A Logic for Conditional Strategic Reasoning

Valentin Goranko
Stockholm University
(based on joint work with **Fengkui Ju @ Beijing Normal University**)

Logic4Peace
April 23, 2022


Make Logic,
not War!

Based on:
Valentin Goranko and Fengkui Ju:
Towards a Logic for Conditional Local Strategic Reasoning,
Proc. of the International Conf. on Logic, Rationality and Interaction
LORI VII, 2019.

Extended journal version:
Valentin Goranko and Fengkui Ju:
A Logic for Conditional Local Strategic Reasoning,
Journal of Logic, Language and Information, 2022.
Online publication:
`link.springer.com/article/10.1007/s10849-022-09357-y`.

Conditional strategic reasoning:

informal discussion

Conditional strategic reasoning: an introduction

We want to formalise the reasoning of (and about) an agent acting in a multi-agent environment, conditional on his/her knowledge of the goals and choices of actions of the other agents.

I focus on a simple case: two agents, Alice and Bob, acting independently, and possibly concurrently with other agents.

Alice, has a goal γ_A to achieve.

Suppose, Alice has several possible choices of action that would possibly, or certainly, guarantee the achievement of her goal.

Bob also has several possible actions and a goal γ_B of his own.

(We ignore for now the other agents, also acting in pursuit of their goals.)

Now, based on his knowledge about Alice's goal and possible choices of actions she may take towards that goal, Bob is to decide on his own choice of action in pursuit of his goal.

A simple scenario: Alice and Bob on campus

Here is a simple illustrating scenario.

Alice and Bob are students at DownTown University.

Alice is coming to campus today, to meet with her supervisor.

Bob wants to meet with Alice somewhere on campus today.

Alice may not know that, and they may have no direct communication.

Bob may, or may not, know what Alice is going to do on campus.

Now, using his knowledge of what, where, and when Alice intends to do, Bob wants to come up with a plan of where to meet her.

Conditional strategic reasoning: Bob's reasoning about Alice's actions

This calls for a conditional strategic reasoning of the type:

For some/every action of Alice that guarantees achievement of her goal γ_A*, Bob has/does not have an action of his own to guarantee achievement of his goal* γ_B*.*

We focus on local conditional strategic reasoning, which only refers to the immediate actions of the agents, not to their *long-term global strategies*.

Remark: standard logics for multi-agent strategic reasoning, such as Coalition logic (CL) and the Alternating time temporal logic (ATL), only capture *unconditional reasoning*, where the agents' actions / strategies are to succeed against *any behaviour* of the other agents.

A more expressive logic for *conditional strategic reasoning* is needed here.

Bob's reasoning
Case 1: Bob does not know Alice's goal or actions

Depending on Bob's knowledge about Alice's goal and of her possible or expected choices of action, there can be several possible cases for Bob's reasoning.

The simplest: *Bob does not know Alice's goal, or her available actions*, and therefore has no a priori expectations about her choice of action.

Then, Bob can only make sure that γ_B will occur if he has an action to make γ_B true, *regardless* of how Alice (and all others) act.

(*E.g., if Bob is standing by the only entrance of the campus, then he is sure to meet Alice when she comes, no matter what she will do there.*)

This can be simply expressed in Coalition Logic, as $[B]\gamma_B$.

Bob's reasoning
Case 2: Bob knows Alice's goal and possible actions

Suppose now, that Bob knows Alice's goal, as well as *all possible actions* of Alice that can ensure the satisfaction of her goal.

So, Bob knows that Alice will perform one of these actions, but possibly *does not know 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.*)

We can express Bob's conditional ability to achieve his goal as follows:

> " *Whichever way Alice acts towards achieving her goal* γ_A*,*
> *Bob can act so as to ensure achievement of his goal* γ_B*.*"

This claim can no longer be expressed in Coalition Logic, except in the special case when the occurrence of Alice's goal guarantees occurrence of Bob's goal, too. (Thus, Bob need not do anything about that.)

That can again be simply expressed in CL as $[\emptyset](\gamma_A \to \gamma_B)$.

Reactive and proactive ability

Again, Bob's conditional ability to achieve his goal means:

> " *Whichever way Alice acts towards achieving her goal* γ_A*,*
> *Bob can act so as to ensure achievement of his goal* γ_B*."*

This admits two different readings:

as a reactive ability, and as a proactive ability.

(In the LORI paper these were called *ability de dicto* and *ability de re*.)

Bob's reactive ability

Suppose now, that Bob will know Alice's choice of action when he is to choose his action.

Then, Bob's ability to achieve his goal is reactive, meaning that *for every action* of Alice that ensures her goal γ_A, Bob has an action of his, *possibly dependent on Alice's action*, that would also ensure his goal γ_B.

(E.g., suppose that Alice's supervisor tells Bob where and when he is going to meet with Alice. Then Bob can wait for Alice at the respective place.)

