Ability. There are several possible interpretations of A's ability to bring about property γ, formalized by formula $\langle\langle A\rangle\rangle\gamma$, under imperfect information:

1. There exists a behavior specification σ_A (not necessarily executable!) for agents in A such that, for every execution of σ_A, γ holds;
2. There is a uniform strategy s_A such that, for every execution of s_A, γ holds *(A has objective ability to enforce γ)*;
3. A knows that there is a uniform s_A such that, for every execution of s_A, γ holds *(A has a strategy "de dicto" to enforce γ)*;
4. There is a uniform s_A such that A knows that, for every execution of s_A, γ holds *(A has a strategy "de re" to enforce γ)*.

[19] For a formal argument, see [2,57].

Note that the above interpretations form a sequence of increasingly stronger levels of ability – each next one implies the previous ones.

Case 4 is arguably most interesting, as it formalizes the notion of agents in *A knowing how to play in order to enforce* γ. However, the statement "A knows that every execution of s_A satisfies γ" is precise only if A consists of a single agent a. Then, we take into account the paths starting from states indistinguishable from the current one according to a (i.e., $\bigcup_{q' \text{ with } q \sim_a q'} \mathsf{outcome\_plays}(q', s_a)$). In case of multiple agents, there are several different "modes" in which they can know the right strategy. That is, given strategy s_A, coalition A can have:

- *Common knowledge* that s_A is a winning strategy. This requires the least amount of additional communication when coordinating a joint strategy: it is sufficient that agents in A agree upon a total order over their collective strategies before the game; then, during the game, they always choose the maximal strategy (with respect to this order) out of all the strategies that they commonly identify as winning;
- *Mutual knowledge* that s_A is a winning strategy: everybody in A knows that s_A is winning;
- *Distributed knowledge* that s_A is a winning strategy: if the agents share their knowledge at the current state, they can identify the strategy as winning;
- *"Leader"*: the strategy can be identified by an agent $a \in A$;
- *"Headquarters committee"*: the strategy can be identified by a subgroup $A' \subseteq A$;
- *"Consulting company"*: the strategy can be identified by another group B;
- ...other cases are also possible.

Expressing Levels of Ability: Constructive Knowledge. The issue of expressing various knowledge-related levels of ability through a suitable combination of strategic and epistemic logics has attracted significant attention. Most extensions (or refinements) of ATL, proposed as solutions, cover only some of the possibilities, albeit in an elegant way [2,66,78]. Others, such as [58,60], offer a more general treatment of the problem at the expense of an overblown logical language. *Constructive Strategic Logic* (CSL), proposed by Jamroga and Ågotnes [57], aims at a solution which is both general and elegant. However, there is a price to pay. In CSL, formulae are interpreted over *sets of states* rather than single states. We write $\mathcal{M}, Q \models \langle\langle A \rangle\rangle \varphi$ to express the fact that A must have a strategy which is successful for all "opening" states from Q. New epistemic operators $\mathbb{K}_i, \mathbb{E}_A, \mathbb{C}_A, \mathbb{D}_A$ for "practical" or "constructive" knowledge yield the set of states for which a single evidence (i.e., a successful strategy) should be presented (instead of checking if the required property holds in each of the states separately, like standard epistemic operators do).

Formally, the semantics of CSL (in its broadest syntactic variant CSL*) over concurrent epistemic game models is defined by the following clauses:

$\mathcal{M}, Q \models p$ iff $p \in \pi(q)$ for every $q \in Q$;
$\mathcal{M}, Q \models \neg\varphi$ iff $\mathcal{M}, Q \not\models \varphi$;
$\mathcal{M}, Q \models \varphi \wedge \psi$ iff $\mathcal{M}, Q \models \varphi$ and $\mathcal{M}, Q \models \psi$;

$\mathcal{M}, Q \models \langle\langle A \rangle\rangle \varphi$ iff there is a uniform strategy s_A such that $\mathcal{M}, \lambda \models \varphi$ for every $\lambda \in \cup_{q \in Q}$ $\mathsf{outcome\_plays}(q, s_A)$;

$\mathcal{M}, Q \models \mathbb{K}_i \varphi$ iff $\mathcal{M}, \{q' \mid \exists_{q \in Q}\ q \sim_i q'\} \models \varphi$;
$\mathcal{M}, Q \models \mathbb{C}_A \varphi$ iff $\mathcal{M}, \{q' \mid \exists_{q \in Q}\ q \sim_A^C q'\} \models \varphi$;
$\mathcal{M}, Q \models \mathbb{E}_A \varphi$ iff $\mathcal{M}, \{q' \mid \exists_{q \in Q}\ q \sim_A^E q'\} \models \varphi$;
$\mathcal{M}, Q \models \mathbb{D}_A \varphi$ iff $\mathcal{M}, \{q' \mid \exists_{q \in Q}\ q \sim_A^D q'\} \models \varphi$.

The semantic clauses for path subformulae are the same as in ATL^*. Additionally, we define that $\mathcal{M}, q \models \varphi$ iff $\mathcal{M}, \{q\} \models \varphi$.

