Towards a Logic for Conditional Local Strategic Reasoning *

Valentin Goranko[1] and Fengkui Ju[2]

[1] Stockholm University, Sweden, and University of Johannesburg, South Africa***
valentin.goranko@philosophy.su.se
[2] School of Philosophy, Beijing Normal University
fengkui.ju@bnu.edu.cn

Abstract. We consider systems of rational agents who act in pursuit of their individual and collective objectives and we study the reasoning of an agent or an external observer about the consequences from the expected choices of action of the other agents based on their objectives, in order to assess the reasoners ability to achieve his own objective.
To formalize such reasoning we introduce new modal operators of conditional strategic reasoning and use them to extend Coalition Logic in order to capture variations of conditional strategic reasoning. 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.

Keywords: Concurrent game models · Conditional strategic reasoning · Coalition Logic · 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 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.

Put in a more general perspective, we consider a system of agents acting independently, and possibly concurrently, in pursuit of their individual and collective goals and we analyse the reasoning of an agent (or, just an observer)

*** Visiting professorship
* The work of Valentin Goranko was supported by a research grant 2015-04388 of the Swedish Research Council. The work of Fengkui Ju was supported by the Major Program of the National Social Science Foundation of China (NO. 17ZDA026). We thank the reviewers for some helpful remarks.

about the possible local actions (at the current state only) of the other agents and their possible effect for realising the outcome of interest for the observer or, respectively, for enabling the reasoner to act towards achieving his own goal.

Our contributions. In this paper we identify several distinct cases of conditional strategic reasoning of an observer or an active agent, depending on his knowledge about the objectives and possible actions of the other agents. To formalize such reasoning we introduce new modal operators of conditional strategic reasoning and use them to extend Coalition Logic to capture variations of conditional strategic reasoning. 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.

Related work. 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 ([10], [11]) and its temporal extension ATL ([4]), and evolving towards increasingly expressive formalisms, such as Strategy Logic (cf. [9]) (cf. [5] and [3] for overviews of the area). Still, we are aware of very few works that deal more explicitly with *conditional* strategic reasoning in the sense of the present paper, with perhaps the closest being the recent [8], to which the present work relates both conceptually and technically. In the literature there has been some work on reasoning about agents' goals (cf. [6]).

Structure of the paper Section 2 provides an informal discussion on conditional strategic reasoning, motivating the further technical work. Section 3 introduces several modal operators formalising patterns conditional strategic reasoning and uses them to introduce the new logic ConStR as an extension of Coalition Logic with these operators. Section 4 introduces the matching notion of bisimulation for that logic and discuss briefly it expressiveness. We end with brief concluding remarks in Section 5.

2 Conditional strategic reasoning: an informal discussion

Suppose that Alice has an objective α to achieve (say, pick a book from a friend). Suppose also that Alice has several possible choices of an action (or strategy) [3] that would possibly, or certainly, guarantee the achievement of her objective.

[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.

2.1 Conditional reasoning of an observer about an 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 β. More generally, we 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 a *passive observer's conditional strategic reasoning* about statements of the type:

"*Some/every action of Alice that guarantees achievement of* α *also guarantees/enables occurrence of the (desired or expected) outcome* β".

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 β.

Bob's reasoning, case 1: Bob knows nothing about Alice. Suppose that *Bob does not know Alice's objective*, and therefore has no a priori expectations about her choice of action. E.g., if Alice is coming to the university and Bob is standing by the only entrance of the campus, he will know for sure that he is going to meet Alice, no matter what she will do there.

Then, Bob can *only* claim for sure that the outcome β will occur if β is inevitable, regardless of how Alice (and all others) will act. This can be expressed in Coalition Logic CL (cf. [11]) simply as $[\emptyset]\beta$.

Bob's reasoning, case 2: Bob only know Alice's objective. 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 pick some book, but does not know where and when.

Then, Bob can only claim that the outcome β will occur for sure if β is true *on every possible course of events ("play") on which* α *is true.* (E.g., Bob knows that the book is in the library, and β is the event "Alice enters the library building".) This can be expressed as a conditional $\alpha \to \beta$, in the right context. Depending on how the conditional is interpreted, there are different cases:

- $\alpha \to \beta$ is a material implication, with unconstrained context.
  This can be expressed in CL simply as $[\emptyset](\alpha \to \beta)$.
- $\alpha \to \beta$ is a proper conditional, with a somehow constrained context.
  In general, this cannot be expressed in CL anymore, but it can possibly be expressed in a suitably extended language and in a suitably updated model. Here we will not pursue this line, but will leave it to a follow-up work.