This claim cannot be expressed in CL, so a new operator is needed for it.

Bob's proactive ability

Suppose now Bob will not know Alice's action when he is to choose his.

(*E.g., all that Bob knows is that Alice will meet with her supervisor either in his office, or in the lecture room, or in the café, but does not know where.*)

Now, for Bob to ensure that his goal will be achieved, he must have a *uniform* choice of action to make γ_B true, when applied with any action of Alice that ensures the truth of her goal γ_A.

(*E.g., suppose that all meeting places for Alice and her supervisor are in the same building, so Bob can wait for her at the only entrance of that building.*)

That is, Bob must have a proactive ability to achieve his goal.

This cannot be expressed in CL, either, so a new operator is needed again.

Remark: The notions of proactive and reactive ability respectively correspond to α-*effectivity* and β-*effectivity* in game theory.

Case 3: Bob's reasoning assuming Alice's cooperation

Suppose now that Alice also knows Bob's goal and can choose to cooperate with Bob by selecting a suitable action σ_{A}, that would not only guarantee achievement of her goal γ_A, but will also enable Bob to supplement σ_{A} with an action σ_{B} of his, which would then also guarantee achievement of his goal γ_B.

(We also assume that Alice knows enough about Bob's possible actions.)

This scenario cannot be formalised in CL, either.

Multi-agent concurrent game models informally

The dynamics of the agents' actions and interaction is modelled by multi-agent transition systems, a.k.a. concurrent game models.

Informally, these are described as follows.

▷ Agents (players) act in a common environment (the "system") by taking actions in a discrete succession of rounds.

▷ At any moment the system is in a current state.

▷ At the current state each agent independently chooses an action from a set of available actions.

▷ All agents apply their actions simultaneously.

▷ The resulting collective action determines an outcome (successor) state, which becomes the new current state.

▷ The same happens at that successor state, etc.

The logic for conditional strategic reasoning

ConStR

The logic for conditional strategic reasoning ConStR

ConStR is a modal logic featuring 3 binary modal operators, formalising reactive abilities, proactive abilities, and abilities under cooperation.

In what follows, A and B are different agents
but, more generally, they can be any two coalitions of agents.

Modal operators for conditional strategic reasoning: O_α

$[\mathrm{A}]_\alpha(\phi; \langle \mathrm{B} \rangle \psi)$, intuitively meaning:

B has an action σ_B
such that if A applies any action that guarantees the truth of ϕ,
then B can guarantee the truth of ψ by applying σ_B.

This operator formalises the notion of an agent's *proactive ability*, i.e. α-*effectivity*.

Modal operators for conditional strategic reasoning: O_β

$[\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \rangle \psi)$, intuitively meaning:

for any action σ_A of A
that guarantees the truth of ϕ when applied by A,
there is an action σ_B that guarantees ψ if additionally applied by B.

This operator formalises the notion of an agent's *reactive ability*,
i.e. β-*effectivity*.

Modal operators for conditional strategic reasoning: O_c

$\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi)$, intuitively meaning:

the agent A has an action σ_{A} such that,
when applied, it guarantees the truth of ϕ
and enables the agent B to apply an action σ_{B}
that guarantees the truth of ψ when additionally applied by B.

(When A and B are coalitions, 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, where A presumably knows the goal of B and can choose to cooperate with B by selecting a suitable action.

When $\psi = \top$, this is equivalent to the Coalition Logic operator $[\mathrm{A}]\,\phi$.

$[\mathrm{A}]\,\phi$ can also be easily expressed by means of each of O_α and O_β.

The logic ConStR formally

Agt: a fixed finite nonempty set of agents,
AP: a fixed countable set of atomic propositions.

Formulae of ConStR, where $p \in \mathrm{AP}$ and $\mathrm{A}, \mathrm{B} \subseteq \mathrm{Agt}$:

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

Some definable strategic operators in ConStR

- The dual operator $[[\mathrm{A}]]_{\mathsf{c}}(\phi; [\mathrm{B}]\psi) := \neg\langle\langle \mathrm{A}\rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B}\rangle\neg\psi)$ says that every (joint) action of A that guarantees the truth of ϕ, would prevent B from acting additionally so as to guarantee ψ.

  Formalises the conditional reasoning with conflicting goals of A and B.

- $[\mathrm{A}]_{\beta}(\phi|\psi) := [\mathrm{A}]_{\beta}(\phi; \langle\emptyset\rangle\psi)$: for any (joint) action of A, if it guarantees ϕ to be true then it guarantees ψ to be true, too.

  This operator formalises the case of Bob's reasoning as an observer.

- $\langle \mathrm{A}\rangle_{\beta}(\phi|\psi) := \neg[\mathrm{A}]_{\beta}(\phi|\neg\psi)$: there is a joint action of A that guarantees ϕ to be true and enables ψ to be true, too.

