



A Logic for Conditional Local Strategic Reasoning

Valentin Goranko^{1,2} · Fengkui Ju³

Accepted: 18 February 2022 / Published online: 25 April 2022
© The Author(s) 2022, corrected publication 2022

Abstract

We consider systems of rational agents who act and interact in pursuit of their individual and collective objectives. We study and formalise the reasoning of an agent, or of an external observer, about the expected choices of action of the other agents based on their objectives, in order to assess the reasoner’s ability, or expectation, to achieve their own objective. To formalize such reasoning we extend Pauly’s Coalition Logic with three new modal operators of conditional strategic reasoning, thus introducing the Logic for Local Conditional Strategic Reasoning ConStR. We provide formal semantics for the new conditional strategic operators in concurrent game models, introduce the matching notion of bisimulation for each of them, prove bisimulation invariance and Hennessy–Milner property for each of them, and discuss and compare briefly their expressiveness. Finally, we also propose systems of axioms for each of the basic operators of ConStR and for the full logic.

Keywords Conditional strategic reasoning · Concurrent games · Coalition Logic · Proactive and reactive abilities · Bisimulations · Expressiveness

1 Introduction

Consider the following scenario. Alice and Bob are students at DownTown University. Alice is coming to campus today, and has some agenda to complete. Bob wants to meet Alice somewhere on campus today. She does not know that (maybe, even does

This paper is a revised and extended version of Goranko and Ju (2019).

✉ Valentin Goranko
valentin.goranko@philosophy.su.se

Fengkui Ju
fengkui.ju@kul.pl

¹ Department of Philosophy, Stockholm University, Stockholm, Sweden

² School of Mathematics, University of the Witwatersrand (visiting professorship), Johannesburg, South Africa

³ Faculty of Philosophy, The John Paul II Catholic University of Lublin, Lublin, Poland

not know Bob) and they have no communication. Bob may, or may not, know what Alice is going to do on campus, or where and at what time she will go during the day. Using his knowledge of what, where, and when Alice intends to do today, Bob wants to come up with a plan of how (where and when) to meet her.

From a more general perspective, we consider a scenario of agents acting independently, and possibly concurrently, in pursuit of their individual and collective goals and we analyse the reasoning of an agent (or, observer) about the possible *local* actions (at the current state only) of the other agents and their effect for realising or enabling the outcome of interest for the reasoner.

Related work and motivation. The kind of strategic reasoning discussed here is within the conceptual thrust motivating the research on logic-based strategic reasoning over the past two decades, starting with Coalition Logic CL (Pauly, 2001, 2002), its temporal extension, the alternating-time temporal logic ATL (Alur et al., 2002), its epistemic extension ATEL (van der Hoek & Wooldridge, 2004), and gradually evolving towards increasingly expressive formalisms, such Strategy Logic SL (Mogavero et al., 2014; cf. 2017). See Ågotnes et al. (2015) and van Benthem et al. (2015) for overviews of the area. Most of these logical systems (except SL where the agents' strategies are explicitly named in the language) assume arbitrary or adversarial behaviour of the agents outside of the proponent coalitions in CL, ATL, and ATEL. Also, the knowledge of the agents involved in ATEL refers to truths (of formulae in the language) at the current state, rather than to their knowledge about each others' objectives and available actions. Thus, these logics formalise *absolute/unconstrained* strategic reasoning—usually, by an external observer—about the unconditional strategic abilities of agents and coalitions to achieve their goals. However, such unconstrained strategic reasoning is seldom applicable in practice, except in purely adversarial zero-sum games. Usually, all agents acting in the system (except for the environment, or an absolute adversary) have their own goals and act in pursuit of their fulfilment, rather than just to prevent the proponents from achieving their goals. This calls for a more refined strategic reasoning, *conditional* on the agents' knowledge of the opponents' goals and possible available actions to achieve them, which is the proposal of the present paper. It should be noted that there is a recent line of research on *rational synthesis* (Fisman et al., 2010; Kupferman et al., 2016) and *rational verification* (Gutierrez et al., 2017; Wooldridge et al., 2016), which does take into account all agents' goals, but aims at designing stable strategy profiles (Nash equilibria) that only guarantee the satisfaction of the goal of one special agent (the proponent, representing the system), whereas all others are supposed to act rationally and accept the proposed solution, whether it satisfies their own goals or not. Thus, our work takes an essentially different perspective and has quite different objectives. We are aware of few other works that deal more directly and explicitly with *conditional* strategic reasoning in a sense akin to the present paper. Besides the earlier, conference version (Goranko & Ju, 2019) of this work, perhaps the closest to it in spirit is the recent (Goranko & Enqvist, 2018), to which the present work relates both conceptually and technically, as well as the conceptually related work (Naumov & Yuan, 2019), which presents a logic which can express statement of the type: “The coalition B has a strategy to achieve their goal once they know the strategy of the coalition A, no matter what the strategy is”. If the epistemic ingredient

in it is considered implicit, this statement is expressible by our operator for ‘reactive strategic ability’ O_β . We also note that an axiomatic system is proposed and proved complete in Naumov and Yuan (2019), which shares some basic axioms with our axiom system for O_β presented in Sect. 6.3, but differs from it on others.

Our contributions. In this work we identify several patterns of conditional strategic reasoning of an observer or an active agent, depending on his/her knowledge about the objectives and possible actions of the other agents. To formalize such reasoning we extend Coalition Logic (Pauly, 2001, 2002) with three new modal operators of conditional strategic reasoning, thus introducing the Logic for Local Conditional Strategic Reasoning ConStR. We provide formal semantics for the new conditional strategic operators, introduce the matching notion of bisimulation for each of them and discuss and compare briefly their expressiveness. We then also propose systems of axioms for each of the basic operators of ConStR and for the whole logic, without stating yet completeness claims (for lack of space, these are left to future work).

Structure of the Paper. Section 2 provides some preliminaries on concurrent game models and the coalition logic CL. Then, Sect. 3 presents an informal discussion on conditional strategic reasoning, motivating the further technical work. Section 4 introduces three modal operators formalising patterns of conditional strategic reasoning and the new logic ConStR as an extension of Coalition Logic with these operators. Section 5 introduces the matching notion of bisimulation for that logic and discuss briefly its expressiveness. In Sect. 6 we propose systems of axioms for each of the basic operators of ConStR and for the full logic. We end with brief concluding remarks in Sect. 7.

2 Preliminaries

Multi-agent game models. We fix a finite set of **agents** $\mathbb{A}gt = \{a_1, \dots, a_n\}$ and a set of **atomic propositions** Π . Subsets of $\mathbb{A}gt$ will also be called **coalitions**.

Definition 1 A **game model**¹ for $\mathbb{A}gt$ and Π is a tuple

$$\mathcal{M} = (S, \{\Sigma_a\}_{a \in \mathbb{A}gt}, g, V)$$

where S is a non-empty set of **states**; each Σ_a is a non-empty set of possible **actions** of agent a ; $V : \Pi \rightarrow \mathcal{P}(S)$ is a **valuation** of the atomic propositions from Π in S ; and g is a **game map** that assigns to each $s \in S$ a strategic game form $g(s) = (\Sigma_{a_1}^s, \dots, \Sigma_{a_n}^s, o_s)$, where each $\Sigma_{a_i}^s \subseteq \Sigma_{a_i}$ is a non-empty set of actions available to player a_i at s , and

$$o_s : \Sigma_{a_1}^s \times \dots \times \Sigma_{a_n}^s \rightarrow S$$

¹ These game models are essentially equivalent to concurrent game models used in Alur et al. (2002).

is a **local outcome function** assigning to any **action profile** $\sigma \in \Sigma_{a_1}^s \times \cdots \times \Sigma_{a_n}^s$ the **outcome state** $o_s(\sigma)$ produced by σ when applied at $s \in S$. The set $\Sigma_{a_1}^s \times \cdots \times \Sigma_{a_n}^s$ of **action profiles available at s** will be denoted by Act_s .

Now, the **global outcome function** in \mathcal{M} is the partial mapping

$$O : S \times \Sigma_{a_1} \times \cdots \times \Sigma_{a_n} \dashrightarrow S$$

defined by $O(s, \sigma) = o_s(\sigma)$, whenever $\sigma \in \text{Act}_s$.

