

Representing Strategies

Hein Duijf

Jan Broersen*

Quite some work in the ATL-tradition uses the differences between various types of strategies (positional, uniform, perfect recall) to give alternative semantics to the same logical language. This paper contributes to another perspective on strategy types, one where we characterise the differences between them on the syntactic (object language) level. This is important for a more traditional knowledge representation view on strategic content. Leaving differences between strategy types implicit in the semantics is a sensible idea if the goal is to use the strategic formalism for model checking. But, for traditional knowledge representation in terms of object language level formulas, we need to extend the language. This paper introduces a strategic STIT syntax with explicit operators for knowledge that allows us to characterise strategy types. This more expressive strategic language is interpreted on standard ATL-type concurrent epistemic game structures. We introduce rule-based strategies in our language and fruitfully apply them to the representation and characterisation of positional and uniform strategies. Our representations highlight crucial conditions to be met for strategy types. We demonstrate the usefulness of our work by showing that it leads to a critical reexamination of coalitional uniform strategies.

1 Introduction

To make way for strategic reasoning on the syntactic level we need to know how to represent the various types of strategies that have been proposed in the literature. In this paper, our aim in particular will be to provide syntactic counterparts for the various strategy types proposed in the ATL-tradition, such as positional strategies [2] and uniform strategies [10]. Our proposed language, an extension of strategic STIT including temporal and epistemic modalities and action types, is sufficiently expressive for representing positional and uniform strategies.¹ Whereas the introduction of positional and uniform strategies are intuitively appealing, our characterizations elucidate the underlying conditions to be met for these semantically defined strategy types. Although emulating ability-modalities akin to the ATL-tradition is not the main objective here, some hints are provided that guide such a future endeavour.

Our representation of uniform positional strategies enhances the understanding of a *coalition's* uniform strategy. Our result naturally invites different ways to distribute relevant strategic knowledge. This observation complements [1] where it is argued that a coalition's uniform strategy is imprecise only in the “mode” of the coalition's knowledge, referring either to common, distributive or mutual knowledge.

To further explain our view and approach, we will subsequently answer the following three questions:

1. what *is* a strategy?
2. how to represent the *performance* of a strategy?
3. how to *characterize* a rule-based strategy?

*Both authors gratefully acknowledge financial support from the ERC-2013-CoG project REINS, nr. 616512

¹Recently, a syntactic characterization of uniform strategies in *Epistemic Strategy Logic* was presented [11]. Our representation differs in that we use rule-based strategies.

what is a strategy?

It is remarkable how strong the notion of a strategy varies throughout the literature on strategic reasoning. In ATL frameworks² a strategy is a mapping that assigns an action type to every finite history of system states; in STIT frameworks³, since acting means restricting the possible futures, a strategy is identified with the futures it allows; finally Dynamic Logics⁴ take strategies to be temporally extended act structures of sequences of (atomic) actions. This paper starts with the ATL conception of strategy types and ultimately provides representations thereof.

how to represent the performance of a strategy?

To express that a strategy is actually performed, we use insights from STIT frameworks. Other frameworks express strategy performance only implicitly, safely tucked away under path quantifiers (ATL) or in quantifiers in the central modalities (Dynamic logic). This follows from the observation that the main operators $\langle\langle C \rangle\rangle\varphi$ and $[\alpha]\varphi$ in these systems are interpreted as ‘coalition C is able to ensure φ ’ and ‘after executing action α , φ holds’. This reveals that ATL enables one to reason about strategic ability and Dynamic Logics support reasoning about the results of actions, but not reasoning about the *performance* of actions or strategies here and now.

how to characterize a rule-based strategy?

Strategies are typically communicated in the form of condition-action rules. Therefore it makes sense to also logically represent them in that form using a suitable language. Recently, ATL has been extended to enable reasoning about rule-based strategies in [17]. One of the main themes there is the representation of (semantic) strategies by formulas of their proposed language, which includes rule-based strategies. A different study in [16] evaluates formulas at game-strategy pairs thereby combining aspects of *Game Logic*⁵ and strategic reasoning. They propose a multi-sorted language to express the *structure* of the strategy. The logic includes two types of conditional strategies $[\psi \mapsto a]^i$ and $\pi \mapsto \sigma$, which are interpreted as ‘player i chooses move a whenever ψ holds’ and ‘player i sticks to the specification given by σ if on the history of play, all moves made by \bar{i} conform to π ’, respectively. The ψ in the first formula is restricted to boolean combinations of propositional letters, and the second formula represents that the other players *have acted* in accordance with π . We, however, do not want to commit ourselves to these restrictions in the language or regarding the conditions; Most importantly, to represent uniform strategies epistemic conditions have to be allowed.

The paper is organized as follows: In Section 2 the well-known Concurrent Epistemic Game Models (see [2] and [8]) are introduced to provide the basis for our semantics. In contrast to the usual ATL*-divide between path and state formulae, a key idea from G.STRAT⁶ to evaluate formulas against tuples consisting of the state, path and *the current strategy profile* is implemented in Section 3. In parallel, a logical language is introduced including action types and temporal, epistemic and agency operators. As is common in STIT frameworks, the structure of a strategy can be described by specifying the (temporal) properties it ensures. This allows us to view a rule-based strategy as a set of condition-effect rules in Section 4, which are then used in Section 5 to investigate how this logical language can be fruitfully applied to represent strategy types. Certain types of rule-based strategies are used to represent positional and uniform strategies, thereby uncovering crucial conditions for these semantic strategy types. Some novel implications are drawn on coalition’s uniform strategies by appealing, not to different “modes” of coalitional knowledge but, to ways of distributing the relevant strategic knowledge. We conclude in

²See the seminal work [2] and extensions such as CATL [7], and *Strategy Logic* [6] and [12].

³See the seminal work [3], [9] and the recent extension to strategic action [5].

⁴See [4].

⁵See the original work [14] and the overview [15].

⁶G.STRAT is first introduced in [5]. It is an extension of basic STIT frameworks to a strategic and multi-agent setting.