A nice feature of CSL is that standard knowledge operators can be defined using constructive knowledge, e.g., as $K_a \varphi \equiv \mathbb{K}_a \langle\langle \emptyset \rangle\rangle \varphi \mathcal{U} \varphi$[20]. It is easy to see that $\mathcal{M}, q \models \mathbb{K}_a \langle\langle \emptyset \rangle\rangle \varphi \mathcal{U} \varphi$ iff $\mathcal{M}, q' \models \varphi$ for every q' such that $q \sim_a q'$.

We point out that in CSL:

1. $\mathbb{K}_a \langle\langle a \rangle\rangle \varphi$ refers to agent a having a strategy *"de re"* to enforce φ (i.e. having a successful uniform strategy and knowing the strategy);
2. $K_a \langle\langle a \rangle\rangle \varphi$ refers to agent a having a strategy *"de dicto"* to enforce φ (i.e. knowing only that *some* successful uniform strategy is available);
3. $\langle\langle a \rangle\rangle \varphi$ expresses that agent a has a uniform strategy to enforce φ *from the current state* (but not necessarily even knows about it).

Thus, $\mathbb{K}_a \langle\langle a \rangle\rangle \varphi$ captures the notion of a's knowing *how to play* to achieve φ, while $K_a \langle\langle a \rangle\rangle \varphi$ refers to knowing only *that a successful play is possible*. This extends naturally to abilities of coalitions, with $\mathbb{C}_A \langle\langle A \rangle\rangle \varphi, \mathbb{E}_A \langle\langle A \rangle\rangle \varphi, \mathbb{D}_A \langle\langle A \rangle\rangle \varphi$ formalizing common, mutual, and distributed knowledge how to play, $\mathbb{K}_a \langle\langle A \rangle\rangle \varphi$ capturing the "leader" scenario, and so on (and similarly for different levels of knowledge "de dicto"). We conclude this topic with the following example.

Example 16 (Market Scenario). Consider an industrial company that wants to start production, and looks for a good strategy when and how it should do it. The market model is depicted in Fig. 7. The economy is assumed to run in simple cycles: after the moment of bad economy (bad-market), there is always a good time for small and medium enterprises (s&m), after which the market tightens and an oligopoly emerges. At the end, the market gets stale, and we have stagnation and bad economy again.

The company c is the only agent whose actions are represented in the model. The company can wait (action *wait*) or decide to start production: either on its own (*own-production*), or as a subcontractor of a major company (*subproduction*). Both decisions can lead to either loss or success, depending on the current market conditions. However, the company management cannot recognize the market conditions: bad market, time for small and medium enterprises, and oligopoly market look the same to them, as the epistemic links for c indicate.

The company can call the services of two marketing experts. Expert 1 is a specialist on oligopoly, and can recognize oligopoly conditions (although she cannot

[20] We cannot replace $\varphi \mathcal{U} \varphi$ by φ when the latter is a path formula, as then $\langle\langle \emptyset \rangle\rangle \varphi$ would not be a formula of CSL.

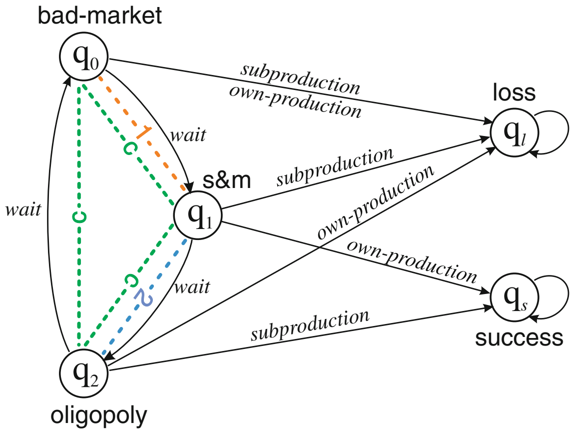

Fig. 7. Simple market: model $\mathcal{M}_{mark}$

distinguish between bad economy and s&m market). Expert 2 can recognize bad economy, but he cannot distinguish between other types of market. The experts' actions have no influence on the actual transitions of the model, and are omitted from the graph in Fig. 7. It is easy to see that the company cannot identify a successful strategy on its own: for instance, for the small and medium enterprises period, we have that $\mathcal{M}_{mark}, q_1 \models \neg\mathbb{K}_c\langle\langle c\rangle\rangle\mathcal{F}\,\mathsf{success}$. It is not even enough to call the help of a single expert: $\mathcal{M}_{mark}, q_1 \models \neg\mathbb{K}_1\langle\langle c\rangle\rangle\mathcal{F}\,\mathsf{success}\wedge\neg\mathbb{K}_2\langle\langle c\rangle\rangle\mathcal{F}\,\mathsf{success}$, or to ask the experts to independently work out a common strategy: $\mathcal{M}_{mark}, q_1 \models \neg\mathbb{E}_{\{1,2\}}\langle\langle c\rangle\rangle\mathcal{F}\,\mathsf{success}$. Still, the experts can propose the right strategy if they join forces and share available information: $\mathcal{M}_{mark}, q_1 \models \mathbb{D}_{\{1,2\}}\langle\langle c\rangle\rangle\mathcal{F}\,\mathsf{success}$.