Given a coalition $C \subseteq \text{Agt}$, a **joint action** for C in the model \mathcal{M} is a tuple of individual actions $\sigma_C \in \prod_{a \in C} \Sigma_a$. For any such joint action σ_C that is available at $s \in S$, we define the **set of outcome states from σ_C at s** :

$$\text{Out}[s, \sigma_C] = \{u \in S \mid \exists \sigma \in \text{Act}_s : \sigma|_C = \sigma_C \ \& \ o_s(\sigma) = u\}$$

where $\sigma|_C$ is the restriction of σ to C . Note that the empty tuple σ_\emptyset is the only available joint action for the empty coalition \emptyset at any state.

The basic logic for coalitional strategic reasoning CL. The Coalition Logic CL was introduced in Pauly (2001), cf. also Pauly (2002). CL extends the classical propositional logic with *coalitional strategic modal operators* $[C]$, for any coalition of agents C . The formulae of CL are defined as follows:

$$\varphi := p \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid [C]\varphi$$

We will write $[i]$ instead of $[\{i\}]$. The intuitive reading of $[C]\varphi$ is:

“The coalition C has a joint action that ensures an outcome (state) satisfying φ , regardless of how all other agents act.”

The semantics of CL is defined in terms of the notion of **truth of a CL-formula ψ at a state s of a game model \mathcal{M}** , denoted $\mathcal{M}, s \models \psi$, by induction on formulae, via the key clause:

$$\mathcal{M}, s \models [C]\varphi \Leftrightarrow \text{there exists a joint action } \sigma_C \text{ available at } s, \text{ such that } \mathcal{M}, u \models \varphi \text{ for each } u \in \text{Out}[s, \sigma_C]$$

Thus, $[C]\varphi$ formalises a claim of the ability of the agent/coalition C to choose a suitable (joint) action to ensure achieving the goal φ *regardless of how all other agents choose to act*, and therefore without assuming that the agents in C know the goal(s) of the remaining agents and their available actions to achieve these goals.

The notion of *bisimulation* that guarantees truth invariance of all CL-formulae was first defined in Alur et al. (1998) for the so called ‘alternating transition systems’ (equivalent to a special case of concurrent game models), then independently in Pauly (2001) for the abstract game models defined there, and later in Ågotnes et al. (2007) for concurrent game models, which definition we give here.

Definition 2 (CL-bisimulation) Let $\mathcal{M} = (S, \{\Sigma_a\}_{a \in \text{Agt}}, g, V)$ be a concurrent game model. A binary relation $\mathcal{R} \subseteq S^2$ is a **CL-bisimulation in \mathcal{M}** if it satisfies the following conditions for every pair of states (s_1, s_2) such that $s_1 \mathcal{R} s_2$ and for every coalition C :

Atom equivalence: For every $p \in \Pi$: $s_1 \in V(p)$ iff $s_2 \in V(p)$.

Forth: For every joint action σ_C^1 of C at s_1 , there is a joint action σ_C^2 of C at s_2 such that for every $u_2 \in \text{Out}[s_2, \sigma_C^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_C^1]$ such that $u_1 \mathcal{R} u_2$.

Back: Like **Forth**, but with 1 and 2 swapped.

CL-bisimulation is defined here *within* a model. It readily extends to CL-bisimulation *between* models, by treating both as parts of a single model.

*The Alternating-time Temporal Logic ATL**. The Alternating-time Temporal Logic ATL*, proposed by Alur et al. (2002), is an extension of CL with temporal operators. Its featured operator is $\langle\langle C \rangle\rangle\phi$, denoting the claim that C has a joint strategy that guarantees the satisfaction of ϕ , where ϕ is a ‘temporal objective’, i.e. a (path) formula beginning with a temporal operator ‘nexttime’ X, ‘always’ G, or ‘until’ U. The logic CL embeds into ATL* as the fragment extending propositional logic only with combinations of strategic and temporal operators of the type $\langle\langle C \rangle\rangle X$, cf. Goranko (2001). We only mention ATL* here for the sake of some further references, but the present paper will not make any essential use of that logic, and no familiarity with it, nor even with its fragment ATL, is required.

3 Conditional Strategic Reasoning: An Informal Discussion

²Recall the scenario with the students Alice and Bob. Suppose that Alice has an objective γ_A to achieve—say, to meet with her supervisor Carl on campus today. Suppose also that Alice has several possible choices of an action (or a ‘strategy’) ³ that would possibly, or certainly, guarantee the achievement of her objective. In our example, suppose these choices are: meeting in the supervisor’s office, or in the library, or at the campus café.

3.1 Observer’s Conditional Reasoning About An Outcome from Agent’s Actions

Let us first consider the case where Bob is just an observer who is not acting, but only reasoning about the consequences from Alice’s possible actions with respect to the occurrence of another—intended or not—outcome event γ_B . For instance, suppose that Bob is interested in meeting with Alice on campus today—let us call that event γ_B —and is sitting in the campus café and reasoning about whether Alice will happen to come to the café, thus enabling the event γ_B (recall that Alice may not know about bob expecting her there, or at all). More generally, we can also assume that there are other agents, besides Alice, also acting in pursuit of their own goals, and Bob is reasoning about their individual and collective choices of action and the consequences from these choices. This leads to an *observer’s conditional strategic reasoning* about claims of the type:

² The reader who is only interested in the technical part of the paper can skip this section without essential loss.

³ In this paper we focus on local reasoning, about once-off actions, but in this section the word ‘action’ can be conceived in a wider sense, and may mean either a once-off action, or a global strategy guiding the long term behaviour of the agent.

“Some/every action of Alice that guarantees achievement of γ_A also guarantees/enables occurrence of the outcome γ_B ”.

Depending on Bob’s knowledge about Alice’s objective and of her expected choices of action there can be several possible cases for Bob’s reasoning about the expected occurrence of the outcome γ_B .

3.1.1 Observer Bob’s Reasoning, Case 1: Bob Knows Nothing About Alice

Suppose that *Bob does not know Alice’s objective* γ_A , and therefore has no a priori expectations about her choice of action. In our example, suppose that Bob only knows that Alice is coming to campus today, but not why and where on campus she is going. Then, Bob can *only* claim for sure that the outcome γ_B will occur if γ_B is inevitable, regardless of how Alice (and all others) will act. For instance, if Bob knows that Alice is coming to campus and he is standing by the only entrance of the campus, then he will know for sure that he is going to meet Alice (γ_B will occur), no matter what she will do there. This claim can be expressed in Coalition Logic CL simply as $[\emptyset]\gamma_B$.

3.1.2 Observer Bob’s Reasoning, Case 2: Bob Only Know Alice’s Goal

Suppose now that Bob *does know Alice’s objective* and knows that Alice can guarantee the achievement of that objective and will act towards that, but Bob does not know *how exactly* Alice might act. E.g., Bob knows that Alice is coming to campus to meet with her supervisor Carl but does not know where and when. Then, Bob can only claim that the outcome γ_B will occur for sure if γ_B is true *on every possible course of events* (“play”) *on which γ_A is true*. For instance, if Bob knows that Alice’s supervisor will be working in his office for the whole day, and he is sitting in the corridor, next to Carl’s office, then he knows that he will meet with Alice (γ_B will occur) no matter when Alice comes to meet with Carl (i.e. no matter how γ_A occurs).

This can be expressed in CL simply as $[\emptyset](\gamma_A \rightarrow \gamma_B)$ and reflects the case when γ_A can occur in various, possibly unintended ways, but its occurrence always implies the occurrence of γ_B (e.g. if Bob is with Carl throughout the day, then even if Alice bumps into Carl accidentally, Bob will still meet her).

3.1.3 Observer’s Reasoning, Case 3: Bob Knows Alice’s Goal and Actions

Suppose now that Bob not only knows Alice’s objective, but also *knows all possible actions (or, strategies)* of Alice that can ensure the satisfaction of her objective γ_A , and *knows that Alice will perform one of them*, but *does not know to which one*. (E.g., Bob knows that Alice, who is coming to campus to meet with her supervisor, can meet with him either in his office, or in the library, or in the café.)

Now, for Bob to claim that the outcome γ_B will occur for sure, it suffices to know that *each action of Alice that guarantees γ_A* will also guarantee γ_B . (E.g., suppose that all possible meeting places for Alice and her supervisor are in the main building and Bob is waiting at the only entrance of the main building.)

Here the conditional “If γ_A then γ_B ” has a suitably constrained context, specifying that γ_A can occur only because the agent (Alice) takes a deliberate action to bring about

γ_A . This can no longer be expressed in CL and requires introducing a new strategic operator.