Section 6, proof sketches are to be found in the appendix.

2 Concurrent Game Models

In this section we introduce Concurrent Game Models (see [2]) and Concurrent Epistemic Game Models (see [8]). Our treatment will be roughly in line with [1] although our notation on histories will differ to neatly support the syntactic approach in the next section.

Definition 1 (Concurrent game structure and model). *A concurrent game structure (CGS) is a tuple $\mathcal{S} = (\text{Ag}, \text{St}, \text{Act}, \text{act}, \text{out})$ which consists of:*

1. *a finite, non-empty set of agents $\text{Ag} = \{1, \dots, K\}$; the subsets of Ag are called coalitions;*
2. *a non-empty set of states St ;*
3. *a non-empty set of action types Act ;*
4. *an action manager function $\text{act} : \text{Ag} \times \text{St} \rightarrow \mathcal{P}(\text{Act})$ assigning to every player i and state q a non-empty set of actions available for execution by i at q .
An action profile is a tuple of actions $\alpha = \langle \alpha_1, \dots, \alpha_K \rangle \in \text{Act}^K$. The action profile is executable at the state q if $\alpha_i \in \text{act}(i, q)$ for every $i \in \text{Ag}$. We denote by $\text{act}(q)$ the subset of $\prod_{i \in \text{Ag}} \text{act}(i, q)$ consisting of all action profiles executable at the state q .*
5. (a) *a transition function out that assigns an outcome state $\text{out}(q, \alpha)$ to every state q and every action profile $\alpha \in \text{act}(q)$.⁷*
(b) *for every two states q, q' there is at most one $\alpha \in \text{act}(q)$ such that $\text{out}(q, \alpha) = q'$, i.e. if a state transition is labelled, then it has a unique label.*

A concurrent game model (CGM) is a CGS endowed with a labeling $L : \text{St} \rightarrow \mathcal{P}(\text{Prop})$ of the states with sets of atomic propositions from a fixed set Prop . As usual, the labeling describes which atomic propositions are true at a given state.

We fix a concurrent game structure $\mathcal{S} = (\text{Ag}, \text{St}, \text{Act}, \text{act}, \text{out})$. An action profile is used to determine a successor of a state using the transition function out . The set of the available action profiles is denoted by $\text{act}(q)$, consequently the set of possible successors of q is the set of states $\text{out}(q, \alpha)$ where α ranges over $\text{act}(q)$. An infinite sequence $\lambda = q_0 q_1 q_2 \dots$ of states from St is called a *play* if q_{k+1} is a successor of q_k for all positions $k \geq 0$. $\lambda[k]$ denotes the k -th component q_k in λ , and $\lambda[0, k]$ denotes the initial sequence, or *history*, $q_0 \dots q_k$ of λ .

A *perfect recall strategy for an agent i* is a function s_i that maps every history $\lambda[0, k]$ to an action type $s_i(\lambda[0, k]) \in \text{act}(\lambda[k])$. A *positional (aka memoryless) strategy for an agent i* is a function s_i that maps every state q to an action type. A *perfect recall strategy for a coalition $C \subseteq \text{Ags}$* , also called a coalitional strategy for C , is a function s_C mapping each agent $i \in C$ to a perfect recall strategy $s_C(i)$. Positional strategies for coalitions are defined analogously. A *strategy profile s* is a coalitional strategy for Ag . A coalitional strategy s_C *extends* s'_C , notation $s_C \sqsupseteq s'_C$, if and only if $C \supseteq C'$ and $s'_C(i) = s_C(i)$ for every $i \in C$. Given a strategy profile, we often write s_C for the coalitional strategy for C satisfying $s_C \sqsubseteq s$.

⁷In [1, p. 553 – our boldfacing, emphasis in original] they introduce out as “a *transition function* out that assigns a **unique outcome state** $\text{out}(q, \alpha)$ to every state q and every action profile α which is executable at q ”. We are puzzled about what this required uniqueness adds to the fact that out is a function. Although they do not explicitly mention what is meant, we think this uniqueness is represented in (b). The examples that they consider are in line with (b). This addition is crucial for our main Proposition 1(3) as it implies that two strategies are play-equivalent if and only if they are identical.

The set $\text{out_plays}(\lambda[0,k], s_C)$ of outcome plays of a strategy s_C for C at a history $\lambda[0,k]$ is the set of all plays $\lambda' = q_0q_1\cdots$ such that $\lambda'[0,k] = \lambda[0,k]$ and, for every $l \geq k$, there is an action profile $\alpha = \langle \alpha_1, \dots, \alpha_K \rangle$ satisfying $\alpha_i = s_C(i)(\lambda'[l])$ for all $i \in C$ and $q_{l+1} = \text{out}(q_l, \alpha)$.

Our formal results rely on the notion of *play-equivalence*:

Definition 2 (Play-equivalence). *Let λ be a play, let k be a position, let C a coalition, and let s_C, s'_C perfect recall strategies. We say that s_C and s'_C are play-equivalent at $\lambda[0,k]$ if and only if $\text{out_plays}(\lambda[0,k], s_C) = \text{out_plays}(\lambda[0,k], s'_C)$.*

It is standard to model the agent's incomplete information by extending concurrent game models:

Definition 3 (Concurrent Epistemic Game Model). *A concurrent epistemic game model (CEGM) is a tuple $\mathcal{M} = (\text{Ag}, \text{St}, \{\sim_i \mid i \in \text{Ag}\}, \text{Act}, \text{act}, \text{out}, L)$ which consists of a CGM $(\text{Ag}, \text{St}, \text{Act}, \text{act}, \text{out}, L)$ and indistinguishability relations $\sim_i \subseteq \text{St} \times \text{St}$, one for each agent. Furthermore, it is assumed that (1) \sim_i is an equivalence relation, and (2) $q \sim_i q'$ implies that $\text{act}(q, i) = \text{act}(q', i)$.*

These indistinguishability relations are straightforwardly extended to histories by: $\lambda[0,k] \sim_i \lambda'[0,k']$ iff $k = k'$ and for every $l \leq k$ we have $\lambda[l] \sim_i \lambda'[l]$. Then we introduce a third strategy type:

a *uniform strategy for an agent i* is a perfect recall strategy s_i satisfying: for all histories $\lambda[0,k], \lambda'[0,k'] \in \text{St}^+$, if $\lambda[0,k] \sim_i \lambda'[0,k']$ then $s_i(\lambda[0,k]) = s_i(\lambda'[0,k'])$.