Bob's reasoning, case 3: Bob knows Alice's objective and possible actions. Suppose now Bob not only knows Alice's objective, but also *knows all possible actions / strategies* of Alice that can ensure the satisfaction of her objective α, and knows that Alice will commit to one of them, but *does not know to which one*. (E.g., Bob knows that Alice is coming to campus to meet with her supervisor and she can meet with him either in his office, or in the lecture room, or in the café.)

Now, for Bob to claim that the outcome β will occur for sure, he must know that *each* action of Alice that guarantees α will also guarantee β. (E.g., suppose that all possible meeting places are in the main building and β is the event "Alice enters the main building".) This can no longer be expressed in CL and requires introducing a new strategic operator.

Bob's reasoning, case 4: Bob knows Alice's action. Lastly, suppose that *Bob knows the specific action which she is taking in order to guarantee the achievement of her objective*. Then, Bob can claim that the outcome β will occur for sure, as long as that specific action of Alice guarantees the satisfaction of β. Again, this claim could be interpreted either in the same (original) model, or in a respectively updated one, obtained by preserving only the plays that are enabled by that action if Alice's strategic commitment is assumed. The latter corresponds to *reasoning with strategy contexts* which we will not discuss here, but in a follow-up work.

2.2 Conditional reasoning of an agent about another agent's actions

Suppose now that Bob is not just an observer, but also an acting agent, who has the outcome β as his own goal. Suppose also that there may be other agents, besides Alice and Bob, 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.

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 β. This calls for an *agent's conditional strategic reasoning* about statements of the type:
"*For some/every action of Alice that guarantees achievement of α, Bob has/does not have an action of his own to guarantee achievement of his objective β*".

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 Section 2.1 accordingly applies here, too. However, now in the reasoning case 3 the statement

> "*Bob knows that whichever way Alice acts towards achieving the objective α, he can act so as to bring about achievement of his objective β.*"

admits two different readings, *de dicto* and *de re*, which we discuss here.

Bob's reasoning, case 3: *de dicto* reading. In the *de dicto* reading, where Bob only knows that Alice has committed to act so as to achieve α, but, as far as he knows, Alice has not yet chosen her action, or her choice will remain unknown to Bob.

In this case Bob must consider all possible courses of events (plays) that can occur as a result of Alice acting towards achieving α and reason about whether he can act *uniformly* for each of them in a way that would bring about β, without knowing which of them will take place. (E.g., in our running story from Section 2.1, Bob can choose to wait for Alice at the only entrance of the main 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 α, and is looking for an action that will bring about β on each of these outcome states.

Bob's reasoning, case 3: *de re* reading. This is the reading where *for every action* of Alice that ensures α, Bob is looking for an action of his, *possibly dependent on Alice's action* that would also ensure the occurrence of β (possibly in different ways for the different actions). More formally, each of Alice's actions that would guarantee α generates a set of possible outcome states, and for each of them Bob is looking for an action that will bring about β on that set of outcome states.

For example, suppose Bob knows that Alice has agreed with her friend Charlie on a meeting on campus today and there are two options: to meet in the campus café or to meet in the library; both options are ok for Charlie and Alice is yet to decide on either option. Note, that the sentence "*Alice has decided to meet with Charlie on campus today*" is true in either case. However, the sentence "*Alice has decided to meet with Charlie in the café today, or Alice has decided to meet with Charlie in the library today*" should not be regarded as true (yet). After Alice makes her choice, this sentence becomes true, too. But even then, from Bob's perspective, the same distinction applies depending on whether or not he knows Alice's choice, so he has to take into account both options when deciding for himself on what to do.

3 A logic for conditional strategic reasoning

3.1 Preliminaries

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

Definition 1 (Multi-agent game model). *A* ***game model***[4] *for* $\mathbb{A}\mathrm{gt}$ *and* Π *is a tuple*

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

[4] These game models are essentially equivalent to concurrent game models used in [4].