Lastly, suppose that *Bob also knows the specific action which Alice is taking in order to guarantee the achievement of her goal*. Then, Bob can claim that the outcome γ_B will occur for sure, as long as *that specific action* of Alice guarantees the satisfaction of γ_B . To formalise that one needs explicit names for actions, but in our logic we will be able to state something stronger, viz. that *every specific action of Alice that guarantees γ_A will also bring about γ_B* .

3.2 Conditional Reasoning of an Agent About Another Agent's Actions

Suppose now that Bob is not just a passive observer, but an acting agent, who has the outcome γ_B as his own goal. There may be other agents, besides Alice and Bob, also acting in pursuit of their own goals, and Bob is reasoning about their expected choices of action and the consequences from these choices. Now, Bob is to decide—based on his reasoning about Alice's (and other agents') possible choices of actions—on his own action in pursuit of γ_B . This calls for an *agent's conditional strategic reasoning* about statements of the type:

“For some/every action of Alice that guarantees achievement of γ_A , Bob has an action of his own to guarantee achievement of his objective γ_B ”.

We call this *local conditional strategic reasoning*, as it only refers to the immediate actions of the agents, not about their *global strategies*. Respectively, the outcomes from the local action profiles are just successor states, while in the general case they are (finite or possibly infinite) *plays*. The global conditional strategic reasoning will be treated in a follow-up work.

Each of the cases considered in Sect. 3.1 accordingly applies here, too, with the only difference being that now Bob is to choose a suitable action of his own. Besides, there are several additional cases to consider regarding the possible choice of action of Bob.

3.2.1 Agent Bob's Reasoning, Case 1a: Assuming Alice's Cooperation

First, suppose, in addition, that Alice also knows Bob's objective and can choose to cooperate with Bob by selecting a suitable action σ_A that would not only guarantee achievement of her objective but will also enable Bob to supplement σ_A with an action σ_B which would then also guarantee achievement of his objective, too. (So, we also assume that Alice knows enough about Bob's possible actions.) We refer the reader to Example 1 for a formal model illustrating the agent's ability assuming cooperation from the other agent.

3.2.2 Agent Bob's Reasoning, Case 1b: Not Assuming Alice's Cooperation

Now, suppose Bob cannot count on Alice's cooperation. Still, the statement

“Whichever way Alice acts towards achieving the objective γ_A , Bob can act so as to bring about achievement of his objective γ_B .”

admits two different readings, which we respectively call ‘*proactive ability*’ and ‘*reactive ability*’ which we discuss below.⁴

3.2.3 Agent Bob’s Reasoning, Case 2: Reactive Ability

In this case, *for every action* of Alice that ensures γ_A Bob is to choose *reactively* an action of his, *generally dependent on Alice’s action*, that would also ensure the occurrence of γ_B (possibly in different ways for the different actions). This essentially corresponds to the case where Bob knows Alice’s action at the time of choosing his own action, and for every choice of action of Alice that guarantees γ_A , Bob’s respective choice will also bring about γ_B . For instance, in our running example, if Bob knows where Alice is going to meet with her supervisor, he can choose respectively where to go and wait for her.

More formally, each of Alice’s actions that would guarantee γ_A generates a set of *possible* outcome states (plays), and for each such set Bob is looking for a respective action that will bring about γ_B on that set of outcome states.

3.2.4 Agent Bob’s Reasoning, Case 3: Proactive Ability

The case of *proactive ability* is when Bob only knows that Alice has committed to act so as to achieve her goal γ_A (to meet with Carl at one of the 3 possible meeting places), but does not know the action that Alice has chosen and her choice will remain unknown to Bob at the time when he is to choose his action aiming at satisfying γ_B (meet Alice).

In this case Bob must consider all possible courses of events (plays) that can occur as a result of Alice acting towards achieving γ_A and reason about whether he can choose *proactively and uniformly* one action that would bring about γ_B regardless of which action Alice may choose to apply in order to achieve her goal γ_A . For instance, in our running story, assuming that all meeting places are in the main building, Bob can choose to wait for Alice at the only entrance of that building. Formally speaking, in this case, based on his knowledge Bob considers the set of states in the model which is the union of all sets of outcome states enabled by the specific actions of Alice that would guarantee γ_A , and is looking for an action that will bring about γ_B on each of these outcome states.

We refer the reader to Example 2 for a formal model illustrating the concepts of proactive and reactive abilities of an agent and their difference.

Top sum up, the *proactive–reactive* ability distinction applies to Bob depending on whether or not he knows Alice’s choice at the time when he is to make his own choice of action. If he knows Alice’s choice at that time, his reasoning is about reactive ability, else it is a proactive ability reasoning. We note that the notions of proactive and reactive ability respectively correspond to the notions of α -*effectivity* and β -*effectivity* in game theory (cf. e.g. Abdou & Keiding, 1991).

Lastly, an important point: even though knowledge of the agent about the other’s goals and possible actions is essential, it will not feature in our formal logical language,

⁴ In Goranko and Ju (2019) these were called respectively ‘*ability de re*’ and ‘*ability de dicto*’.

nor in the formal semantics, but only in the external reasoners' analysis of which case of conditional strategic ability applies.

4 The Logic of Conditional Strategic Reasoning ConStR

4.1 Modal Operators for Conditional Strategic Reasoning

Given coalitions A and B and joint actions σ_A for A and σ_B for B , we say that σ_B is **consistent with** σ_A if σ_B coincides with σ_A on $A \cap B$.

We now introduce new operators for conditional strategic reasoning, for any coalitions A and B , with intuitive semantics corresponding to the three reasoning cases in Sect. 3.2, as follows.

$(O_c) \langle A \rangle_c(\phi; \langle B \rangle \psi)$ says that A has a joint action σ_A which, when applied, guarantees the truth of ϕ and enables B to apply a joint action σ_B that is consistent with σ_A and guarantees ψ when *additionally* applied by B , in sense that all agents in A act according to σ_A and those in $B \setminus A$ act according to σ_B .⁵ This operator formalises the agent's reasoning Case 1 discussed in Sect. 3.2, where A knows the objective of B and can choose to cooperate with B by selecting a suitable action.

$(O_\alpha) [A]_\alpha(\phi; \langle B \rangle \psi)$ says that the coalition $B \setminus A$ has an action $\sigma_{B \setminus A}$ such that if A applies any action that guarantees the truth of ϕ , then $B \setminus A$ can guarantee the truth of ψ by applying additionally the action $\sigma_{B \setminus A}$.

This operator formalises a claim of the ability of the agent/coalition B to choose a suitable (joint) action so as to achieve the goal ψ assuming that A acts so as to achieve the goal ϕ , if B is to choose their (joint) action *before* A chooses their (joint) action, or before B learns the action of A . This corresponds to the notion of agent's *proactive ability* discussed in Sect. 3.2.4, respectively to the game-theoretic notion of α -*effectivity*, hence the notation.

$(O_\beta) [A]_\beta(\phi; \langle B \rangle \psi)$ says that for any joint action σ_A of A that guarantees the truth of ϕ , when applied by A there is an action σ_B that is consistent with σ_A and guarantees ψ when additionally applied by B .⁶ This operator formalises a claim of the ability of the agent/coalition B to choose a suitable (joint) action so as to achieve the goal ψ assuming that A acts so as to achieve the goal ϕ , if B is to choose their (joint) action *after* B learns the (joint) action of A . This corresponds to the notion of agent's *reactive ability* discussed in Sect. 3.2.3, respectively to the game-theoretic notion of β -*effectivity*, hence the notation.

⁵ We note that $\langle A \rangle_c(\phi; \langle B \rangle \psi)$ is equivalent to $\langle\langle A \rangle\rangle(\phi \wedge \langle\langle B \setminus A \rangle\rangle \psi)$ in ATL*.

⁶ Note that $[A]_\beta(\perp; \langle B \rangle \psi)$ is vacuously true for any A, B , and ψ , as then there cannot be such joint actions σ_A that enable satisfying \perp . This may sound odd, but it is no special phenomenon in ConStR, as the same effect occurs in FOL with universal quantification over an empty set of objects.

4.2 The Language of ConStR

We fix a finite nonempty set of agents $\mathbb{A}gt$ and a countable set of atomic propositions Π . The formulae of ConStR, where $p \in \Pi$ and $A, B \subseteq \mathbb{A}gt$ are defined as follows:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \langle A \rangle_c(\phi; \langle B \rangle\phi) \mid [A]_\alpha(\phi; \langle B \rangle\phi) \mid [A]_\beta(\phi; \langle B \rangle\phi)$$

4.3 Some Definable Operators and Expressions in ConStR

The following can be easily seen from the informal semantics above, and can also be easily verified with the formal semantics introduced further.