A *uniform coalitional strategy for coalition C* is a function s_C mapping each agent $i \in C$ to a uniform strategy $s_C(i)$.

In the remainder, we mean “perfect recall strategies” when writing “strategies”, unless otherwise specified.

3 Strategic language

In the previous section we outlined the models that provide the basis for the semantics of our logical enterprise. In the current section we introduce our logical framework, which is inspired by [5].⁸ First, we introduce the syntax of the logical language. Second, we present the truth conditions of the logical formalism. It is crucial that we evaluate formulas with respect to tuples $\langle k, \lambda, s \rangle$ which include *the current strategic course of action s* (inspiration from [5]). Finally, some crucial observations on the resulting logical formalism are presented by reviewing the underlying models.

Definition 4 (Syntax). *Fix a set of propositional letters Prop , a finite set of agents Ags , and a set of action types Σ . The formulas of the language \mathcal{L} are given by:*

$$\varphi ::= p \mid \alpha_C \mid \varphi \wedge \varphi \mid \neg\varphi \mid \text{X}\varphi \mid \text{G}\varphi \mid \square\varphi \mid [C \text{ sstit}]\varphi \mid K_i\varphi,$$

where p ranges over Prop , C ranges over subsets of Ags , and α_C ranges over Σ^C .

Given a CEGM $\mathcal{M} = (\text{Ag}, \text{St}, \{\sim_i \mid i \in \text{Ag}\}, \text{Act}, \text{act}, \text{out}, L)$ with $\text{Ag} = \text{Ags}$ and $\text{Act} = \Sigma$, these formulas will be evaluated at tuples $\langle k, \lambda, s \rangle$ consisting of a strategy profile s , a play λ such that $\lambda \in \text{out_plays}(\lambda[0], s)$, and a position k . This means that the truth of formulas is evaluated with respect to a current state $\lambda[k]$, a current history $\lambda[0,k]$, a current future $\lambda[k, \infty)$, and a current strategy profile s . Obviously, by incorporating the current strategy profile into the worlds of evaluation we get a semantic explication of the performance of a strategy.

⁸Whereas [5] uses models based on the STIT-tradition, here Concurrent Game Models are used for interpreting the language.

The central agency operator is the modality $[C \text{ sstit}] \varphi$ which stands for ‘the coalition C strategically sees to it that φ holds’. Relative to a tuple $\langle k, \lambda, s \rangle$ the modality $[C \text{ sstit}] \varphi$ is interpreted as ‘the coalition C is in the process of executing s_C thereby ensuring the (temporal) condition φ ’. In addition, the language includes temporal modalities $X\varphi$ and $G\varphi$ which are interpreted, relative to a tuple $\langle k, \lambda, s \rangle$, as ‘ φ holds in the next moment after $\lambda[k]$ on λ ’ and ‘ φ holds on all future moments after $\lambda[k]$ on λ ’, respectively. In contrast to this longitudinal dimension of time, the language includes a temporal modality $\Box\varphi$ for historical necessity. The modality $\Box\varphi$ is interpreted, relative to a tuple $\langle k, \lambda, s \rangle$, as ‘ φ holds on any tuple at $\lambda[0, k]$ ’. This highlights that the truth of $\Box\varphi$ does not depend on the dynamic aspects represented by the current future and the current strategy profile, we call such formulas *moment-determinate*. Finally, we include epistemic modalities $K_i\varphi$, one for each agent, which are interpreted as ‘agent i knows that φ ’. The presented syntax and semantics are formally connected by the truth conditions for the syntactic clauses:

Definition 5 (Semantics). *Let $\mathcal{M} = (\text{Ag}, \text{St}, \{\sim_i \mid i \in \text{Ag}\}, \text{Act}, \text{act}, \text{out}, L)$ be a CEGM with $\text{Ag} = \text{Ags}$ and $\text{Act} = \Sigma$. The points of evaluation for our logical formulas are tuples $\langle k, \lambda, s \rangle$ consisting of a strategy profile s , a play $\lambda \in \text{out_plays}(\lambda[0], s)$, and a position k . The truth conditions are given by inductive definitions (suppressing the model \mathcal{M} and not listing standard propositional truth conditions):*

$$\langle k, \lambda, s \rangle \models \alpha_C \Leftrightarrow s_C(\lambda[0, k]) = \alpha_C$$

$$\langle k, \lambda, s \rangle \models X\varphi \Leftrightarrow \langle k+1, \lambda, s \rangle \models \varphi$$

$$\langle k, \lambda, s \rangle \models G\varphi \Leftrightarrow \text{for each } l \geq k : \langle l, \lambda, s \rangle \models \varphi$$

$$\langle k, \lambda, s \rangle \models \Box\varphi \Leftrightarrow \text{for every strategy } s' \text{ and every } \lambda' \in \text{out_plays}(\lambda[0, k], s') \text{ we have } \langle k, \lambda', s' \rangle \models \varphi$$

$$\langle k, \lambda, s \rangle \models [C \text{ sstit}] \varphi \Leftrightarrow \text{for every } s' \sqsupseteq s_C \text{ and every } \lambda' \in \text{out_plays}(\lambda[0, k], s') \text{ we have } \langle k, \lambda', s' \rangle \models \varphi$$

$$\langle k, \lambda, s \rangle \models K_i\varphi \Leftrightarrow \text{for every } \lambda', k' \text{ such that } \lambda'[0, k'] \sim_i \lambda[0, k] \text{ and any } s' \text{ we have } \langle k', \lambda', s' \rangle \models \varphi.^9$$