This is not true anymore for bad market, i.e., $\mathcal{M}_{mark}, q_0 \models \neg\mathbb{D}_{\{1,2\}}\langle\langle c\rangle\rangle\mathcal{F}\,\mathsf{success}$, because c is a memoryless agent, and it has no uniform strategy to enforce $\mathsf{success}$ from q_0 at all. However, the experts can suggest a more complex scheme that involves consulting them once again in the future, as evidenced by $\mathcal{M}_{mark}, q_0 \models \mathbb{D}_{\{1,2\}}\langle\langle c\rangle\rangle\mathcal{X}\ \mathbb{D}_{\{1,2\}}\langle\langle c\rangle\rangle\mathcal{F}\,\mathsf{success}$.

7 Deductive Systems and Logical Decision Problems

7.1 Validity and Satisfiability in ATL and ATL*

Characterizing the valid and, dually, the satisfiable formulae of a given logic by means of sound and complete deductive systems is a fundamental logical problem. Few such deductive systems have been developed so far for the logics discussed here, and these are mostly axiomatic systems. We will briefly present the one for ATL.

Axiomatic Systems for CL and ATL. In Pauly [71–73] it was shown that the conditions of liveness, safety, superadditivity, and $\mathbb{A}$gt-maximality in Definition 4 can be captured by a few simple axiom schemes presented below, where $A_1, A_2 \subseteq \mathbb{A}\mathrm{gt}$ are any coalitions of players:

1. Complete set of axioms for classical propositional logic.
2. $\langle\langle \mathbb{A}\mathrm{gt} \rangle\rangle \mathcal{X} \top$
3. $\langle\langle A \rangle\rangle \mathcal{X} \bot$
4. $\neg \langle\langle \emptyset \rangle\rangle \mathcal{X} \varphi \rightarrow \langle\langle \mathbb{A}\mathrm{gt} \rangle\rangle \mathcal{X} \neg \varphi$
5. $\langle\langle A_1 \rangle\rangle \mathcal{X} \varphi \wedge \langle\langle A_2 \rangle\rangle \mathcal{X} \psi \rightarrow \langle\langle A_1 \cup A_2 \rangle\rangle \mathcal{X} (\varphi \wedge \psi)$ for any disjoint $A_1, A_2 \subseteq \mathbb{A}\mathrm{gt}$

These, together with the inference rules Modus Ponens and monotonicity:

$$\frac{\varphi \rightarrow \psi}{\langle\langle A \rangle\rangle \mathcal{X} \varphi \rightarrow \langle\langle A \rangle\rangle \mathcal{X} \psi}$$

provide a sound and complete axiomatization of the valid formulae of CL, see [71,73].

The temporal operators $\mathcal{G}$ and $\mathcal{U}$ satisfy the following validities in ATL that define them recursively as fixed points of certain monotone operators:

($\mathbf{FP}_{\mathcal{G}}$) $\langle\langle A \rangle\rangle \mathcal{G} \varphi \leftrightarrow \varphi \wedge \langle\langle A \rangle\rangle \mathcal{X} \langle\langle A \rangle\rangle \mathcal{G} \varphi$,
($\mathbf{GFP}_{\mathcal{G}}$) $\langle\langle \emptyset \rangle\rangle \mathcal{G} (\theta \rightarrow (\varphi \wedge \langle\langle A \rangle\rangle \mathcal{X} \theta)) \rightarrow \langle\langle \emptyset \rangle\rangle \mathcal{G} (\theta \rightarrow \langle\langle A \rangle\rangle \mathcal{G} \varphi)$,
($\mathbf{FP}_{\mathcal{U}}$) $\langle\langle A \rangle\rangle \psi \, \mathcal{U} \, \varphi \leftrightarrow \varphi \vee (\psi \wedge \langle\langle A \rangle\rangle \mathcal{X} \langle\langle A \rangle\rangle \psi \, \mathcal{U} \, \varphi)$,
($\mathbf{LFP}_{\mathcal{U}}$) $\langle\langle \emptyset \rangle\rangle \mathcal{G} ((\varphi \vee (\psi \wedge \langle\langle A \rangle\rangle \mathcal{X} \theta)) \rightarrow \theta) \rightarrow \langle\langle \emptyset \rangle\rangle \mathcal{G} (\langle\langle A \rangle\rangle \psi \, \mathcal{U} \, \varphi \rightarrow \theta)$.

It was proved in Goranko and van Drimmelen [47] that these axioms added to Pauly's axioms for CL, plus the rule $\langle\langle \emptyset \rangle\rangle \mathcal{G}$*-Necessitation:*

$$\frac{\varphi}{\langle\langle \emptyset \rangle\rangle \mathcal{G} \varphi}.$$

provide a sound and complete axiomatization for the validities of ATL.

No explicit complete axiomatizations for ATL*, nor for any of the variations of ATL with incomplete information, are known yet.