where S is a non-empty set of ***states****; each Σ_a is a non-empty set of possible* ***actions*** *of agent a; $V : \Pi \to \mathcal{P}(S)$ 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 $g(s) = (\Sigma^s_{a_1}, ...\Sigma^s_{a_n}, o_s)$, where each $\Sigma^s_{a_i} \subseteq \Sigma_{a_i}$ is a non-empty set of actions available to player a_i at s, and*

$$o_s : \Sigma^s_{a_1} \times ... \times \Sigma^s_{a_n} \to S$$

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

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

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

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

Given a coalition $\mathrm{C} \subseteq \mathbb{A}\mathrm{gt}$, a ***joint action*** *for C in the model $\mathcal{M}$ is a tuple of individual actions $\sigma_\mathrm{C} \in \prod_{a \in \mathrm{C}} \Sigma_a$. 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:*

$$\mathsf{Out}[s, \sigma_\mathrm{C}] = \{u \in S \mid \exists \sigma \in \mathrm{Act}_s :\ \sigma|_\mathrm{C} = \sigma_\mathrm{C} \ \&\ o_s(\sigma) = u\}$$

where $\sigma|_\mathrm{C}$ is the restriction of σ to C. Note that the empty tuple $\sigma_\emptyset$ is the only available joint action for the empty coalition $\emptyset$ at any state.

The basic logic for coalitional strategic reasoning CL. Coalition Logic (CL) was introduced in [10], cf. also [11]. CL extends the classical propositional logic with *coalitional strategic modal operators* $[\mathrm{C}]$, for any coalition of agents C. Formulae of CL:

$$\varphi := p \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \ [\mathrm{C}]\,\varphi$$

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

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

Semantics of CL. The formulae of CL are interpreted in game models (GM). The semantics is defined in terms of **truth of a CL-formula ψ at a state s of a GM** $\mathcal{M}$, denoted $\mathcal{M}, s \vDash \psi$, by induction on formulae, via the key clause:
$\mathcal{M}, s \models [\mathrm{C}]\,\phi$ iff there exists a joint action σ_C available at s, such that $\mathcal{M}, u \models \phi$ for each $u \in \mathsf{Out}[s, \sigma_\mathrm{C}]$.

We note that $[\mathrm{C}]\,\phi$ formalises a claim of the ability of the agent/coalition C to choose a suitable (joint) action so as to achieve the goal ϕ *regardless of how all other agents choose to act*, and therefore regardless of whether the agents in C know the goal(s) of the remaining agents. This subsumes Cases 1 and 2 of Bob's reasoning, discussed in Section 2.1.

Alternating bisimulations. The notion of bisimulation that guarantees truth invariance of all CL-formulae was first defined in [10] for the abstract game models defined there, and later (under the name "alternating bisimulations") in [1], to which we refer the reader for the definition and proof of bisimulation invariance of all ATL-formulae (incl. all CL-formulae).

3.2 The logic of conditional strategic reasoning ConStR

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 $\mathrm{A} \cap \mathrm{B}$.

Operators for conditional strategic reasoning of ConStR.

We now introduce new operators for conditional strategic reasoning, for any coalitions A and B with intuitive semantics as follows.

(O_c) $\langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi)$: A has a joint action σ_{A} such that, when applied, it 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 $\mathrm{B} \setminus \mathrm{A}$ act according to σ_{B}.

This operator formalises a scenario (not discussed in Section 2, but still basic for conditional strategic reasoning, see further) where A knows the objective of B and can choose to cooperate with B by selecting a suitable action.

(O_{dr}) $[\mathrm{A}]_{\mathsf{dr}}(\phi; \langle \mathrm{B} \rangle \psi)$: 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.

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 *de re* reading of Case 3 of Bob's reasoning, discussed in Section 2.2. In particular, the case when B is not informed about the goal of A, but has to choose their action after learning the action of A is formalised by $[\mathrm{A}]_{\mathsf{dr}}(\top; \langle \mathrm{B} \rangle \psi)$.

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

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 *de dicto* reading of Case 3 of Bob's reasoning, discussed in Section 2.2. In particular, the case when B is not informed about the goal of A and has to choose a joint action before A has chosen their action is formalised by $[\mathrm{A}]_{\mathsf{dd}}(\top; \langle \mathrm{B} \rangle \psi)$.

Language of ConStR. We fix a finite nonempty set of agents $\mathbb{A}\mathrm{gt}$ and a countable set of atomic propositions Π.