With the semantics in place, we gather some crucial observations:

- Because the truth of a propositional letter only depends on the current state, it is not surprising that the truth of any propositional formula only depends on the current state. Therefore, we will often write $\lambda[k] \models \varphi$ instead of $\langle k, \lambda, s \rangle \models \varphi$ for a propositional formula φ . This is connected to the familiar divide in ATL* syntax between state and path formulas.
- As mentioned before, a formula φ is moment-determinate if $\models \varphi \leftrightarrow \Box\varphi$. The truth of such formulas only depends on the history, so we will often write $\lambda[0, k] \models \varphi$ instead of $\langle k, \lambda, s \rangle \models \varphi$ for such formulas.
- Formulas of the form $K_i\varphi$ are moment-determinate. In particular, an agent does not know what he is doing.
- The truth of a coalitional action type only depends on the current strategy of that coalition, i.e. $\langle k, \lambda, s \rangle \models \alpha_C \leftrightarrow [C \text{ sstit}] \alpha_C$.

⁹Little is known about properties of this logic and such a formal inquiry would lead us to far astray from the current enterprise. To guide some of the formal intuitions of the reader we mention some validities without proof or conceptual motivation: \Box and $[i \text{ sstit}]$ are S5-modalities, the $[i \text{ sstit}]$ -operator is monotone in its agency argument, i.e. for $A \subseteq B$ we have $\models [A \text{ sstit}]p \rightarrow [B \text{ sstit}]p$, the temporal part is the standard discrete linear temporal logic containing X and G , some interaction principles are $\models \Box Xp \rightarrow X\Box p$, $\models [C \text{ sstit}]Xp \rightarrow X[C \text{ sstit}]p$, $\models [C \text{ sstit}]\Box p \leftrightarrow \Box p$, $\models [C \text{ sstit}]X[C \text{ sstit}]p \leftrightarrow [C \text{ sstit}]Xp$, $\models [C \text{ sstit}]G[C \text{ sstit}]p \leftrightarrow [C \text{ sstit}]Gp$, and $\diamond[A \text{ sstit}]p \wedge \diamond[B \text{ sstit}]q \rightarrow \diamond[A \cup B \text{ sstit}](p \wedge q)$ for disjoint coalitions A and B (independence of agency).

- The formula $\diamond\alpha_C$ expresses that a coalitional action type α_C is executable at a state.
- Observe that only the truth conditions for the $[C \text{ sstit}]$ -operator and the action types α_C involve the current C -strategy profile. It is clear that adding (at least one of) these is necessary to express that a certain strategy is performed. The action types are inherited from a bottom-up perspective on strategies with the action types as atomic building blocks. In contrast, the $[C \text{ sstit}]$ -operator incorporates a top-down view in that a strategy is described by the properties it ensures.

4 Rule-based strategies

A rule-based strategy consists of rules. Such a rule is composed of a condition and an effect, thereby incorporating the intuition that a rule is triggered under certain conditions and has a certain effect:

Definition 6 (Rule-based strategies). *A rule-based strategy RS is a finite set of condition-effect rules $\{c_1 \mapsto e_1, \dots, c_N \mapsto e_N\}$, which represents that condition c_n triggers effect e_n . We denote the set of conditions occurring in such a rule-based strategy RS by $\text{Cond}(RS)$, likewise the effects by $\text{Eff}(RS)$.*

Performing a rule-based strategy means that in case a rule is triggered one ensures that the corresponding effect is realized:

Definition 7 (Performing rule-based strategies). *Henceforth we fix a CEGM $(\text{Ag}, \text{St}, \{\sim_i \mid i \in \text{Ag}\}, \text{Act}, \text{act}, \text{out}, L)$ with $\text{Ag} = \text{Ags}$ and $\text{Act} = \Sigma$. Let $RS = \{c_1 \mapsto e_1, \dots, c_N \mapsto e_N\}$ be a rule-based strategy. First, we say that a coalition acts according to rule-based strategy RS at $\langle k, \lambda, s \rangle$, denoted by $\langle k, \lambda, s \rangle \models [C \text{ acc}]RS$, if and only if*

$$\langle k, \lambda, s \rangle \models [C \text{ sstit}] \bigwedge_{n \leq N} (c_n \rightarrow [C \text{ sstit}]e_n).$$

Second, we say that a coalition C performs RS at $\langle k, \lambda, s \rangle$, denoted by $\langle k, \lambda, s \rangle \models [C \text{ perf}]RS$, if and only if

$$\langle k, \lambda, s \rangle \models [C \text{ sstit}]G[C \text{ acc}]RS.$$

The formula $[C \text{ acc}]S$ is interpreted, relative to a tuple $\langle k, \lambda, s \rangle$, as ‘coalition C is in the process of executing strategy s_C thereby ensuring that the conditionals are met’. Informally, it means that coalition C is currently performing a strategy that ensures that in case a condition holds he performs a strategy ensuring the corresponding effect.

Although the nested $[C \text{ sstit}]$ operator may be puzzling at first sight, it makes perfect sense. To argue in favour we break the formula down. A rule of the rule-based strategy is formalized as $c \rightarrow [C \text{ sstit}]e$, but one should not forget that here we intend to formulate that a coalition is acting according to such a rule-based strategy. This is expressed by the second $[C \text{ sstit}]$ operator, which guarantees that one is acting accordingly not only at the current play, but also at all plays in $\text{out_plays}(\lambda[0, k], s_C)$. To formalize that a coalition is performing a rule-based strategy, we add the G operator to express that it is henceforth acting according to strategy RS .

There are two ways in which a rule-based strategy can be unsatisfactory: (a) the agent might not be able to perform a certain rule-based strategy, or (b) the action description given by a certain rule-based strategy can be underspecified. So a rule-based strategy can be viewed as a partial perfect recall strategy which is defined at a history if and only if it is possible to act accordingly and there is but one way to do so.¹⁰ To investigate this more thoroughly, we continue in our logical framework.