Decidability and Decision Methods for ATL and ATL*. A fundamental algorithmic problem in logic is whether a given logical formula is satisfiable in any model for the given logic, or dually, whether its negation is valid in the given semantics. A constructive procedure for testing satisfiability is of practical importance because it can be used to construct (to synthesize) models from formal logical specifications. Sound and complete axiomatic systems provide only semi-decision methods for testing validity, respectively non-satisfiability, while complete algorithmic decision methods exist only for logics with a decidable validity/satisfiability problem. The decidability of that problem in ATL, with **EXPTIME**-complete worst-case complexity of the decision algorithm, was first proved in van Drimmelen [39] (see also Goranko and van Drimmelen [47] for detailed proofs) by proving a bounded-branching tree-model property and using alternating tree automata, under the assumption that the number of agents is fixed. The **EXPTIME**-completeness of ATL satisfiability was later re-proved in Walther et al. [80] without the assumption of fixed number of agents. An optimal and practically implementable tableau-based constructive decision method for testing satisfiability in ATL was developed in Goranko and Shkatov [46]. Later, the decidability and **2EXPTIME**-complete complexity of the satisfiability problem for ATL* was proved in Schewe [77] using alternating tree automata.

7.2 Model Checking of ATL and ATL*

Model checking is another fundamental logical decision problem. It calls for a procedure that determines whether a given formula is true in a given model. For such procedures to be algorithmically implementable, the model must be finite, or effectively (finitely) presented. We briefly discuss model checking ATL and ATL* under the different semantic variants considered in this chapter. We focus on the main technical issues that arise in that area, namely the computational complexity of the model checking algorithms, as a measure of the inherent complexity of the underlying semantics of the logic. The relevant complexity results are summarized in Fig. 8.

	Ir	IR	ir	iR
ATL	**P**	**P**	$\mathbf{\Delta_2^P}$	Undecidable†
ATL*	**PSPACE**	**2EXPTIME**	**PSPACE**	Undecidable

Fig. 8. Overview of the exact complexity results for model checking in explicit models of formulae from the logic in the respective row with the semantics given in the column.

A deterministic polynomial-time model checking algorithm for $\mathsf{ATL}_{\mathsf{ir}}$ (and thus $\mathsf{ATL}_{\mathsf{iR}}$) is presented in Alur et al. [8]. The algorithm is based on the fixpoint characterizations of strategic-temporal modalities:

$$\langle\langle A\rangle\rangle\mathcal{G}\,\varphi \leftrightarrow \varphi \wedge \langle\langle A\rangle\rangle\mathcal{X}\,\langle\langle A\rangle\rangle\mathcal{G}\,\varphi$$
$$\langle\langle A\rangle\rangle\varphi_1\,\mathcal{U}\,\varphi_2 \leftrightarrow \varphi_2 \vee (\varphi_1 \wedge \langle\langle A\rangle\rangle\mathcal{X}\,\langle\langle A\rangle\rangle\varphi_1\,\mathcal{U}\,\varphi_2).$$

The perfect information assumption allows to compute a winning strategy step-by-step (if it exists). In the case of $\langle\langle A\rangle\rangle\mathcal{G}\,\varphi$, for example, the procedure starts with all states in which φ holds and subsequently removes states in which there is no joint action for team A to guarantee to end up in one of the states in which φ holds. Let us refer to the resulting set of states as Q_1. In the next step it is checked whether for each state in Q_1 there is a joint action of team A which guarantees to remain in Q_1. States in which such a joint action does not exist are removed from Q_1. This procedure is applied recursively until a fixed point is reached. The formula $\langle\langle A\rangle\rangle\mathcal{G}\,\varphi$ is true in all the remaining states.

The deterministic **2EXPTIME** algorithm for model checking $\mathsf{ATL}^*_{\mathsf{IR}}$ makes use of a sophisticated tree automaton construction, see [8].

Algorithms for the the remaining settings based on memoryless strategies employ model checking algorithms for CTL and CTL* (model checking is **P**-complete and **PSPACE**-complete, respectively, see Clarke et al. [35]). The key observation is that there are only finitely many *memoryless* strategies and that a strategy can be guessed in non-deterministic polynomial time. Then, all transitions not possible given the guessed strategy profile are removed from the model and the resulting *temporal* model checking problem is solved (cf. [78]). For illustration, consider $\mathsf{ATL}_{\mathsf{iR}}$ and the formula $\langle\langle A\rangle\rangle\mathcal{G}\,\varphi$. First, we guess a memoryless uniform strategy s_A of coalition A. It is easy to see that the validation whether the strategy profile is uniform or not can be done in deterministic polynomial-time. Afterwards, all transitions not possible according to s_A are removed from the model as well as all transition labels. It remains to check whether $\mathcal{G}\,\varphi$ holds on *all* possible behaviors/paths in the resulting purely temporal model. The latter corresponds to CTL model checking of $\mathsf{A}\mathcal{G}\,\varphi$. However, since φ may contain *nested* cooperation modalities we need to proceed *bottom-up* which shows that the problem is contained in $\mathbf{P}^{\mathbf{NP}^{\mathbf{P}}} = \mathbf{\Delta_2^P}$. Similarly, we obtain **PSPACE** algorithms for $\mathsf{ATL}^*_{\mathsf{Ir}}$ and $\mathsf{ATL}^*_{\mathsf{ir}}$: guess strategies and solve the CTL^* model checking problem, which can be done in $\mathbf{P}^{\mathbf{NP}^{\mathbf{PSPACE}}} = \mathbf{PSPACE}$.

We note that model checking of ATL in the case of imperfect information and memory-based strategy is undecidable, cf. [8,37].