The formulae of ConStR, where $p \in \Pi$ and $\mathrm{A}, \mathrm{B} \subseteq \mathbb{A}\mathrm{gt}$ are defined as follows:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \phi) \mid [\mathrm{A}]_{\mathsf{dr}}(\phi; \langle \mathrm{B} \rangle \phi) \mid [\mathrm{A}]_{\mathsf{dd}}(\phi; \langle \mathrm{B} \rangle \phi)$$

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 dual operator $\neg\langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \neg\psi)$ 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.
- $[\mathrm{A}]_{\mathsf{c}}(\phi | \psi) := [\mathrm{A}]_{\mathsf{dr}}(\phi; \langle \emptyset \rangle \psi)$: for any joint strategy of A, if it guarantees ϕ to be true then it guarantees ψ to be true, too. This operator formalises Case 2 of Bob's reasoning as an observer (rather than an acting agent), discussed in Section 2.1.
- $\langle \mathrm{A} \rangle_{\mathsf{c}}(\phi | \psi) := \neg[\mathrm{A}]_{\mathsf{c}}(\phi | \neg\psi)$: 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, $[\mathrm{C}](\phi; \psi_1, \ldots, \psi_k)$, introduced in [8], viz. $\langle \mathrm{A} \rangle_{\mathsf{c}}(\phi | \psi) \equiv [A](\phi; \psi)$. Moreover, $\langle \mathrm{A} \rangle_{\mathsf{c}}(\phi | \psi)$ is also definable as $\langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \overline{\mathrm{A}} \rangle \psi)$, where $\overline{\mathrm{A}} = \mathbb{A}\mathrm{gt} \setminus \mathrm{A}$.
- The coalitional strategic operator $[\mathrm{A}]$ from CL is a special case of the above: $[\mathrm{A}]\phi := \langle \mathrm{A} \rangle_{\mathsf{c}}(\phi | \top)$, meaning "A has a joint action to ensure the truth of ϕ"[5].
- $\langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi)$ is definable in terms of the "group protecting coalitional operator" GIP, introduced in [8]: $\langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \equiv \langle\!\!| A \triangleright \phi, \mathrm{A} \cup \mathrm{B} \triangleright \psi |\!\!\rangle$. Nevertheless, it now has a different motivation and intuitive interpretation.

Semantics of ConStR. Given coalitions $\mathrm{A}, \mathrm{B} \subseteq \mathbb{A}\mathrm{gt}$ and joint actions σ_{A} for A and σ_{B} for B, we define $\sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}$ to be the joint action for $\mathrm{A} \cup \mathrm{B}$ which equals to σ_{A} when restricted to A and equals to $\sigma_{\mathrm{B}}|_{\mathrm{B} \setminus \mathrm{A}}$ when restricted to $\mathrm{B} \setminus \mathrm{A}$. Note $\sigma_{\mathrm{A}} \uplus \sigma_{\emptyset} = \sigma_{\mathrm{A}}$ for any $\mathrm{A} \subseteq \mathbb{A}\mathrm{gt}$.

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

$\mathcal{M}, s \Vdash \langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \;\Leftrightarrow\;$ A has a joint action σ_{A}, such that $\mathcal{M}, u \Vdash \phi$ for every $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}}]$ and B has a joint action σ_{B} such that $\mathcal{M}, u \Vdash \psi$ for every $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}]$.

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

$\mathcal{M}, s \Vdash [\mathrm{A}]_{\mathsf{dr}}(\phi; \langle \mathrm{B} \rangle \psi) \Leftrightarrow$ for every joint action σ_{A} of A such that $\mathcal{M}, u \Vdash \phi$ for every $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}}]$, B has a joint action σ_{B} (generally, dependent on σ_{A}) such that $\mathcal{M}, u \Vdash \psi$ for every $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}]$.

$\mathcal{M}, s \Vdash [\mathrm{A}]_{\mathsf{dd}}(\phi; \langle \mathrm{B} \rangle \psi) \Leftrightarrow$ B has a joint action σ_{B} such that for every joint action σ_{A} of A, if $\mathcal{M}, u \Vdash \phi$ for each $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}}]$, then $\mathcal{M}, u \Vdash \psi$ for each $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}]$.

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

In the Appendix we provide a few simple examples illustrating the semantics of the strategic operators introduced here.

4 Bisimulations and expressiveness of ConStR

4.1 Bisimulations for CSR

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_{dr}, and O_{dd}[6]. We only define ConStR-bisimulation within a game model, which generalises to ConStR-bisimulation *between* game models, by treating both as parts of their disjoint union.