¹⁰This resembles the informal notion of deterministic strategies in [17, p. 204]: “move recommendations are always unique” for deterministic strategies.

Can a given rule-based strategy RS be viewed as a partial strategy? There is a straightforward way to attempt this whenever the conditions are moment-determinate, i.e. $\mathcal{M} \models c \leftrightarrow \Box c$ for each $c \in \text{Cond}(RS)$.¹¹ Using RS , we define a *partial* coalition strategy $s_C^{RS} : \text{St}^+ \rightarrow \text{Act}^C$ by:

$$s_C^{RS}(\lambda[0, k]) = \alpha_C \quad \text{iff} \quad \lambda[0, k] \models \Diamond[C \text{ acc}]RS \wedge \Box([C \text{ acc}]RS \rightarrow [C \text{ sstit}]\alpha_C).$$

The first conjunct says that the coalition is able to act accordingly, whereas the second conjunct says that performing action profile α_C is the only way to do so.

Clearly, this partial coalition strategy is defined at a history $\lambda[0, k]$ if the following conditions hold:

- (1) there is a $c \in \text{Cond}(RS)$ such that $\lambda[0, k] \models c$,
- (2) there is at least one α_C such that $\lambda[0, k] \models \Diamond[C \text{ sstit}]\alpha_C \wedge \Box([C \text{ sstit}]\alpha_C \rightarrow [C \text{ acc}]RS)$, and
- (3) there is at most one α_C such that $\lambda[0, k] \models \Diamond[C \text{ sstit}]\alpha_C \wedge \Box([C \text{ sstit}]\alpha_C \rightarrow [C \text{ acc}]RS)$.

The failure of (1) and the failure of (3) signify that the rule-based strategy is underspecified either because no rule has been triggered or because there are multiple ways to act accordingly. The failure of (2), however, indicates a practical inconsistency or a conflict in the rule-based strategy RS , because it implies that there is no way to act accordingly.

Before proceeding, we extend $\text{out_plays}(-, -)$ to pertain also to partial strategies:

Definition 8. Let s_C be a partial perfect recall strategy. We define

$$\text{out_plays}(\lambda[0, k], s_C) := \{\lambda' \in \text{St}^\omega \mid \lambda' \sqsupset \lambda[0, k] \text{ and for every } l \geq k \text{ we have: if } s_C(\lambda'[0, l]) \text{ is defined then } \lambda[l+1] \in \text{out_set}(\lambda[0, l], s_C)\}.$$

To investigate the perfect recall strategies represented by a rule-based strategy, we use the partial strategy it defines:

Proposition 1. Let RS be a rule-based strategy with only moment-determinate conditions. Let s_C^{RS} be the partial coalition strategy defined by RS . Then

1. for any profile $\langle k', \lambda', s' \rangle$ such that s_C^{RS} is defined on $\lambda'[0, k']$ we have that the following are equivalent: (a) $\langle k', \lambda', s' \rangle \models [C \text{ acc}]RS$ and (b) $s_C^{RS}(\lambda'[0, k']) = s'_C(\lambda'[0, k'])$.
2. Let $\langle k', \lambda', s' \rangle$ be a profile such that s_C^{RS} is defined on all histories in $\text{out_plays}(\lambda'[0, k'], s'_C)$. Then $\langle k', \lambda', s' \rangle \models [C \text{ perf}]RS$ if and only if s_C^{RS} and s'_C are play-equivalent at $\lambda'[0, k']$.

This establishes a crucial connection between the syntactic notion of performing a rule-based strategy and the semantic notion of a (partial) perfect recall strategy. In the following subsections we use this link to represent positional and uniform strategies up to play-equivalence by rule-based strategies.

The notion of play-equivalence stems from the STIT views in our formalism. From a STIT perspective a strategy is identified by the futures it allows, so two play-equivalent strategies not only *appear* to be same strategy, they *are* the same strategy.

5 Representation Results

5.1 Representing positional strategies

In this subsection, we prove that rule-based strategies can be used to represent positional strategies. For that purpose we introduce a specific type of rule-based strategies:

¹¹If the conditions are not moment-determinate “a choice of an agent, at a given point of a play, may depend on choices other agents can make in the future or in counterfactual plays” (cf. the study on “behavioral strategies” in [13, p. 149]).

Definition 9. A proposition-action strategy RS for a coalition C is a rule-based strategy such that the conditions c and effects e are respectively of the form:

$$c ::= p \mid c \wedge c \mid \neg c \quad e ::= \alpha_C,$$

where p ranges over $Prop$ and α_C ranges over Σ^C .

Because the effects are of the form α_C , there can be at most one way to act according to a proposition-action strategy whenever one of the conditions is triggered. This motivates our definition of *completeness*; a notion that plays a key role in our findings in the correspondence between proposition-action strategies and positional strategies:

Definition 10 (Completeness). We say that a rule-based strategy $RS = \{c_1 \mapsto e_1, \dots, c_N \mapsto e_N\}$ is complete at $\lambda[0, k]$ if and only if $\lambda[0, k] \models \Box G \bigvee_{n \leq N} c_n$, i.e. in any possible future point one of the conditions in S is triggered.¹²

Proposition 2. Let s be a strategy profile, and let $\lambda \in \text{out_plays}(\lambda[0], s)$. Suppose there is a proposition-action strategy RS for C that is complete at $\lambda[0]$ such that $\langle 0, \lambda, s \rangle \models [C \text{ perf}]RS$. Then there is a positional strategy \hat{s}_C that is play-equivalent to s_C at $\lambda[0]$.

This shows that performing a complete proposition-action strategy implies that one is performing a strategy that is play-equivalent to a positional strategy. In a sense, this means that the strategies represented by complete proposition-action strategies are positional strategies.