For a more detailed overview of the complexities of model checking of these logics we refer the reader to [8,25] for ATL and ATL^*, [27] for ATL^+, a computationally better behaved fragment of ATL^*, and to [32,61–63] for more powerful and recent extensions of ATL^*.

8 Concluding Remarks: Brief Parallels with Other Logical Approaches to Strategic Reasoning

While strategic reasoning is a highly involved and complex form of reasoning, requiring strong logical and analytic skills, its seems rather surprising that until the 1980 s formal logic was seldom employed to either analyze or facilitate strategic reasoning. However, with the ongoing invasion of logic into game theory and multi-agent systems over the past 20 years, its role in both doing and analyzing strategic reasoning has become increasingly more instrumental and recognized. Logic has been successfully applied to several rather different aspects of strategic reasoning and the variety of logical systems presented and discussed here gives a good overall picture of only one of the logical approaches to strategic reasoning, viz. reasoning *about objective strategic abilities of players and coalitions pursuing a specific goal, in competitive concurrent multi-player games where the remaining players are regarded as (temporary) adversaries as far as achieving of that goal is concerned.* As mentioned in the introduction, there are several other related logic-based approaches to strategic reasoning and most of them are treated in other chapters of this book.

– *Logics of agencies, abilities and actions.* Philosophical approaches to developing logics of agency and ability, include early works of von Wright and Kanger and more recent ones by Brown [23], Belnap and Perloff [13], and Chellas

[34]. In particular, Brown [23] proposes a modal logic with non-normal modalities formalising the idea that the modality for ability has a more complex, existential-universal meaning (the agent has *some* action or choice, such that *every* outcome from executing that action (making that choice) achieves the aim), underlying the approaches to formalizing agents' ability presented here.

– STIT *logics*. These originate from the work of Belnap [13] introducing the operator *seeing to it that*, abbreviated to "STIT". The approach to strategic reasoning taken in the STIT family of logics of agency, discussed in Broersen and Herzig [21], is the closest to the one presented in this chapter and we will provide a more detailed parallel with it now. To begin with, both the STIT-based and the ATL-based approaches assume that agents act simultaneously and independently. The main conceptual difference between the family of STIT logics of agency and ATL-like logics is that the former delve into more philosophical issues of agency and emphasize the *intentional aspect* of the agents' strategies, whereas the latter take a more pragmatic view on agents and focus on the *practical effects* of their strategic abilities and choices, disregarding desires, intentions, and other less explicit attitudes. The basic STIT operator is similar to the one-step strategic modality of CL while the intended meaning of the "strategic" version of the STIT operator, SSTIT, comes very close to the intended meaning of the strategic operator $\langle\langle\rangle\rangle$ in ATL. The main technical difference between these logics is in the semantics, which is rather more general and abstract in the case of STIT as compared to ATL. Strategies in ATL models are explicit rules mapping possible game configurations to prescribed actions, whereas strategies in STIT models are implicit and essentially represented by the respective plays ('histories') that they can enable. More precisely, the formal semantics of the SSTIT operator defines 'histories' as abstract objects representing the possible courses of events. Agents' strategies are abstract sets of histories satisfying some requirements, of which the most essential one is that every strategy profile of the set of all agents intersects in a single history. This semantics essentially extends the original semantics for ATL based on "alternating transition systems", subsequently replaced by the more concrete and – in our view – more realistic semantics based on concurrent game models, presented here[21]. Due to the expressiveness of the language of STIT/SSTIT and the generality of its semantics, it naturally embeds ATL* with complete information, as well as a number of its variations considered here, as demonstrated in [20, 22]. The price to pay for that expressiveness, as it should be expected, is the generally intractable and usually undecidable complexity of STIT logics.

– *Logics for compositional reasoning about strategies*, initiated by Parikh [69] and discussed and extended in this book by Paul, Ramanujam and Simon [70], is another approach, conceptually close to the present, where strategies are treated as first-class citizens to which an endogenous, structural view is

[21] Yet, the SSTIT semantic structures relate quite naturally to path effectivity models introduced and characterized in [44], and these could provide a more feasible semantics for SSTIT.

applied, and "the study of rationality in extensive form games largely takes a functional view of strategies". In a way, this approach relates to the ATL-based one like the Propositional Dynamic Logic PDL relates to the temporal logics LTL and CTL as alternative approaches to reasoning about programs.

– *Logics of knowledge and beliefs.* As we have noted repeatedly, strategic reasoning is intimately related to players' knowledge and information. One of the deepest and most successful manifestations of logical methods in strategic reasoning is the doxastic-epistemic treatment of the concepts of individual and common rationality in game theory. This approach is treated in-depth and from different perspectives in the chapters by Bonanno [18], Pacuit [68] and Perea [74] of this book, as well as in Baltag, Smets et al. [11,12], etc. As stated in the chapter by Pacuit [68], this approach is not so focused on strategies and strategic abilities per se, but rather on the process of rational deliberation that leads players to their strategic choices and the latter are crucially dependent on the players' mutual rationality assumptions, rather than on demand for success against any – rational, adversarial, or simply random – behaviour of the others.
– *Logics for social choice theory*, discussed in the chapter by van Eijck [40] of this book, focuses on logical modeling of specific strategic abilities that arise on social choice scenarios such as voting.
– *Dynamic epistemic logic.* The relation of the ATL-based family of logics with Dynamic epistemic logic (DEL) [10,38] is more distant and implicit. DEL does not purport to reason explicitly about strategic abilities of agents, but it does provide a framework for such reasoning, in terms of which *epistemic objectives* agents can achieve by performing various epistemic actions, represented by action models.
– Lastly, for broader and more conceptual perspectives on the subject we refer the reader to the rest of this book and to van Benthem [15].