Definition 2 (ConStR-bisimulation). *Let $\mathcal{M} = (S, \{\Sigma_a\}_{a \in \mathbb{A}\mathrm{gt}}, g, V)$ be a game model. A binary relation $\beta \subseteq S^2$ is a ConStR-**bisimulation in** $\mathcal{M}$ if it satisfies the following conditions for every pair of states (s_1, s_2) such that $s_1 \beta s_2$ 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 Figure 1)*

- A-**Forth$_{\mathsf{c}}$:** *For any joint action σ^1_{A} of* A *at s_1 there is a joint action σ^2_{A} of* A *at s_2, such that:*
  - A-**LocalBack$_{\mathsf{c}}$:** *For every $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}}]$ there exists $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}}]$ such that $u_1 \beta u_2$.*
  - B-**Forth$_{\mathsf{c}}$:** *For every joint action σ^1_{B} of* B *at s_1 there is a joint action σ^2_{B} of* B *at s_2, such that:*
    - **($\mathrm{A} \uplus \mathrm{B}$)-LocalBack$_{\mathsf{c}}$:** *For every $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}} \uplus \sigma^2_{\mathrm{B}}]$ there exists $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}} \uplus \sigma^1_{\mathrm{B}}]$ such that $u_1 \beta u_2$.*
- A-**Back$_{\mathsf{c}}$:** *Like* A-***Forth**$_c$, but with 1 and 2 swapped.*

O_{dr}-bisimulation:

- A-**Forth$_{\mathsf{dr}}$:** *For any joint action σ^1_{A} of* A *at s_1 there is a joint action σ^2_{A} of* A *at s_2, such that:*

[6] 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 [8].

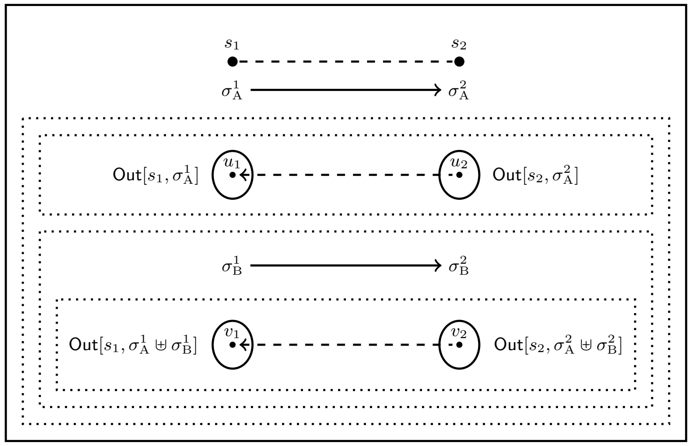

Fig. 1. The A-**Forth**$_{\mathsf{c}}$ half of O_c-bisimulation

A-**LocalBack**$_{\mathsf{dr}}$**:** *For every* $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}}]$ *there exists* $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}}]$ *such that* $u_1 \beta u_2$*.*

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

(A $\uplus$ B)-**LocalForth**$_{\mathsf{dr}}$**:** *For every* $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}} \uplus \sigma^1_{\mathrm{B}}]$ *there exists* $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}} \uplus \sigma^2_{\mathrm{B}}]$ *such that* $u_1 \beta u_2$*.*

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

O_{dd}**-bisimulation:**

B-**Forth**$_{\mathsf{dd}}$**:** *For any joint action* σ^1_{B} *of* B *at* s_1 *there is a joint action* σ^2_{B} *of* B *at* s_2*, such that:*

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

(A)-**LocalForth**$_{\mathsf{dd}}$**:** *For every* $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}}]$ *there exists* $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}}]$ *such that* $u_1 \beta u_2$*.*

(A $\uplus$ B)-**LocalBack**$_{\mathsf{dd}}$**:** *For every* $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}} \uplus \sigma^2_{\mathrm{B}}]$ *there exists* $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}} \uplus \sigma^1_{\mathrm{B}}]$ *such that* $u_1 \beta u_2$*.*

B-**Back**$_{\mathsf{dd}}$**:** *Like* B-***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 bisimulation* β *in* $\mathcal{M}$ *such that* $s_1 \beta s_2$*.*

Proposition 1 (ConStR**-bisimulation invariance). *Let* β *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* $s_1 \beta s_2$*:* $\mathcal{M}, s_1 \models \theta$ *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.

For the strategic operators, we only check here the case of $\theta = \langle\langle \mathrm{A}\rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B}\rangle\psi)$, assuming that the claim holds for ϕ and ψ. The cases of $[\mathrm{A}]_{\mathsf{dr}}(\phi; \langle \mathrm{B}\rangle\psi)$ and $[\mathrm{A}]_{\mathsf{dd}}(\phi; \langle \mathrm{B}\rangle\psi)$ are quite analogous. Their proofs are omitted for lack of space.

(Case O_c) Let $\theta = \langle\langle \mathrm{A}\rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B}\rangle\psi)$, assuming that the claim holds for ϕ and ψ.