The converse does not hold in general, which can be shown by providing a CGM containing a positional strategy differing at two propositionally equivalent states. So to prove the converse we have to restrict our investigation to CGMs in which enough states are propositionally definable:

Proposition 3. Let s_C be a positional C -strategy, and let $\lambda \in \text{out_plays}(\lambda[0], s_C)$. Suppose that $s_C(\text{St}) \subseteq \text{Act}^C$ is finite and that every $s_C^{-1}(\alpha_C) \subseteq \text{St}$ is propositionally definable. Then there is a proposition-action strategy RS for C that is complete at $\lambda[0]$ such that $\langle 0, \lambda, s \rangle \models [C \text{ perf}]RS$.

This shows that, under certain semantic constraints, a given positional strategy is represented by a complete proposition-action strategy. Thereby we can move strategic reasoning in the semantics about positional strategies to reasoning on the syntactic level about proposition-action strategies. In conclusion, we show that in a common class of CGMs, complete proposition-action strategies correspond to positional strategies up to play-equivalence:

Corollary 1. Let s be a perfect recall strategy for C , and let $\lambda \in \text{out_plays}(\lambda[0], s_C)$. Let St be finite and let every state be propositionally definable. Then the following are equivalent

1. there is a proposition-action strategy RS for C such that $\langle k, \lambda, s \rangle \models [C \text{ perf}]RS \wedge \Box G \bigvee_{c \in \text{Cond}(RS)} c$,
2. at $\lambda[0]$ the perfect recall strategy s is play-equivalent to a positional strategy.

This corollary uncovers that completeness is an underlying condition for positional strategies. Although this discovery is unsurprising and intuitive, it shows that our language is able to express such underlying intuitions.

5.2 Representing uniform strategies

Here we prove that rule-based strategies are useful for representing uniform strategies, focussing on individuals' uniform strategies, using a type of rule-based strategies:

¹²Compare [17, p. 207]: “a complete strategy provides the player with a “complete guideline” that always provides the player with one or more suggestions how to act when it is his move”.

Definition 11. A knowledge-action strategy RS for a coalition C is a rule-based strategy such that the conditions are of the form $K_i c$ with c and the effects e respectively of the form:

$$c ::= p \mid c \wedge c \mid \neg c \qquad e ::= \alpha_C,$$

where p ranges over Prop and α_C ranges over Σ^C .

Proposition 4. Let $\langle 0, \lambda, s \rangle$ be a tuple. Define $H = \{\lambda'[0, k'] \mid \lambda'[0] \sim_i \lambda[0]\}$. Suppose there is a knowledge-action strategy RS for agent i such that

1. agent i knows that RS is complete at $\lambda[0]$, i.e. $\langle 0, \lambda, s \rangle \models K_i \Box G \bigvee_{K_i c \in \text{Cond}(RS)} K_i c$;
2. agent i knows that she is henceforth able to act according to RS , i.e. $\langle 0, \lambda, s \rangle \models K_i G \Diamond [i \text{ acc}] RS$;
3. agent i performs strategy S , i.e. $\langle 0, \lambda, s \rangle \models [i \text{ perf}] RS$.

Then there is a uniform strategy \hat{s}_i on H that is play-equivalent to s_i .

This result is similar to Proposition 2 on positional strategies. Here we see that whenever an agent performs a knowledge-action strategy of which he knows both that it is complete and that he is henceforth able to act accordingly, the agent is performing a strategy that is play-equivalent to a uniform strategy. This means that, under certain syntactically representable epistemic conditions, a syntactically characterised knowledge-action strategy corresponds with a uniform strategy in the semantic structures.

The converse does not hold, as can be shown by providing a CEGM containing a uniform strategy that differs at two distinguishable propositionally equivalent states. But the mismatch runs deeper because of the restrictions on the conditions of knowledge-action strategies. To obtain a correspondence result in line with Corollary 1 we believe that the language has to be extended with temporal modalities referring to the past and the conditions of knowledge-action strategies have to be modified accordingly. To keep the current exposition accessible we leave this for another occasion. In spite of these simplifications, a representation result for uniform positional strategies can be proven:

Proposition 5. Let $\langle 0, \lambda, s \rangle$ be a profile. Define $H = \{\lambda'[0, k'] \mid \lambda'[0] \sim_i \lambda[0]\}$. Suppose s_i is a uniform positional strategy for agent i on H . Suppose that $s_i(\text{St}) \subseteq \text{Act}$ is finite and that every $s_i^{-1}(\alpha_i)$ is propositionally definable.

Then there is a knowledge-action strategy RS for i such that

1. agent i knows that RS is complete at $\lambda[0]$, i.e. $\langle 0, \lambda, s \rangle \models K_i \Box G \bigvee_{K_i c \in \text{Cond}(RS)} K_i c$;
2. agent i knows that she is henceforth able to act according to RS , i.e. $\langle 0, \lambda, s \rangle \models K_i G \Diamond [i \text{ acc}] RS$;
3. $\langle 0, \lambda, s \rangle \models [i \text{ perf}] RS$.

Note that this proposition starts with a uniform *positional* strategy. This result establishes that, under certain semantic restrictions, a uniform positional strategy is represented by a knowledge-action strategy of which one knows both that it is complete and that one can henceforth act accordingly. This shows that strategic reasoning in the semantics on uniform positional strategies can be diverted to reasoning on the syntactic level about knowledge-action strategies.

Our representation result reveals crucial underlying conditions for uniform positional strategies, which are expressible in our language. Indeed, the previous proposition shows that, under certain model restrictions, performing a uniform positional strategy implies that one is performing a knowledge-action strategy and one knows both that this knowledge-action strategy is complete and that one is henceforth able to act accordingly. Revealing such underlying conditions enhances our understanding and triggers further questions; two of such inquiries are discussed below.