Acknowledgments. We are grateful to the participants in the Workshop on *Modelling Strategic Reasoning* held in February 2012 in the Lorentz Center, Leiden, and particularly to Nicolas Troquard and Dirk Walther, as well as to the anonymous reviewers, for their valuable comments and suggestions. Wojciech Jamroga acknowledges the support of the National Research Fund (FNR) Luxembourg under the project GALOT (INTER/DFG/12/06), as well as the support of the 7th Framework Programme of the European Union under the Marie Curie IEF project ReVINK (PIEF-GA-2012-626398). The final work of Valentin Goranko on this chapter was done while he was an invited visiting professor at the Centre International de Mathématiques et Informatique de Toulouse.

References

1. Abdou, J., Keiding, H.: Effectivity Functions in Social Choice Theory. Kluwer, Netherlands (1991)
2. Ågotnes, T.: Action and knowledge in alternating-time temporal logic. Synthese **149**(2), 377–409 (2006). Section on Knowledge, Rationality and Action

3. Ågotnes, T., Goranko, V., Jamroga, W.: Alternating-time temporal logics with irrevocable strategies. In: Samet, D. (ed.) Proceeding of TARK XI, pp. 15–24. ACM, New York (2007)
4. Ågotnes, T., Goranko, V., Jamroga, W.: Strategic commitment and release in logics for multi-agent systems (extended abstract). Technical report IfI-08-01, Clausthal University of Technology (2008)
5. Ågotnes, T., Walther, D.: A logic of strategic ability under bounded memory. J. Logic Lang. Inform. **18**(1), 55–77 (2009)
6. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. In: Proceedings of the FOCS 1997, pp. 100–109. IEEE Computer Society Press (1997)
7. Kupferman, O., Alur, R., Henzinger, T.A.: Alternating-time temporal logic. In: de Roever, W.-P., Pnueli, A., Langmaack, H. (eds.) COMPOS 1997. LNCS, vol. 1536, pp. 23–60. Springer, Heidelberg (1998)
8. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. J. ACM **49**, 672–713 (2002)
9. Baltag, A.: A logic for suspicious players. Bull. Econ. Res. **54**(1), 1–46 (2002)
10. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge and private suspicions. Technical report SEN-R9922, CWI, Amsterdam (1999)
11. Baltag, A., Smets, S.: A qualitative theory of dynamic interactive belief revision. In: Bonanno, G., van der Hoek, W., Wooldridge, M. (eds.) Logic and the Foundation of Game and Decision Theory (LOFT7), volume 3 of Texts in Logic and Games, pp. 13–60. Amsterdam University Press, Amsterdam (2008)
12. Baltag, A., Smets, S., Zvesper, J.A.: Keep 'hoping' for rationality: A solution to the backward induction paradox. Synthese **169**(2), 301–333 (2009)
13. Belnap, N., Perloff, M.: Seeing to it that: a canonical form for agentives. Theoria **54**, 175–199 (1988)
14. van Benthem, J.: Rational dynamics and epistemic logic in games. In: Vannucci, S. (ed.) Logic, Game Theory and Social Choice III, pp. 19–23. ILLC, Amsterdam (2003)
15. van Benthem, J.: Logic in Games. MIT Press, Cambridge (2013)
16. van Benthem, J.: Logic of strategies: What and how? In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) Models of Strategic Reasoning. LNCS, vol. 8972, pp. 321–332. Springer, Heidelberg (2015)
17. Bonanno, G.: Modal logic and game theory: two alternative approaches. Risk Decis. Policy **7**, 309–324 (2002)
18. Bonanno, G.: Reasoning about strategies and rational play in dynamic games. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) Models of Strategic Reasoning. LNCS, vol. 8972, pp. 34–62. Springer, Heidelberg (2015)
19. Brihaye, T., Da Costa, A., Laroussinie, F., Markey, N.: ATL with strategy contexts and bounded memory. Technical report LSV-08-14, ENS Cachan (2008)
20. Broersen, J.: CTL.STIT: enhancing ATL to express important multi-agent system verification properties. In: Proceedings of the AAMAS 2010, pp. 683–690 (2010)
21. Broersen, J., Herzig, A.: Using STIT theory to talk about strategies. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) Models of Strategic Reasoning. Lecturer Notes in Computer Science, vol. 8972, pp. 137–173. Springer, Heidelberg (2015)
22. Broersen, J., Herzig, A., Troquard, N.: Embedding alternating-time temporal logic in strategic STIT logic of agency. J. Log. Comput. **16**(5), 559–578 (2006)
23. Brown, M.A.: On the logic of ability. J. Philos. Logic **17**, 1–26 (1988)
24. Bulling, N.: Modelling and Verifying Abilities of Rational Agents. Ph.D. thesis, Clausthal University of Technology (2010)