Suppose, $\mathcal{M}, s_1 \models \theta$. Then A has a joint action σ^1_{A} 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$_{\mathsf{c}}$**, there is a joint action σ^2_{A} of A at s_2, such that, by A-**LocalBack$_{\mathsf{c}}$**, for each $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}}]$ there exists $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}}]$ such that $u_1 \beta u_2$. By the choice of σ^1_{A}, $\mathcal{M}, u_1 \models \phi$ for each $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}}]$. It follows, by the inductive hypothesis applied to ϕ, that $\mathcal{M}, u_2 \models \phi$ for each $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}}]$. Moreover, B has a joint action σ^1_{B} at s_1 such that, when applied by B, in addition to A applying σ^1_{A}, it guarantees ψ, i.e. $\mathcal{M}, u_1 \models \psi$ for each $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}} \uplus \sigma^1_{\mathrm{B}}]$. By condition B-**Forth$_{\mathsf{c}}$**, there is a joint action σ^2_{B} of B at s_2, such that, by $(\mathrm{A} \uplus \mathrm{B})$**-LocalBack$_{\mathsf{c}}$**, for every $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}} \uplus \sigma^2_{\mathrm{B}}]$ there exists $u_1 \in \mathsf{Out}[s_1, \sigma^1_{\mathrm{A}} \uplus \sigma^1_{\mathrm{B}}]$ such that $u_1 \beta u_2$. Therefore, by the inductive hypothesis applied to ψ, $\mathcal{M}, u_2 \models \psi$ for each $u_2 \in \mathsf{Out}[s_2, \sigma^2_{\mathrm{A}} \uplus \sigma^2_{\mathrm{B}}]$. Thus, $\mathcal{M}, s_2 \models \theta$. The converse is similar, using A-**Back$_{\mathsf{c}}$**.

Proposition 2 (Hennessy-Milner property). *For any finite game model $\mathcal{M}$ there is a* ConStR*-bisimulation β in $\mathcal{M}$, such that for any pair $s_1, s_2 \in \mathcal{M}$, $s_1 \beta s_2$ holds iff s_1 and s_2 are* ConStR*-equivalent (satisfy the same* ConStR*-formulae).*

Proof. (Sketch) One direction follows from Prop. 1. For the converse, it suffices to prove that the relation of ConStR-equivalence is itself a ConStR-bisimulation in $\mathcal{M}$. 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, $\chi(s_1) \wedge \chi(s_2) \equiv \bot$ 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) $\bigwedge_{A,B \subseteq \mathbb{A}\mathrm{gt}} \big\{ \langle\langle \mathrm{A}\rangle\rangle_{\mathsf{c}}(\chi(Z); \langle \mathrm{B}\rangle \chi(Y)) \mid \exists \sigma \in \mathrm{Act}_s : \mathsf{Out}[s, \sigma|_A] = Z \text{ and } \mathsf{Out}[s, \sigma|_{(A \cup B)}] = Y \big\}$

(2) $\bigwedge_{A,B \subseteq \mathbb{A}\mathrm{gt}} \big\{ [\mathrm{A}]_{\mathsf{dr}}(\chi(Z); \langle \mathrm{B}\rangle \chi(Y)) \mid \forall \sigma \in \mathrm{Act}_s : \mathsf{Out}[s, \sigma|_A] \subseteq Z \text{ implies } \mathsf{Out}[s, \sigma'|_{(A \cup B)}] \subseteq Y \text{ for some } \sigma' \in \mathrm{Act}_s \text{ such that } \sigma'|_A = \sigma|_A \big\}$

(3) $\bigwedge_{A,B \subseteq \mathbb{A}\mathrm{gt}} \big\{ [\mathrm{A}]_{\mathsf{dd}}(\chi(Z); \langle \mathrm{B}\rangle \chi(Y)) \mid \exists \sigma \in \mathrm{Act}_s : \forall \sigma' \in \mathrm{Act}_s \text{ if } \mathsf{Out}[s, \sigma'|_A] \subseteq Z \text{ and } \sigma'|_{(\mathrm{B} \setminus \mathrm{A})} = \sigma|_{(\mathrm{B} \setminus \mathrm{A})} \text{ then } \mathsf{Out}[s, \sigma'|_{(A \cup B)}] \subseteq Y \big\}$

4.2 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, where $\equiv$ is the logical equivalence in* ConStR[7].

[7] Even though we state the non-definability claims for CL, they apply likewise even to ATL*, because all formulae of ATL* are invariant under alternating bisimulations.