- The coalitional operator from CL is definable as a special case of each of O_c , O_α , O_β as follows:

$$(O_c) [A]\phi := \langle A \rangle_c(\phi; \langle A \rangle\phi) \text{ or } [A]\phi := \langle A \rangle_c(\phi; \langle A \rangle\top);$$

$(O_\alpha) [A]\phi := [\emptyset]_\alpha(\top; \langle A \rangle\phi)$. This expresses the case when B is not informed about the goal of A and has to choose proactively a joint action, before A has chosen their action. Thus, it indeed claims an unconditional ability of B to choose an action that guarantees ϕ .

$(O_\beta) [A]\phi := [\emptyset]_\beta(\top; \langle A \rangle\phi)$, or $[A]\phi := [\bar{A}]_\beta(\top; \langle A \rangle\phi)$; (the empty coalition has only one strategy, and it guarantees the satisfaction of \top)

Thus, the cases of observer's reasoning discussed in Sects. 3.1.1 and 3.1.2 are readily formalisable in ConStR.

- The dual to O_c operator $\neg\langle A \rangle_c(\phi; \langle B \rangle\neg\psi)$ says that every joint action of A that, when applied, guarantees the truth of ϕ , would prevent B from acting additionally so as to guarantee ψ . This formalises the conditional reasoning scenario where the goals of A and B are conflicting and where Bob can establish that whichever way A acts towards their goal, that would block B from acting to guarantee achievement of its goal.
- $[A]_\beta(\top; \langle B \rangle\psi)$ essentially formalises the case when B is not informed about the goal of A, but has to choose their action after learning the action of A.
- On the other hand, $[A]_c(\phi|\psi) := [A]_\beta(\phi; \langle \emptyset \rangle\psi)$, also equivalent to $[A]_\alpha(\phi; \langle \emptyset \rangle\psi)$, says that for any joint strategy of A, if it guarantees ϕ to be true, then it guarantees ψ to be true, too. That formalises the case in Sect. 3.1.3 of reasoning of an observer who knows both the goal ϕ and the possible actions of A, about the occurrence of outcome ψ .
- $\langle A \rangle_c(\phi|\psi) := \neg[A]_c(\phi|\neg\psi)$ says that there is a joint strategy of A that guarantees ϕ to be true and enables ψ to be true, too. Note that it is equivalent to a special case of the “socially friendly coalitional operator” SF, $[C](\phi; \psi_1, \dots, \psi_k)$, introduced in Goranko and Enqvist (2018), viz. $\langle A \rangle_c(\phi|\psi) \equiv [A](\phi; \psi)$. Moreover, $\langle A \rangle_c(\phi|\psi)$ is also definable as $\langle A \rangle_c(\phi; \langle \bar{A} \rangle\psi)$, where $\bar{A} = \mathbb{A}gt \setminus A$.
- The coalitional operator $[A]$ from CL is a special case of the above: $[A]\phi := \langle A \rangle_c(\phi|\top)$, meaning “A has a joint action to ensure the truth of ϕ ”.⁷

⁷ NB: We have preserved the box-like notation for $[A]$ from CL, even though it is not consistent with ours.

- $\langle A \rangle_c(\phi; \langle B \rangle \psi)$ is definable in terms of the “group protecting coalitional operator” GP, introduced in Goranko and Enqvist (2018): $\langle A \rangle_c(\phi; \langle B \rangle \psi) \equiv \llbracket A \triangleright \phi, A \cup B \triangleright \psi \rrbracket$. Nevertheless, it now has a different motivation and intuitive interpretation.

4.4 Formal Semantics of ConStR

Given coalitions $A, B \subseteq \mathbb{A}gt$ and joint actions σ_A for A and σ_B for B , we define $\sigma_A \uplus \sigma_B$ to be the joint action for $A \cup B$ which equals to σ_A when restricted to A and equals to $\sigma_B|_{B \setminus A}$ when restricted to $B \setminus A$. Thus, in particular, $\sigma_A \uplus \sigma_B = \sigma_A$ for any $B \subseteq A \subseteq \mathbb{A}gt$.

Now, let $\mathcal{M} = (S, \{\Sigma_a\}_{a \in \mathbb{A}gt}, g, V)$ be a game model. The formal semantics of ConStR extends the one of CL to the new operators as follows:


- $\mathcal{M}, s \Vdash \langle A \rangle_c(\phi; \langle B \rangle \psi) \Leftrightarrow A$ has a joint action σ_A , such that $\mathcal{M}, u \Vdash \phi$ for every $u \in \text{Out}[s, \sigma_A]$ and B has a joint action σ_B such that $\mathcal{M}, u \Vdash \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_B]$.
- $\mathcal{M}, s \Vdash [A]_\alpha(\phi; \langle B \rangle \psi) \Leftrightarrow B$ has a joint action σ_B such that for every joint action σ_A of A , if $\mathcal{M}, u \Vdash \phi$ for every $u \in \text{Out}[s, \sigma_A]$, then $\mathcal{M}, u \Vdash \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_B]$.
- $\mathcal{M}, s \Vdash [A]_\beta(\phi; \langle B \rangle \psi) \Leftrightarrow$ for every joint action σ_A of A such that $\mathcal{M}, u \Vdash \phi$ for every $u \in \text{Out}[s, \sigma_A]$, B has a joint action σ_B (generally, dependent on σ_A) such that $\mathcal{M}, u \Vdash \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_B]$.

Remark The semantics of each of the operators above can be re-stated to consider joint actions for $B \setminus A$ rather than the whole B . For instance, it can be easily verified for the latter operator, that $\mathcal{M}, s \Vdash [A]_\alpha(\phi; \langle B \rangle \psi)$ iff $B \setminus A$ has a joint action $\sigma_{B \setminus A}$ such that for every joint action σ_A of A , if $\mathcal{M}, u \Vdash \phi$ for every $u \in \text{Out}[s, \sigma_A]$, then $\mathcal{M}, u \Vdash \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_{B \setminus A}]$.

4.5 Some Examples


Here we provide a few simple examples illustrating the semantics of ConStR.

Example 1 The game model \mathcal{M} below has two players, a and b . Each has two actions at state s_0 : a_1, a_2 , resp. b_1, b_2 .



It is easy to see that $\mathcal{M}, s_0 \Vdash \langle a \rangle_c(p; \langle b \rangle q)$, while $\mathcal{M}, s_0 \not\Vdash [b]q$. Thus, an agent may have only conditional ability to achieve its goal.

Example 2 The game model \mathcal{M} below has two players, a and b. The agent a has 3 actions at state s_0 , a_1, a_2, a_3 , and b has 2 actions, b_1, b_2 .



Note that:

- $\mathcal{M}, s_0 \models [a]_\beta(p; \langle b \rangle q)$. Indeed, agent a has two actions at state s_0 to ensure p : a_1 and a_2 . For each of them, b has an action to ensure q , viz.: choose b_2 if a chooses a_1 , and choose b_1 if a chooses a_2 .
 - $\mathcal{M}, s_0 \not\models [a]_\alpha(p; \langle b \rangle q)$. Indeed, neither b_1 nor b_2 ensures q against both choices a_1 and a_2 of a. Thus, b does not have a uniform action to ensure q against any action of a that ensures p .
- Therefore, $[a]_\alpha(p; \langle b \rangle q)$ and $[a]_\beta(p; \langle b \rangle q)$ are semantically different.
- However, if the outcomes of (a_2, b_1) and (a_2, b_2) are swapped, then $[a]_\alpha(p; \langle b \rangle q)$ becomes true at s_0 in the resulting model.
 - (\mathcal{M}, s_0) does not satisfy the ATL* formula $\llbracket a \rrbracket (Xp \rightarrow \langle\langle b \rangle\rangle Xq)$ (where $\llbracket C \rrbracket \phi := \neg \langle\langle C \rangle\rangle \neg \phi$), hence it is not equivalent to $[a]_\beta(p; \langle b \rangle q)$.

5 Bisimulations and Expressiveness of ConStR

5.1 Bisimulations for ConStR

The definition of ConStR-bisimulation involves, besides atomic equivalence, 3 nested Forth and Back conditions, for each of the respective new operators $O_c, O_\alpha,$ and O_β .⁸ As the definition of CL-bisimulation given in Sect. 2, we only define ConStR-bisimulation within a game model, which generalises to ConStR-bisimulation between game models easily. Note that the nested back-and-forth conditions are needed because of the patterns of quantification in the semantic definitions of the new strategic operators: first quantification over the actions of A and B, and then over the outcomes generated by these actions.