25. Bulling, N., Dix, J., Jamroga, W.: Model checking logics of strategic ability: Complexity. In: Dastani, M., Hindriks, K.V., Meyer, J.-J.C. (eds.) Specification and Verification of Multi-Agent Systems, pp. 125–159. Springer, New York (2010)
26. Bulling, N., Jamroga, W.: Rational play and rational beliefs under uncertainty. In: Proceedings of AAMAS 2009, pp. 257–264. ACM Press, Budapest, May 2009
27. Bulling, N., Jamroga, W.: Verifying agents with memory is harder than it seemed. AI Commun. **23**(4), 389–403 (2010)
28. Bulling, N., Jamroga, W.: Comparing variants of strategic ability: How uncertainty and memory influence general properties of games. Auton. Agent. Multi-Agent Syst. **28**(3), 474–518 (2014)
29. Bulling, N., Jamroga, W., Dix, J.: Reasoning about temporal properties of rational play. Ann. Math. Artif. Intell. **53**(1–4), 51–114 (2009)
30. Bulling, N., Jamroga, W., Popovici, M.: ATL* with truly perfect recall: Expressivity and validities. In: Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014), pp. 177-182. IOS Press, Prague, August 2014
31. Chatterjee, K., Piterman, N., Henzinger, T.A.: Strategy logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 59–73. Springer, Heidelberg (2007)
32. Chatterjee, K., Henzinger, T.A., Piterman, N.: Strategy logic. Inf. Comput. **208**(6), 677–693 (2010)
33. Chellas, B.: Modal Logic: An Introduction. Cambridge University Press, Cambridge (1980)
34. Chellas, B.: On bringing it about. J. Philos. Logic **24**(6), 563–571 (1995)
35. Clarke, E.M., Emerson, E.A., Sistla, A.P.: Automatic verification of finite-state concurrent systems using temporal logic specifications. ACM Trans. Program. Lang. Syst. **8**(2), 244–263 (1986)
36. Diaconu, R., Dima, C.: Model-checking alternating-time temporal logic with strategies based on common knowledge is undecidable. Appl. Artif. Intell. **26**(4), 331–348 (2012)
37. Dima, C., Tiplea, F.L.: Model-checking ATL under imperfect information and perfect recall semantics is undecidable (2011). CoRR, abs/1102.4225
38. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Springer, Dordecht (2008)
39. van Drimmelen, G.: Satisfiability in alternating-time temporal logic. In: Proceedings of LICS 2003, pp. 208–217. IEEE Computer Society Press (2003)
40. van Eijck, J.: Strategies in social software. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) Models of Strategic Reasoning. LNCS, vol. 8972, pp. 292–317. Springer, Heidelberg (2015)
41. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning about Knowledge. MIT Press, London (1995)
42. Goranko, V.: Coalition games and alternating temporal logics. In: van Benthem, J. (ed.) Proceedings of TARK VIII, pp. 259–272. Morgan Kaufmann, San Francisco (2001)
43. Goranko, V., Jamroga, W.: Comparing semantics of logics for multi-agent systems. Synthese **139**(2), 241–280 (2004)
44. Goranko, V., Jamroga, W.: State and path coalition effectivity models for logics of multi-player games. In: Proceedings of AAMAS 2012 (2012)
45. Goranko, V., Jamroga, W., Turrini, P.: Strategic games and truly playable effectivity functions. In: Proceedings of AAMAS 2011, pp. 727–734 (2011)
46. Goranko, V., Shkatov, D.: Tableau-based decision procedures for logics of strategic ability in multiagent systems. ACM Trans. Comput. Log. **11**(1), 1–48 (2009)