Does ensuring a property φ by performing a uniform strategy *entail* that knowing that performing this uniform strategy ensures that property? According to our representation result, this translates to questioning whether $[i \text{ sstit}]\varphi \wedge [i \text{ perf}]RS \wedge K_i \Box G \bigvee_{K_i c \in \text{Cond}(RS)} K_i c \wedge K_i G \Diamond [i \text{ acc}]RS$ logically entails $K_i \Box ([i \text{ perf}]RS \rightarrow [i \text{ sstit}]\varphi)$ (where RS is a knowledge-action strategy). It turns out that this indeed fails since “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.” [1, p. 574] Uniform strategies are therefore not faithful to the expectation that “the agent has enough control and knowledge to identify and execute a strategy that enforces [a certain property] φ .” [1, p. 574] An agent has this control and knowledge if and only if there is a knowledge-action strategy RS satisfying $K_i \Box G \bigvee_{K_i c \in \text{Cond}(RS)} K_i c \wedge K_i G \Diamond [i \text{ acc}]RS \wedge K_i \Box ([i \text{ perf}]RS \rightarrow [i \text{ sstit}]\varphi)$. This discussion highlights the flexibility of our syntactical approach to correct the flaw of uniform strategies.

What *is* a coalition’s uniform strategy? Formally, it is a tuple of individuals’ uniform strategies; intuitively, it is intended to capture a coalition’s control and knowledge to identify and execute a strategy that enforces a certain property φ . Does a coalition’s uniform strategy meet this intuition? No, it does not. In [1, pp. 575-576] it is argued that “there are several different “modes” in which [a coalition] can know the right strategy”, pointing to a choice between common, mutual, or distributed knowledge of the right coalitional strategy s_C .¹³ Our results, however, solicit a view, complementing the aforementioned modes, to adequately conceptualize a coalition’s uniform strategies: First, because a coalition’s uniform strategy is merely a tuple of individuals’ uniform strategies, a coalition’s uniform strategy s_C does not require that any member can *identify* the coalitional strategy s_C . Indeed, a coalition’s uniform strategy s_C merely requires every member to know their part s_i of the coalitional strategy s_C . So the *object* of the members’ knowledge differs. Second, since every member knows their part s_i of a coalition’s uniform strategy s_C , it follows that a coalition’s uniform strategy entails that the coalition has distributed knowledge of the right coalitional strategy s_C , positioning a coalition’s uniform strategy between the modes of mutual and distributed knowledge of the right coalitional strategy. Third, this knowledge is, however, distributed in a very particular way, namely by every member knowing *their own part* s_i . For instance, whenever i knows j ’s part s_j and j knows i ’s part s_i , then they have distributed knowledge of $\langle s_i, s_j \rangle$ even though it is not a uniform strategy. A coalition’s uniform strategy hence does not correspond to any of the “modes” in [1]. It seems that this distinctive way of distributing knowledge solicits a comparison, not with the proposed modes in [1] but, with different ways of distributing knowledge of a coalitional strategy s_C .

6 Conclusion

We have shown that, under certain model restrictions, a strategy that is play-equivalent to a positional strategy *corresponds* to a complete proposition-action strategy. Thereby we have established a firm correspondence between a semantic strategy type and a syntactic one.

In our research on individuals’ uniform strategies, we have proven that any knowledge-action strategy of which one knows both that one can henceforth act accordingly and that it is complete *represents* a strategy that is play-equivalent to a uniform strategy. Conversely, under certain semantic restrictions, a uniform positional strategy *is represented by* a knowledge-action strategy of which one knows both that one can henceforth act accordingly and that it is complete. This latter result exposes the implicit

¹³They also mention the option that “the strategy s_C can be identified by” (altered notation) a leader, headquarters committee, or consulting company. Our representation result suggests that the syntactical counterpart of these “modes” is straightforward by replacing K_i ’s with the respective group knowledge in Proposition 5. We will not pursue this suggestion in further detail here.

conditions of uniform positional strategies.

The current enterprise is a crucial first step in facilitating strategic reasoning at the syntactic level. By representing several semantic strategy types and drawing novel conceptual implications we have shown the fruitfulness of our syntactic approach to enhance our understanding of semantic strategy types.

References

- [1] Thomas Ågotnes, Valentin Goranko, Wojciech Jamroga & Michael Wooldridge (2015): *Knowledge and Ability*. In: *Handbook of Epistemic Logic*, College Publications, London, pp. 543–589.
- [2] Rajeev Alur, Thomas A. Henzinger & Orna Kupferman (2002): *Alternating-time temporal logic*. *Journal of the ACM (JACM)* 49(5), pp. 672–713, doi:10.1145/585265.585270.
- [3] Nuel Belnap, Michael Perloff & Ming Xu (2001): *Facing the Future. Agents and Choices in Our Indeterminist World*. Oxford University Press.
- [4] Johan van Benthem (2002): *Extensive games as process models*. *Journal of logic, language and information* 11(3), pp. 289–313, doi:10.1023/A:1015534111901.
- [5] Jan Broersen (2009): *A stit-Logic for Extensive Form Group Strategies*. In: *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 03*, IEEE Computer Society, pp. 484–487, doi:10.1109/WI-IAT.2009.331.
- [6] Krishnendu Chatterjee, Thomas A. Henzinger & Nir Piterman (2007): *Strategy logic*. In: *CONCUR 2007-Concurrency Theory*, Springer, pp. 59–73, doi:10.1016/j.ic.2009.07.004.
- [7] Wiebe van der Hoek, Wojciech Jamroga & Michael Wooldridge (2005): *A logic for strategic reasoning*. In: *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, ACM, pp. 157–164, doi:10.1145/1082473.1082497.
- [8] Wiebe van der Hoek & Michael Wooldridge (2003): *Cooperation, Knowledge, and Time: Alternating-time Temporal Epistemic Logic and its Applications*. *Studia Logica* 75(1), pp. 125–157, doi:10.1023/A:1026171312755.
- [9] John F. Horty (2001): *Agency and deontic logic*. Oxford: Oxford University Press, doi:10.1093/0195134613.001.0001.
- [10] Wojciech Jamroga & Wiebe van der Hoek (2004): *Agents that know how to play*. *Fundamenta Informaticae* 63(2-3), pp. 185–220.
- [11] Sophia Knight & Bastien Maubert (2015): *Dealing with imperfect information in Strategy Logic*. In: *Proceedings of the 3rd International Workshop on Strategic Reasoning*, 20.
- [12] Fabio Mogavero (2013): *Reasoning About Strategies*. In: *Logics in Computer Science*, Atlantis Press, pp. 85–116, doi:10.2991/978-94-91216-95-4_4.
- [13] Fabio Mogavero, Aniello Murano & Luigi Sauro (2014): *A behavioral hierarchy of strategy logic*. In: *Computational Logic in Multi-Agent Systems*, Springer, pp. 148–165, doi:10.1007/978-3-319-09764-0_10.
- [14] Rohit Parikh (1985): *The logic of games and its applications*. *North-Holland Mathematics Studies* 102, pp. 111–139, doi:10.1016/S0304-0208(08)73078-0.
- [15] Marc Pauly & Rohit Parikh (2003): *Game logic-an overview*. *Studia Logica* 75(2), pp. 165–182, doi:10.1023/A:1027354826364.
- [16] Ramaswamy Ramanujam & Sunil Easaw Simon (2008): *Dynamic Logic on Games with Structured Strategies*. In: *KR*, pp. 49–58.
- [17] Dongmo Zhang & Michael Thielscher (2014): *Representing and Reasoning about Game Strategies*. *Journal of Philosophical Logic* 44(2), pp. 203–236, doi:10.1007/s10992-014-9334-6.