Definition 3 (ConStR-bisimulation) Let $\mathcal{M} = (S, \{\Sigma_a\}_{a \in \text{Ag}}, g, V)$ be a game model. A binary relation $\mathcal{R} \subseteq S^2$ is a **ConStR-bisimulation in \mathcal{M}** if it satisfies the following

⁸ Each of these conditions is a respective variation of the bisimulation conditions for the basic strategic operators in the logics SFCL and GPCL defined in Goranko and Enqvist (2018).

conditions for every pair of states (s_1, s_2) such that $s_1 \mathcal{R} s_2$ and for every coalitions A and B:

Atom equivalence: For every $p \in \Pi$: $s_1 \in V(p)$ iff $s_2 \in V(p)$.

O_c -bisimulation: (For illustration, see Fig. 1)

A-Forth_c: For every joint action σ_A^1 of A at s_1 there is a joint action σ_A^2 of A at s_2 , such that:

A-LocalBack_c: For every $u_2 \in \text{Out}[s_2, \sigma_A^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1]$ such that $u_1 \mathcal{R} u_2$.

B-Forth_c: For every joint action σ_B^1 of B at s_1 there is a joint action σ_B^2 of B at s_2 , such that:

(A \uplus B)-LocalBack_c: For every $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$ such that $u_1 \mathcal{R} u_2$.

A-Back_c: Like A-Forth_c, but with 1 and 2 swapped.

O_α -bisimulation: (For illustration, see Fig. 2.)

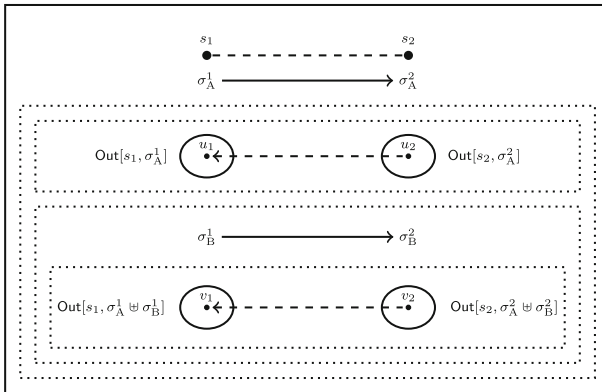


Fig. 1 The A-Forth_c half of O_c -bisimulation

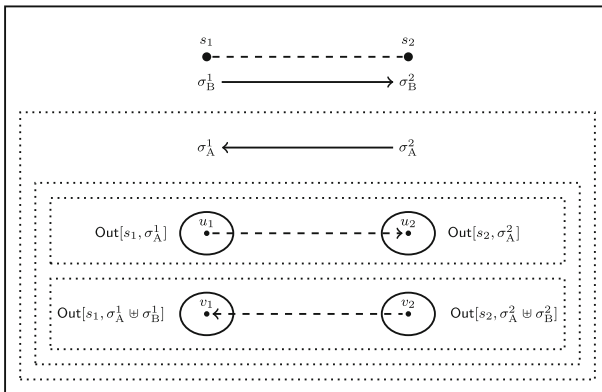


Fig. 2 The B-Forth_α half of O_α -bisimulation

B-Forth $_{\alpha}$: For every joint action σ_B^1 of B at s_1 there is a joint action σ_B^2 of B at s_2 , such that:

A-Back $_{\alpha}$: For every joint action σ_A^2 of A at s_2 there is a joint action σ_A^1 of A at s_1 , such that:

(A)-**LocalForth $_{\alpha}$** : For every $u_1 \in \text{Out}[s_1, \sigma_A^1]$ there exists $u_2 \in \text{Out}[s_2, \sigma_A^2]$ such that $u_1 \mathcal{R} u_2$.

(A \uplus B)-**LocalBack $_{\alpha}$** : For every $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$ such that $u_1 \mathcal{R} u_2$.

B-Back $_{\alpha}$: Like **B-Forth**, but with 1 and 2 swapped.

O $_{\beta}$ -bisimulation: (For illustration, see Fig. 3.)

A-Forth $_{\beta}$: For every joint action σ_A^1 of A at s_1 there is a joint action σ_A^2 of A at s_2 , such that:

A-LocalBack $_{\beta}$: For every $u_2 \in \text{Out}[s_2, \sigma_A^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1]$ such that $u_1 \mathcal{R} u_2$.

B-Back $_{\beta}$: For every joint action σ_B^2 of B at s_2 there is a joint action σ_B^1 of B at s_1 , such that:

(A \uplus B)-**LocalForth $_{\beta}$** : For every $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$ there exists $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$ such that $u_1 \mathcal{R} u_2$.

A-Back $_{\beta}$: Like **A-Forth**, but with 1 and 2 swapped.

States $s_1, s_2 \in \mathcal{M}$ are **ConStR-bisimulation equivalent**, or just **ConStR-bisimilar** if there is a ConStR-bisimulation \mathcal{R} in \mathcal{M} such that $s_1 \mathcal{R} s_2$.

Proposition 1 (ConStR-bisimulation invariance) *Let \mathcal{R} be a ConStR-bisimulation in a game model \mathcal{M} . Then for every ConStR-formula θ and a pair $s_1, s_2 \in \mathcal{M}$ such that $s_1 \mathcal{R} s_2$: $\mathcal{M}, s_1 \models \theta$ iff $\mathcal{M}, s_2 \models \theta$.*

Proof Induction on θ . All boolean cases are straightforward. The cases for the 3 strategic operators are similar, but we will nevertheless check each of them, to ensure that the bisimulation conditions above are correctly defined.

(Case O_c) Let $\theta = \langle A \rangle_c(\phi; \langle B \rangle \psi)$, assuming that the claim holds for ϕ and ψ .

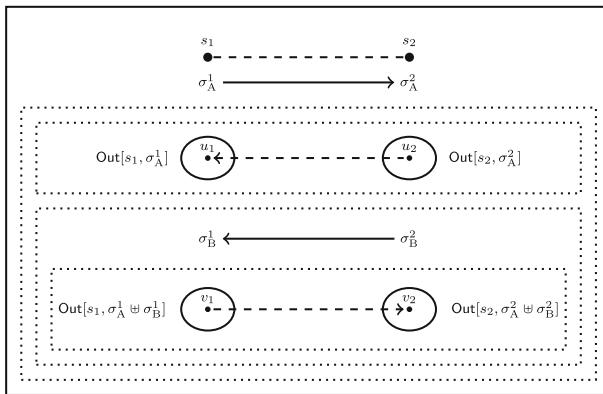


Fig. 3 The A-Forth $_{\beta}$ half of O $_{\beta}$ -bisimulation

Suppose, $\mathcal{M}, s_1 \models \theta$. Then A has a joint action σ_A^1 at s_1 such that, when applied, it guarantees ϕ and enables B to adopt a joint action σ_B that is consistent with σ_A and guarantees ψ when additionally applied by B. By **A-Forth_c**, there is a joint action σ_A^2 of A at s_2 , such that, by **A-LocalBack_c**, for each $u_2 \in \text{Out}[s_2, \sigma_A^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1]$ such that $u_1 \mathcal{R}u_2$. By the choice of σ_A^1 , $\mathcal{M}, u_1 \models \phi$ for each $u_1 \in \text{Out}[s_1, \sigma_A^1]$. It follows, by the inductive hypothesis applied to ϕ , that $\mathcal{M}, u_2 \models \phi$ for each $u_2 \in \text{Out}[s_2, \sigma_A^2]$. Moreover, B has a joint action σ_B^1 at s_1 such that, when applied by B, in addition to A applying σ_A^1 , it guarantees ψ , i.e. $\mathcal{M}, u_1 \models \psi$ for each $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$. By condition **B-Forth_c**, there is a joint action σ_B^2 of B at s_2 , such that, by **(A \uplus B)-LocalBack_c**, for every $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$ such that $u_1 \mathcal{R}u_2$. Therefore, by the inductive hypothesis applied to ψ , $\mathcal{M}, u_2 \models \psi$ for each $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$. Thus, $\mathcal{M}, s_2 \models \theta$.

The converse is similar, using **A-Back_c**.