47. Goranko, V., van Drimmelen, G.: Complete axiomatization and decidability of alternating-time temporal logic. Theoret. Comput. Sci. **353**(1), 93–117 (2006)
48. Guelev, D.P., Dima, C., Enea, C.: An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. J. Appl. Non-Class. Logics **21**(1), 93–131 (2011)
49. Halpern, J.Y.: Reasoning about knowledge: a survey. In: Gabbay, D.M., Hogger, C.J., Robinson, J.A. (eds.) Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 4: Epistemic and Temporal Reasoning, pp. 1–34. Oxford University Press, Oxford (1995)
50. Harrenstein, B.P., van der Hoek, W., Meyer, J.-J.C., Witteveen, C.: A modal characterization of nash equilibrium. Fundam. Informaticae **57**(2–4), 281–321 (2003)
51. Hart, S.: Games in extensive and strategic forms. In: Aumann, R.J., Hart, S. (eds.) Handbook of Game Theory with Economic Applications, vol. 1, pp. 19–40. Elsevier, North-Holland (1992)
52. van der Hoek, W., Jamroga, W., Wooldridge, M.: A logic for strategic reasoning. In: Proceedings of AAMAS 2005, pp. 157–164 (2005)
53. van der Hoek, W., van Otterloo, S., Wooldridge, M.: Preferences in game logics. In: Proceedings of AAMAS-04 (2004)
54. van der Hoek, W., Verbrugge, R.: Epistemic logic: A survey. Game Theory Appl. **8**, 53–94 (2002)
55. van der Hoek, W., Wooldridge, M.: Tractable multiagent planning for epistemic goals. In: Castelfranchi, C., Johnson, W.L. (eds.) Proceeding of AAMAS-02, pp. 1167–1174. ACM Press, New York (2002)
56. van der Hoek, W., Wooldridge, M.: Cooperation, knowledge and time: Alternating-time temporal epistemic logic and its applications. Stud. Logica **75**(1), 125–157 (2003)
57. Jamroga, W., Ågotnes, T.: Constructive knowledge: What agents can achieve under incomplete information. J. Appl. Non-Class. Logics **17**(4), 423–475 (2007)
58. Jamroga, W., van der Hoek, W.: Agents that know how to play. Fundamenta Informaticae **63**(2–3), 185–219 (2004)
59. Jamroga, W., Wooldridge, M.J., van der Hoek, W.: Intentions and strategies in game-like scenarios. In: Bento, C., Cardoso, A., Dias, G. (eds.) EPIA 2005. LNCS (LNAI), vol. 3808, pp. 512–523. Springer, Heidelberg (2005)
60. Jonker, G.: Feasible strategies in Alternating-time Temporal Epistemic Logic. Master thesis, University of Utrecht (2003)
61. Mogavero, F., Murano, A., Perelli, G., Vardi, M.Y.: What makes ATL* decidable? a decidable fragment of strategy logic. In: Ulidowski, I., Koutny, M. (eds.) CONCUR 2012. LNCS, vol. 7454, pp. 193–208. Springer, Heidelberg (2012)
62. Mogavero, F., Murano, A., Vardi, M.Y.: Reasoning about strategies. In: Lodaya, K., Mahajan, M. (eds.) Proceedings of FSTTCS 2010, LIPIcs, pp. 133–144. Atlantis Press, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2010)
63. Vardi, M.Y., Murano, A., Mogavero, F.: Relentful strategic reasoning in alternating-time temporal logic. In: Voronkov, A., Clarke, E.M. (eds.) LPAR-16 2010. LNCS, vol. 6355, pp. 371–386. Springer, Heidelberg (2010)
64. Moulin, H., Peleg, B.: Cores of effectivity functions and implementation theory. J. Math. Econ. **10**(1), 115–145 (1982)
65. Osborne, M., Rubinstein, A.: A Course in Game Theory. MIT Press, Cambridge (1994)
66. van Otterloo, S., Jonker, G.: On epistemic temporal strategic logic. Electron. Notes Theoret. Comput. Sci. **126**, 35–45 (2004)

67. van Otterloo, S., Roy, O.: Verification of voting protocols. Working paper, University of Amsterdam (2005)
68. Pacuit, E.: Dynamic models of rational deliberation in games. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) Models of Strategic Reasoning. LNCS, vol. 8972, pp. 3–33. Springer, Heidelberg (2015)
69. Parikh, R.: The logic of games and its applications. Ann. Discrete Math. **24**, 111–140 (1985)
70. Paul, S., Ramanujam, R., Simon, S.: Automata and compositional strategies in extensive form games. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) Models of Strategic Reasoning. LNCS, vol. 8972, pp. 174–201. Springer, Heidelberg (2015)
71. Pauly, M.: Logic for Social Software. Ph.D. thesis, University of Amsterdam (2001)
72. Pauly, M.: A logical framework for coalitional effectivity in dynamic procedures. Bull. Econ. Res. **53**(4), 305–324 (2001)
73. Pauly, M.: A modal logic for coalitional power in games. J. Logic Comput. **12**(1), 149–166 (2002)
74. Perea, A.: Finite reasoning procedures for dynamic games. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) Models of Strategic Reasoning. LNCS, vol. 8972, pp. 63–90. Springer, Heidelberg (2015)
75. Pinchinat, S.: A generic constructive solution for concurrent games with expressive constraints on strategies. In: Okamura, Y., Yoneda, T., Higashino, T., Namjoshi, K.S. (eds.) ATVA 2007. LNCS, vol. 4762, pp. 253–267. Springer, Heidelberg (2007)
76. Sandholm, T.W.: Distributed rational decision making. In: Weiss, G. (ed.) Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, pp. 201–258. The MIT Press, Cambridge (1999)
77. Schewe, S.: ATL* satisfiability is 2EXPTIME-Complete. In: Walukiewicz, I., Goldberg, L.A., Halldórsson, M.M., Damgård, I., Ingólfsdóttir, A., Aceto, L. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 373–385. Springer, Heidelberg (2008)
78. Schobbens, P.Y.: Alternating-time logic with imperfect recall. Electron. Notes Theoret. Comput. Sci. **85**(2), 82–93 (2004)
79. Troquard, N., Walther, D.: On satisfiability in ATL with strategy contexts. In: del Cerro, L.F., Herzig, A., Mengin, J. (eds.) JELIA 2012. LNCS, vol. 7519, pp. 398–410. Springer, Heidelberg (2012)
80. Walther, D., Lutz, C., Wolter, F., Wooldridge, M.: ATL satisfiability is indeed EXPTIME-complete. J. Logic Comput. **16**(6), 765–787 (2006)
81. Walther, D., van der Hoek, W., Wooldridge, M.: Alternating-time temporal logic with explicit strategies. In: Samet, D. (ed.) Proceedings of the TARK XI, pp. 269–278. Presses Universitaires de Louvain, New York (2007)