A Appendix: Proof sketches of propositions

Proposition 1. *Let RS be a rule-based strategy with only moment-determinate conditions. Let s_C^{RS} be the partial coalition strategy defined by RS . Then*

1. *for any profile $\langle k', \lambda', s' \rangle$ such that s_C^{RS} is defined on $\lambda'[0, k']$ we have that the following are equivalent: (a) $\langle k', \lambda', s' \rangle \models [C \text{ acc}]RS$ and (b) $s_C^{RS}(\lambda'[0, k']) = s'_C(\lambda'[0, k'])$.*
2. *Let $\langle k', \lambda', s' \rangle$ be a profile such that s_C^{RS} is defined on all histories in $\text{out_plays}(\lambda'[0, k'], s'_C)$. Then $\langle k', \lambda', s' \rangle \models [C \text{ perf}]RS$ if and only if s_C^{RS} and s'_C are play-equivalent at $\lambda'[0, k']$.*

Proof sketch: 1. Follows from the fact that s_C^{RS} is defined at a history iff there is exactly one way to act according to RS at that history. 2. Follows straightforwardly from 1. and property 5(b) in Definition 1. \square

Proposition 2. *Let s be a strategy profile, and let $\lambda \in \text{out_plays}(\lambda[0], s)$. Suppose there is a proposition-action strategy RS for C that is complete at $\lambda[0]$ such that $\langle 0, \lambda, s \rangle \models [C \text{ perf}]RS$. Then there is a positional strategy \hat{s}_C that is play-equivalent to s_C at $\lambda[0]$.*

Proof sketch: It is easy to show that the partial strategy defined by RS is defined on and positional for histories in $\text{out_plays}(\lambda[0], s_C)$. This partial positional strategy can be trivially extended to a positional strategy, thereby proving the proposition. \square

Proposition 3. *Let s_C be a positional C -strategy, and let $\lambda \in \text{out_plays}(\lambda[0], s_C)$. Suppose that $s_C(\text{St}) \subseteq \text{Act}^C$ is finite and that every $s_C^{-1}(\alpha_C) \subseteq \text{St}$ is propositionally definable. Then there is a proposition-action strategy RS for C that is complete at $\lambda[0]$ such that $\langle 0, \lambda, s \rangle \models [C \text{ perf}]RS$.*

Proof sketch: Let $\{\alpha_C^1, \dots, \alpha_C^N\} = s_C(\text{St})$, and let us denote the propositional formula defining $s_C^{-1}(\alpha_C^n)$ by ξ_n for each $n \leq N$. The proposition-action strategy $RS := \{\xi_1 \mapsto \alpha_C^1, \dots, \xi_N \mapsto \alpha_C^N\}$ can be used to prove the proposition. \square

Proposition 4. *Let $\langle 0, \lambda, s \rangle$ be a tuple. Define $H = \{\lambda'[0, k'] \mid \lambda'[0] \sim_i \lambda[0]\}$. Suppose there is a knowledge-action strategy RS for agent i such that*

1. *agent i knows that RS is complete at $\lambda[0]$, i.e. $\langle 0, \lambda, s \rangle \models K_i \Box G \bigvee_{K_i c \in \text{Cond}(RS)} K_i c$;*
2. *agent i knows that she is henceforth able to act according to RS , i.e. $\langle 0, \lambda, s \rangle \models K_i G \Diamond [i \text{ acc}]RS$;*
3. *agent i performs strategy S , i.e. $\langle 0, \lambda, s \rangle \models [i \text{ perf}]RS$.*

Then there is a uniform strategy \hat{s}_i on H that is play-equivalent to s_i .

Proof sketch: The partial strategy defined by RS is defined and uniform on histories in H . \square

Proposition 5. *Let $\langle 0, \lambda, s \rangle$ be a profile. Define $H = \{\lambda'[0, k'] \mid \lambda'[0] \sim_i \lambda[0]\}$. Suppose s_i is a uniform positional strategy for agent i on H . Suppose that $s_i(\text{St}) \subseteq \text{Act}$ is finite and that every $s_i^{-1}(\alpha_i)$ is propositionally definable.*

Then there is a knowledge-action strategy RS for i such that

1. *agent i knows that RS is complete at $\lambda[0]$, i.e. $\langle 0, \lambda, s \rangle \models K_i \Box G \bigvee_{K_i c \in \text{Cond}(RS)} K_i c$;*
2. *agent i knows that she is henceforth able to act according to RS , i.e. $\langle 0, \lambda, s \rangle \models K_i G \Diamond [i \text{ acc}]RS$;*
3. *$\langle 0, \lambda, s \rangle \models [i \text{ perf}]RS$.*

Proof sketch: Analogous to Proposition 3. \square