1. $[\mathsf{a}]_{\mathsf{dr}}(p; \langle \mathsf{b} \rangle q) \not\equiv [\mathsf{a}]_{\mathsf{dd}}(p; \langle \mathsf{b} \rangle q)$*.*
2. $\langle\!\langle \mathsf{a} \rangle\!\rangle_{\mathsf{c}}(p; \langle \mathsf{b} \rangle q)$ *is not definable in* CL*.*
3. $[\mathsf{a}]_{\mathsf{c}}(p|q)$ *(and, consequently,* $[\mathsf{a}]_{\mathsf{dr}}(p; \langle \emptyset \rangle q)$*) is not definable in* CL*.*
4. $[\mathsf{b}]_{\mathsf{dd}}(q; \langle \mathsf{a} \rangle p)$ *is not definable in* CL*.*

Proof. The first 3 claims follow respectively from Examples 3, 4, and 5 in the Appendix. The proof of the last claim is analogous.

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

5 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. [9]), 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 first step towards developing a rich technical framework for logic-based conditional strategic reasoning of rational agents. The major further steps and directions include:

1. Complete axiomatization and proof of decidability of the logic ConStR (currently under development).
2. Extending 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.
3. The long term conditional strategic reasoning naturally requires considerations about strategic commitments and model updates (cf. [1], [2]) and, more generally, requires involving strategy contexts in the semantics ([7]).
4. Adding knowledge, explicitly in the language, and implicitly, in the semantics, by assuming that the agents reason and act under imperfect information.
5. 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.

References

1. Ågotnes, T., Goranko, V., Jamroga, W.: Alternating-time temporal logics with irrevocable strategies. In: Samet, D. (ed.) Proc. of TARK XI. pp. 15–24 (2007)
2. Ågotnes, T., Goranko, V., Jamroga, W.: Strategic commitment and release in logics for multi-agent systems (extended abstract). Tech. Rep. IfI-08-01, Clausthal University of Technology (2008)

3. Ågotnes, T., Goranko, V., Jamroga, W., Wooldridge, M.: Knowledge and ability. In: van Ditmarsch, H., Halpern, J., van der Hoek, W., Kooi, B. (eds.) chapter in: Handbook of Epistemic Logic, pp. 543–589. College Publications (2015)
4. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. J. ACM **49**(5), 672–713 (2002)
5. van Benthem, J., Ghosh, S., Verbrugge, R. (eds.): Models of Strategic Reasoning - Logics, Games, and Communities, Lecture Notes in Computer Science, vol. 8972. Springer (2015), https://doi.org/10.1007/978-3-662-48540-8
6. Brafman, R., Tennenholtz, M.: Modeling agents as qualitative decision makers. Artificial Intelligence **94**(1), 217 – 268 (1997)
7. Brihaye, T., Lopes, A.D.C., Laroussinie, F., Markey, N.: ATL with strategy contexts and bounded memory. In: Artëmov, S., Nerode, A. (eds.) Proc. of LFCS'2009. LNCS, vol. 5407, pp. 92–106. Springer (2009)
8. Goranko, V., Enqvist, S.: Socially friendly and group protecting coalition logics. In: Proc. of AAMAS 2018. pp. 372–380 (2018)
9. Mogavero, F., Murano, A., Perelli, G., Vardi, M.Y.: Reasoning about strategies: on the satisfiability problem. Logical Methods in Computer Science **13**(1) (2017)
10. Pauly, M.: Logic for Social Software. Ph.D. thesis, University of Amsterdam (2001)
11. Pauly, M.: A modal logic for coalitional power in games. Journal of Logic and Computation **12**(1), 149–166 (2002)

Appendix: some examples

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.

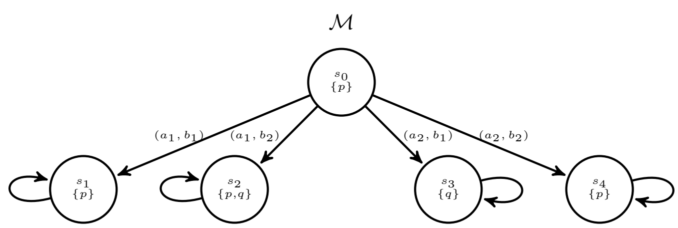

It can be verified that $\mathcal{M}, s_0 \Vdash \langle\!\langle \mathsf{a} \rangle\!\rangle_{\mathsf{c}}(p; \langle \mathsf{b} \rangle q)$, while $\mathcal{M}, s_0 \nVdash [\mathsf{b}]q$. Thus, an agent may have only conditional ability to achieve its goal.