(Case O_α) Let $\theta = [A]_\alpha(\phi; \langle B \rangle \psi)$, assuming the claim holds for ϕ and ψ .

Suppose, $\mathcal{M}, s_1 \models \theta$. Let σ_B^1 be a joint action of B at s_1 satisfying the truth condition of θ . By **B-Forth_α**, there is a joint action σ_B^2 of B at s_2 , such that **A-Back_α** holds. Now, take any joint action σ_A^2 of A at s_2 such that $\mathcal{M}, u_2 \models \phi$ for each $u_2 \in \text{Out}[s_2, \sigma_A^2]$. Then, by **A-Back_α**, there is a joint action σ_A^1 of A at s_1 such that, by **(A)-LocalForth_α**, for every $u_1 \in \text{Out}[s_1, \sigma_A^1]$ there exists $u_2 \in \text{Out}[s_2, \sigma_A^2]$ such that $u_1 \mathcal{R}u_2$. Then, by the inductive hypothesis applied to ϕ , it follows that $\mathcal{M}, u_1 \models \phi$ for each $u_1 \in \text{Out}[s_1, \sigma_A^1]$. By the choice of σ_B^1 , this implies $\mathcal{M}, u_1 \models \psi$ for each $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$. By **(A \uplus B)-LocalBack_α**, for every $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$ such that $u_1 \mathcal{R}u_2$. Therefore, by the inductive hypothesis applied to ψ , we have $\mathcal{M}, u_2 \models \psi$ for each $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$. Thus, σ_B^2 satisfies the truth condition of θ at s_2 . Hence, $\mathcal{M}, s_2 \models \theta$.

The converse direction is analogous, using **B-Back_α**.

(Case O_β) Let $\theta = [A]_\beta(\phi; \langle B \rangle \psi)$, assuming the claim holds for ϕ and ψ .

Suppose, $\mathcal{M}, s_2 \models \theta$. Then, consider any joint action σ_A^1 of A at s_1 such that, when applied by A, it guarantees ϕ . (If no such joint action exist at s_1 , then $\mathcal{M}, s_1 \models \theta$ is vacuously true.) Then, by **A-Forth_β**, there is a joint action σ_A^2 of A at s_2 , such that, by **A-LocalBack_β**, for every $u_2 \in \text{Out}[s_2, \sigma_A^2]$ there exists $u_1 \in \text{Out}[s_1, \sigma_A^1]$ such that $u_1 \mathcal{R}u_2$. Then, by the inductive hypothesis applied to ϕ , it follows that $\mathcal{M}, u_2 \models \phi$ for each $u_2 \in \text{Out}[s_2, \sigma_A^2]$. Therefore, by the assumption for the truth of θ at s_2 , B has a joint action σ_B^2 at s_2 such that, when applied by B, in addition to A applying σ_A^2 , it guarantees the truth of ψ , i.e. $\mathcal{M}, u_2 \models \psi$ for each $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$. Then, by **B-Back_β**, there is a joint action σ_B^1 of B at s_1 , such that, by **(A \uplus B)-LocalForth_β**, for every $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$ there exists $u_2 \in \text{Out}[s_2, \sigma_A^2 \uplus \sigma_B^2]$ such that $u_1 \mathcal{R}u_2$. Therefore, $\mathcal{M}, u_1 \models \psi$ for each $u_1 \in \text{Out}[s_1, \sigma_A^1 \uplus \sigma_B^1]$. Thus, $\mathcal{M}, s_1 \models \theta$.

The converse is analogous, using **A-Back_β**. □

We also obtain the Hennessy–Milner property for ConStR-bisimulations:

Proposition 2 (Hennessy–Milner property) *For any finite game model \mathcal{M} the relation of ConStR-equivalence (satisfaction of the same ConStR-formulae) between states in \mathcal{M} is a ConStR-bisimulation in \mathcal{M} .*

Proof (Sketch) One direction follows from Proposition 1. We now prove the converse. Since \mathcal{M} is finite, there is a mapping χ from \mathcal{M} to the formulae of ConStR

that assigns to each state s in \mathcal{M} its characteristic formula $\chi(s)$, such that s_1, s_2 are ConStR-equivalent if and only if s_1 satisfies $\chi(s_2)$ (and vice versa), iff $\chi(s_1) \equiv \chi(s_2)$. Furthermore, $\chi(s_1) \wedge \chi(s_2) \equiv \perp$ whenever s_1 and s_2 are not ConStR-equivalent. Now, for any set of states Z in \mathcal{M} we define $\chi(Z) := \bigvee_{z \in Z} \chi(z)$. The crucial observation for proving the claim is that every state $s \in \mathcal{M}$ satisfies each of the following formulae, enabling the verification of the respective ConStR-bisimulation conditions:

- (1) $\bigwedge_{A, B \subseteq \text{Act}_s} \{ \langle A \rangle_c(\chi(Z)); \langle B \rangle \chi(Y) \mid \exists \sigma \in \text{Act}_s : \text{Out}[s, \sigma|_A] = Z \text{ and } \text{Out}[s, \sigma|_{(A \cup B)}] = Y \}$
- (2) $\bigwedge_{A, B \subseteq \text{Act}_s} \{ [A]_\alpha(\chi(Z)); \langle B \rangle \chi(Y) \mid \exists \sigma \in \text{Act}_s : \forall \sigma' \in \text{Act}_s \text{ if } \text{Out}[s, \sigma'|_A] \subseteq Z \text{ and } \sigma'|_{(B \setminus A)} = \sigma|_{(B \setminus A)} \text{ then } \text{Out}[s, \sigma'|_{(A \cup B)}] \subseteq Y \}$
- (3) $\bigwedge_{A, B \subseteq \text{Act}_s} \{ [A]_\beta(\chi(Z)); \langle B \rangle \chi(Y) \mid \forall \sigma \in \text{Act}_s : \text{Out}[s, \sigma|_A] \subseteq Z \text{ implies } \text{Out}[s, \sigma'|_{(A \cup B)}] \subseteq Y \text{ for some } \sigma' \in \text{Act}_s \text{ such that } \sigma'|_A = \sigma|_A \}$

□

5.2 Some Remarks on Expressiveness and Definability


Proposition 3 *Let a, b be different agents and p, q be different atomic propositions. Then the following hold.*⁹

1. $[a]_\beta(p; \langle b \rangle q) \not\equiv [a]_\alpha(p; \langle b \rangle q)$.
2. $\langle a \rangle_c(p; \langle b \rangle q)$ is not definable in CL.
3. $[a]_c(p|q)$ (and, consequently, $[a]_\beta(p; \langle \emptyset \rangle q)$) is not definable in CL.
4. $[b]_\alpha(q; \langle a \rangle p)$ is not definable in CL.

Proof The claims follow respectively from Examples 2, 3, 4 and 5. □

The results above generalise to pairwise coalitions in a straightforward way.


Example 3 The game models $\mathcal{M}_1, \mathcal{M}_2$ below involve three players: a, b, c.



⁹ Even though we state the non-definability claims for CL, they apply likewise even to ATL* with its standard, memory-based semantics, because all formulae of ATL* are invariant under CL-bisimulations with respect to that semantics.


Note that (1) The relation $\mathcal{R} = \{(s_i, t_i) \mid i = 0, 1, 2, 3\}$ is an CL-bisimulation between \mathcal{M}_1 and \mathcal{M}_2 ; (2) $\mathcal{M}_1, s_0 \Vdash \langle a \rangle_c(p; \langle b \rangle q)$ but $\mathcal{M}_2, t_0 \not\Vdash \langle a \rangle_c(p; \langle b \rangle q)$.

Example 4 The game models \mathcal{M}_1 and \mathcal{M}_2 below involve two players: a and b.



Note that (1) The relation $\mathcal{R} = \{(s_i, t_i) \mid i = 0, 1, 2, 3\}$ is an CL-bisimulation between \mathcal{M}_1 and \mathcal{M}_2 ; (2) $\mathcal{M}_1, s_0 \Vdash [a]_c(p|q)$ but $\mathcal{M}_2, t_0 \not\Vdash [a]_c(p|q)$.

Example 5 The game models $\mathcal{M}_1, \mathcal{M}_2$ below involve three players: a, b, c.



Note that (1) The relation $\mathcal{R} = \{(s_i, t_i) \mid i = 0, 1, 2, 3\}$ is an CL-bisimulation between \mathcal{M}_1 and \mathcal{M}_2 ; (2) $\mathcal{M}_1, s_0 \Vdash [b]_a(q; \langle a \rangle p)$ but $\mathcal{M}_2, t_0 \not\Vdash [b]_a(q; \langle a \rangle p)$.