Example 2. The game model $\mathcal{M}$ below has two players, a and b.

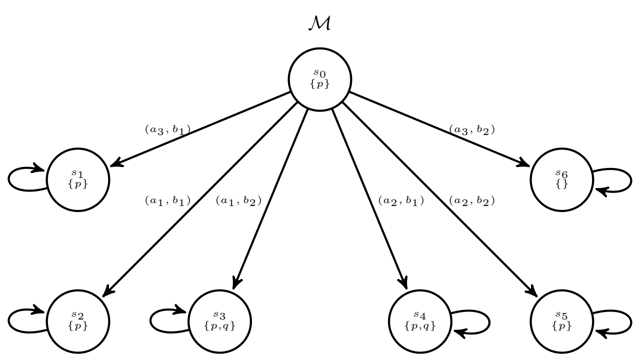

It can be verified that $\mathcal{M}, s_0 \Vdash [\mathsf{a}]_{\mathsf{dr}}(p; \langle \mathsf{b} \rangle q)$. However, $\mathcal{M}, s_0$ does not satisfy the ATL^* formula $[\![\mathsf{a}]\!](\mathrm{X}p \to \langle\!\langle \mathsf{b} \rangle\!\rangle \mathrm{X}q)$, hence these are not equivalent.

Also, $\mathcal{M}, s_0 \nVdash [\mathsf{a}]_{\mathsf{dd}}(p; \langle \mathsf{b} \rangle q)$. However, if the outcomes of (a_2, b_1) and (a_2, b_2) are swapped, then $[\mathsf{a}]_{\mathsf{dd}}(p; \langle \mathsf{b} \rangle q)$ becomes true at s_0 in the resulting model.

Example 3. The game model $\mathcal{M}$ below involves two players: a and b. It can be verified that $\mathcal{M}, s_0 \Vdash [\mathsf{a}]_{\mathsf{dr}}(p; \langle \mathsf{b} \rangle q)$ but $\mathcal{M}, s_0 \nVdash [\mathsf{a}]_{\mathsf{dd}}(p; \langle \mathsf{b} \rangle q)$.

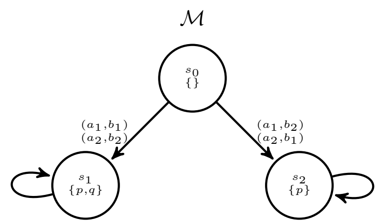

Example 4. The game models $\mathcal{M}_1$ and $\mathcal{M}_2$ below involve three players: a, b, c. It can be verified that:

(1) The relation $\beta = \{(s_i, t_i) \mid i = 0, 1, 2, 3\}$ is an alternating bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2$ (cf. [1]).

(2) $\mathcal{M}_1, s_0 \Vdash \langle\!\langle \mathsf{a} \rangle\!\rangle_{\mathsf{c}}(p; \langle \mathsf{b} \rangle q)$ but $\mathcal{M}_2, t_0 \nVdash \langle\!\langle \mathsf{a} \rangle\!\rangle_{\mathsf{c}}(p; \langle \mathsf{b} \rangle q)$.

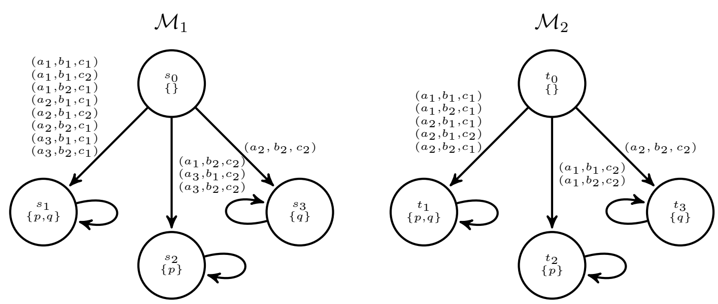

Example 5. The game models $\mathcal{M}_1$ and $\mathcal{M}_2$ below involve two players: a and b. It can be verified that:

(1) The relation $\beta = \{(s_i, t_i) \mid i = 0, 1, 2, 3\}$ is an alternating bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2$ (cf. [1]).

(2) $\mathcal{M}_1, s_0 \Vdash [\mathsf{a}]_{\mathsf{c}}(p|q)$ but $\mathcal{M}_2, t_0 \nVdash [\mathsf{a}]_{\mathsf{c}}(p|q)$.