6 Axiomatic System for ConStR

Here we propose systems of axiom schemes for each of the basic operators of ConStR and for the whole logic, without stating completeness claims; these are left for a follow-up work. Some of these axiom schemes are adapted from the axiomatic systems for SFCL and GPCL presented in Goranko and Enqvist (2018).

6.1 Common Axiom Schemes and Rules for ConStR

The axiomatic system AX_{ConStR} builds on the axiom schemes and rules of the complete axiomatic system for Coalition Logic AX_{CL} , given in Pauly (2001, 2002). Analogues

of these axiom schemes and rules are added for each of the conditional strategic operators occurring in the fragment of ConStR that is to be axiomatized. (Recall that the coalitional operator of CL is a special case of each of O_c , O_α , O_β .) Some of these common axiom schemes and rules will turn out derivable from the special ones added below, but we are not concerned now with minimality of our system.

6.2 Additional Axiom Schemes and Rules for O_c

Axiom schemes:

- (O_c 1) Monotonicity w.r.t. A:
 $\langle A \rangle_c(\phi; \langle B \rangle \psi) \rightarrow \langle A \cup C \rangle_c(\phi; \langle B \rangle \psi)$ for any $C \subseteq \text{Agt}$
- (O_c 2) Monotonicity w.r.t. B:
 $\langle A \rangle_c(\phi; \langle B \rangle \psi) \rightarrow \langle A \rangle_c(\phi; \langle B \cup C \rangle \psi)$ for any $C \subseteq \text{Agt}$
- (O_c 3) $\langle A \rangle_c(\phi; \langle B \rangle \psi) \rightarrow \langle A \cup B \rangle_c((\phi \wedge \psi); \langle \emptyset \rangle \top)$
- (O_c 4) $\langle A \rangle_c(\phi; \langle \emptyset \rangle \psi) \leftrightarrow \langle A \rangle_c((\phi \wedge \psi); \langle \emptyset \rangle \top)$
 (NB: the direction \rightarrow follows from (O_c 3).)
- (O_c 5) $\neg \langle A \rangle_c(\perp; \langle B \rangle \psi)$
- (O_c 6) $\langle A \rangle_c(\phi; \langle B \rangle \psi) \leftrightarrow \langle A \rangle_c(\phi; \langle B \setminus A \rangle \psi)$
- (O_c 7) $\langle A \rangle_c(\phi; \langle B \rangle \psi) \leftrightarrow \langle A \rangle_c(\phi; \langle B \rangle (\phi \wedge \psi))$

Rule of inference: O_c -Monotonicity (O_c -Mon):

$$\frac{\phi \rightarrow \phi', \psi \rightarrow \psi'}{\langle A \rangle_c(\phi; \langle B \rangle \psi) \rightarrow \langle A \rangle_c(\phi'; \langle B \rangle \psi')}$$

6.3 Additional Axiom Schemes and Rules for O_β

Axiom schemes:

- (O_β 1) Monotonicity w.r.t. B:
 $[A]_\beta(\phi; \langle B \rangle \psi) \rightarrow [A]_\beta(\phi; \langle B \cup C \rangle \psi)$ for any $C \subseteq \text{Agt}$.
- (O_β 2) $[A]_\beta(\phi; \langle \emptyset \rangle \phi)$
- (O_β 3) $[A]_\beta(\perp; \langle \emptyset \rangle \psi)$
- (O_β 4) $[\emptyset]_\beta(\top; \langle A \rangle \phi) \rightarrow \neg[A]_\beta(\phi; \langle B \rangle \perp)$
- (O_β 5) $[A]_\beta(\phi; \langle B \rangle \psi) \leftrightarrow [A]_\beta(\phi; \langle B \setminus A \rangle \psi)$
- (O_β 6) $[A]_\beta(\phi; \langle B \rangle \psi) \leftrightarrow [A]_\beta(\phi; \langle B \rangle (\phi \wedge \psi))$

Rule of inference: O_β -Monotonicity (O_β -Mon):

$$\frac{\phi' \rightarrow \phi, \psi \rightarrow \psi'}{[A]_\beta(\phi; \langle B \rangle \psi) \rightarrow [A]_\beta(\phi'; \langle B \rangle \psi')}$$

6.4 Additional Axiom Schemes and Rules for O_α

Axiom schemes: All axioms $(O_\beta 1) - (O_\beta 6)$, rewritten for O_α . In addition:

(O_α^*) Anti-monotonicity w.r.t. A:

$$[A \cup C]_\alpha(\phi; \langle B \rangle \psi) \rightarrow [A]_\alpha(\phi; \langle B \rangle \psi) \text{ for any } C \subseteq \mathbb{A}gt.$$

Rule of inference: O_α -Monotonicity (O_α -Mon):

$$\frac{\phi' \rightarrow \phi, \psi \rightarrow \psi'}{[A]_\alpha(\phi; \langle B \rangle \psi) \rightarrow [A]_\alpha(\phi'; \langle B \rangle \psi')}$$

6.5 Interacting Axioms for ConStR

(ConStR 1) $[A]_\alpha(\phi; \langle B \rangle \psi) \rightarrow [A]_\beta(\phi; \langle B \rangle \psi)$

(ConStR 2) $[\emptyset]_\beta(\top; \langle A \rangle \phi) \wedge [A]_\beta(\phi; \langle B \rangle \psi) \rightarrow \langle A \rangle_c(\phi; \langle B \rangle \psi)$ ¹⁰

Proposition 4 (Soundness) *The following hold for the system AX_{ConStR} .*

1. All axiom schemes are valid in the formal semantics of ConStR.
2. The analogue of the anti-monotonicity axiom with respect to A for O_β :
 $[A \cup C]_\beta(\phi; \langle B \rangle \psi) \rightarrow [A]_\beta(\phi; \langle B \rangle \psi)$, for any A, B, C $\subseteq \mathbb{A}gt$, is not valid.

Proof Checking the soundness of most of the axioms is routine application of the formal semantics and we leave the details to the reader. We will only verify here the anti-monotonicity axiom scheme (O_α^*) , which is less straightforward and plays a special role of distinguishing O_β and O_α .

It suffices to prove the validity of the following instance:

$$[A \cup C]_\alpha(p; \langle B \rangle q) \rightarrow [A]_\alpha(p; \langle B \rangle q), \text{ for any } A, B, C \subseteq \mathbb{A}gt.$$

Consider any CGM \mathcal{M} and a state s in it.

For any formula ϕ we denote $\|\phi\|_{\mathcal{M}} := \{w \in \mathcal{M} \mid \mathcal{M}, w \Vdash \phi\}$.

Suppose $\mathcal{M}, s \Vdash [A \cup C]_\alpha(p; \langle B \rangle q)$. Fix a joint action σ_B witnessing the truth of that antecedent. Thus, for every joint action $\sigma_{A \cup C}$ such that $\text{Out}[s, \sigma_{A \cup C}] \subseteq \|\phi\|_{\mathcal{M}}$ we have that $\text{Out}[s, \sigma_{(A \cup C)} \uplus \sigma_B] \subseteq \|\phi\|_{\mathcal{M}}$.

To show that $\mathcal{M}, s \Vdash [A]_\alpha(p; \langle B \rangle q)$, we use the same joint action σ_B .

Consider any joint action σ_A such that $\text{Out}[s, \sigma_A] \subseteq \|\phi\|_{\mathcal{M}}$.

That implies that for any additional joint action $\sigma_{C \setminus A}$, we have for the resulting joint action $\sigma_{A \cup C} := \sigma_A \uplus \sigma_{C \setminus A}$ that $\text{Out}[s, \sigma_{A \cup C}] \subseteq \|\phi\|_{\mathcal{M}}$.


Then, by the assumption above, we have that $\text{Out}[s, \sigma_{(A \cup C)} \uplus \sigma_B] \subseteq \|\phi\|_{\mathcal{M}}$.

Now, note that every extension of $\sigma_A \uplus \sigma_B$ to a full strategy profile σ' can be obtained by first selecting a joint action σ'_C which coincides with $\sigma_A \uplus \sigma_B$ on $A \cup B$ and is defined according to σ' for the agents in $C \setminus (A \cup B)$. Equivalently, every such strategy profile σ' can be generated by first selecting a joint action $\sigma_{A \cup C}$ which extends σ_A with the joint action $\sigma_{C \setminus A}$ obtained by restricting σ'_C to $C \setminus A$, then adding the actions of the agents from $B \setminus (A \cup C)$ according to $\sigma_A \uplus \sigma_B$, and then adding the remaining actions according to σ' . Thus, σ' can be constructed as an extension of $\sigma_{(A \cup C)} \uplus \sigma_B$, where $\sigma_{(A \cup C)}$ extends σ_A . Hence, the outcome from s of every such σ' is in $\|\phi\|_{\mathcal{M}}$.

¹⁰ An analogue of (ConStR 2) for O_α is easily derivable from (ConStR 1) and (ConStR 2).

Therefore $\text{Out}[s, \sigma_A \uplus \sigma_B] \subseteq \|q\|_{\mathcal{M}}$. Hence, $\mathcal{M}, s \Vdash [A]_{\alpha}(p; \langle B \rangle q)$. Thus, we have shown that $\mathcal{M}, s \Vdash [A \cup C]_{\alpha}(p; \langle B \rangle q) \rightarrow [A]_{\alpha}(p; \langle B \rangle q)$, whence the validity of that formula.

2. The instance $[[\mathbf{a}, \mathbf{c}]]_{\beta}(p; \langle \mathbf{b} \rangle q) \rightarrow [\mathbf{a}]_{\beta}(p; \langle \mathbf{b} \rangle q)$ of the anti-monotonicity principle w.r.t. A for O_{β} is falsified in the example below of a game model \mathcal{M} with three players, \mathbf{a} , \mathbf{b} and \mathbf{c} .



It is easy to verify that $\mathcal{M}, s_0 \Vdash [[\mathbf{a}, \mathbf{c}]]_{\beta}(p; \langle \mathbf{b} \rangle q)$ but $\mathcal{M}, s_0 \not\Vdash [\mathbf{a}]_{\beta}(p; \langle \mathbf{b} \rangle q)$. Hence, $\mathcal{M}, s_0 \not\Vdash [[\mathbf{a}, \mathbf{c}]]_{\beta}(p; \langle \mathbf{b} \rangle q) \rightarrow [\mathbf{a}]_{\beta}(p; \langle \mathbf{b} \rangle q)$. \square

Thus, the anti-monotonicity principle with respect to A is (at least so far) the only axiom scheme in our system, that distinguishes O_{α} from O_{β} .

7 Concluding Remarks: The Road Ahead

First, we note that, while the new strategic operators introduced here can be expressed in a suitable version of Strategy Logic (cf. Mogavero et al. 2017), we choose—for both conceptual and computational reasons—to stay within a purely modal framework where actions and strategies are not explicitly referred and quantified over in the language, but are only present in the semantics.

We regard this work as a step towards developing a rich technical framework for logic-based conditional strategic reasoning of rational agents. The major further steps and directions include:

1. Completeness proofs of the proposed axiomatizations of the 3 main fragments and the entire logic ConStR are currently under construction.
2. We also claim that, for general reasons, ConStR has the finite tree-model property, and is therefore decidable (to be proved in a follow-up paper). We further conjecture that its satisfiability problem is PSPACE-complete—a major argument in favour of the modal approach to formalising conditional strategic reasoning advocated here, as opposed to one based on a version of Strategy Logic. We also leave the question of the precise complexity of model checking subject to further investigation, but conjecture that it will still be tractable in the size of the model, as in CL and ATL.
3. Extension of the framework to a full-fledged, *long term* conditional strategic reasoning, by extending the language with standard temporal operators, to produce an ATL-like extension of ConStR.
4. Long term conditional strategic reasoning naturally requires considerations about strategic commitments and model updates (cf. Ågotnes et al., 2007, 2008) and,

more generally, requires involving strategy contexts in the semantics (Brihaye et al., 2009).

5. Adding knowledge in the semantics, and explicitly in the language, by assuming that the agents reason and act under imperfect information.
6. Last, but most important long-term objective of this project is to model and capture by semantically richer logic-based formalism the *mutually conditional strategic reasoning*, where all agents reason about their strategic choices, conditional on the others' strategic choices, conditional on the reasoners' choices, etc., recursively.

Acknowledgements We thank the reviewers for some helpful remarks. The first author's work was partly supported by a research Grant 2015-04388 of the Swedish Research Council. The second author's work was funded by the Minister of Science and Higher Education within the program under the name 'Regional Initiative of Excellence' in 2019-2022, Project No. 028/RID/2018/19; the amount of funding: 11 742 500 PLN.

Funding Open access funding provided by Stockholm University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abdou, J., & Keiding, H. (1991). *Effectivity functions in social choice*. Springer.
- Ågotnes, T., Goranko, V., & Jamroga, W. (2007). Alternating-time temporal logics with irrevocable strategies. In D. Samet (Ed.) *Proceedings of TARK XI* (pp. 15–24).
- Ågotnes, T., Goranko, V., & Jamroga, W. (2008). *Strategic commitment and release in logics for multi-agent systems (Extended abstract)*. Technical Report IfI-08-01, Clausthal University of Technology.
- Ågotnes, T., Goranko, V., Jamroga, W., & Wooldridge, M. (2015). Knowledge and ability. In H. van Ditmarsch, J. Halpern, W. van der Hoek, & B. Kooi (Eds.), *Handbook of epistemic logic* (pp. 543–589). College Publications.
- Alur, R., Henzinger, T.A., Kupferman, O., & Vardi, M. (1998). Alternating refinement relations. In D. Sangiorgi, & R. de Simone (Eds.), *Proceedings of CONCUR'98*. LNCS 1466. Springer.
- Alur, R., Henzinger, T. A., & Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, 49(5), 672–713.
- Brihaye, T., Lopes, A. D. C., Laroussinie, F., & Markey, N. (2009). ATL with strategy contexts and bounded memory. In S. Artëmov & A. Nerode (Eds.), *Proceedings of LFCS'2009*, LNCS (Vol. 5407, pp. 92–106). Springer.
- Fisman, D., Kupferman, O., & Lustig, Y. (2010). Rational synthesis. In *Proceedings of TACAS* (Vol. 2010, pp. 190–204).
- Goranko, V. (2001). Coalition games and alternating temporal logics. In J. van Benthem (Ed.), *Proceedings of TARK VIII* (pp. 259–272). Morgan Kaufmann.
- Goranko, V., & Enqvist, S. (2018). Socially friendly and group protecting coalition logics. In *Proceedings of AAMAS 2018* (pp. 372–380).
- Goranko, V., & Ju, F. (2019). Towards a logic for conditional local strategic reasoning. In P. Blackburn, E. Lorini, & M. Guo (Eds.), *Logic, rationality, and interaction—7th international workshop, LORI 2019, Chongqing, China, October 18–21, 2019, Proceedings, Lecture Notes in Computer Science* (Vol. 11813, pp. 112–125). Springer.

- Gutierrez, J., Harrenstein, P., & Wooldridge, M. J. (2017). From model checking to equilibrium checking: Reactive modules for rational verification. *Artificial Intelligence*, 248, 123–157.
- Kupferman, O., Perelli, G., & Vardi, M. Y. (2016). Synthesis with rational environments. *Annals of Mathematics and Artificial Intelligence*, 78(1), 3–20.
- Mogavero, F., Murano, A., Perelli, G., & Vardi, M. Y. (2014). Reasoning about strategies: On the model-checking problem. *ACM Transactions on Computational Logic*, 15(4), 34:1-34:47.
- Mogavero, F., Murano, A., Perelli, G., & Vardi, M. Y. (2017). Reasoning about strategies: On the satisfiability problem. *Logical Methods in Computer Science*, 13(1), 1–37.
- Naumov, P., & Yuan, Y. (2019). *Intelligence in strategic games*. CoRR [arXiv:1910.07298](https://arxiv.org/abs/1910.07298)
- Pauly, M. (2001). *Logic for social software*. Ph.D. thesis, University of Amsterdam
- Pauly, M. (2002). A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1), 149–166.
- van Benthem, J., Ghosh, S., & Verbrugge, R. (Eds.). (2015). *Models of strategic reasoning—Logics, games, and communities*, LNCIS (Vol. 8972). Springer.
- van der Hoek, W., & Wooldridge, M. (2004). Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, 75(1), 125–157.
- Wooldridge, M., Gutierrez, J., Harrenstein, P., Marchioni, E., Perelli, G., & Toumi, A. (2016). Rational verification: From model checking to equilibrium checking. In *Proceedings of AAAI 2016* (pp. 4184–4191).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.