  $\langle \mathrm{A}\rangle_{\beta}(\phi|\psi)$ is also definable as $\langle\langle \mathrm{A}\rangle\rangle_{\mathsf{c}}(\phi; \langle\overline{\mathrm{A}}\rangle\psi)$, where $\overline{\mathrm{A}} = \mathsf{Agt}\setminus\mathrm{A}$.

- The strategic operator $[\mathrm{C}]$ from CL is a special case of the above: $[\mathrm{C}]\phi = \langle \mathrm{C}\rangle_{\beta}(\phi|\top)$ means "C has a joint action to ensure the truth of ϕ"

- Likewise, $[\mathrm{A}]$ is expressible as $[\mathrm{C}]\phi = [\emptyset]_{\alpha}(\top; \langle \mathrm{C}\rangle\phi)$.

Axiomatic system and decidability for ConStR

The axiomatic system for ConStR presented in [G. & Ju, JoLLI'2022] involves a list of axioms for each of the strategic operators O_c, O_α, O_β, as well as some interaction axioms.

The completeness proof for that axiomatic system is currently under development.

ConStR has the *bounded tree-model property*, therefore its decidability can be proved by a model-theoretic argument, not relying on the completeness theorem.

Addendum: technical details

Concurrent game models formally

A concurrent game model (CGM) is a tuple

$$\langle \mathbb{A}, \mathsf{St}, \mathsf{Act}, \mathsf{act}, \mathsf{out}, \mathrm{Prop}, \mathsf{L} \rangle$$

where:

- $\mathbb{A}$ is a finite set of agents (players)
- St is a set of system states
- Act is a set of possible actions
- $\mathsf{act} : \mathbb{A} \times \mathsf{St} \to \mathcal{P}(\mathsf{Act})$ – mapping assigning to every agent $\mathbf{i}$ and state s a set $\mathsf{act}(\mathbf{i}, s)$ of actions available to $\mathbf{i}$ at s
- out is the outcome function which, for every state and a tuple of available actions, one for each agent, determines the successor state
- Prop is the set of atomic propositions
- $\mathsf{L} : \mathsf{St} \to \mathcal{P}(\mathrm{Prop})$ is the labeling (state description) function.

The set of possible outcomes from a joint action Ordered join of coalitional actions

Consider a CGM $\mathcal{M} = \langle \mathbb{A}, \mathsf{St}, \mathsf{Act}, \mathsf{act}, \mathsf{out}, \mathrm{Prop}, \mathsf{L} \rangle$.

Given a coalition $C \subseteq \mathsf{Agt}$, a joint action for C in $\mathcal{M}$ is a tuple of individual actions $\sigma_C \in \mathsf{Act}^C$. For any such joint action and state $s \in S$ such that σ_C is available at s, we define its set of possible outcomes:

$$\mathsf{Out}[s, \sigma_C] = \{u \in S \mid \exists \sigma \in \Sigma_s : \ \sigma|_C = \sigma_C \text{ and } \mathsf{out}(s, \sigma) = u\}$$

where $\sigma|_C$ is the restriction of σ to C.

Given coalitions A and B and their joint actions σ_{A} and σ_{B}, the ordered join of σ_{A} and σ_{B} is the joint action $\sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}$ of $\mathrm{A} \cup \mathrm{B}$, defined as follows:
all agents in A act according to σ_{A},
and those in $\mathrm{B} \setminus \mathrm{A}$ act according to σ_{B}.

Formal semantics of ConStR

We define inductively truth of a formula at the state $s \in \mathsf{St}$ in a CGM $\mathcal{M}$:

$\mathcal{M}, s \Vdash p$ iff $p \in \mathsf{L}(s)$; $\mathcal{M}, s \Vdash \top$;

$\mathcal{M}, s \Vdash \neg\phi$ iff $\mathcal{M}, s \nVdash \phi$; $\mathcal{M}, s \Vdash \phi \wedge \psi$ iff $\mathcal{M}, s \Vdash \phi$ and $\mathcal{M}, s \Vdash \psi$;

$\mathcal{M}, s \Vdash \langle\!\langle \mathrm{A} \rangle\!\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi)$ iff 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 $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}]$.

$\mathcal{M}, s \Vdash [\mathrm{A}]_{\beta}(\phi; \langle \mathrm{B} \rangle \psi)$ iff 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 $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}]$.

$\mathcal{M}, s \Vdash [\mathrm{A}]_{\alpha}(\phi; \langle \mathrm{B} \rangle \psi)$ iff 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 $u \in \mathsf{Out}[s, \sigma_{\mathrm{A}} \uplus \sigma_{\mathrm{B}}]$.

Axiomatic system $\mathsf{Ax}_{\mathsf{ConStR}}$: axioms and rules for O_c

I. Common axiom schemes and rules for ConStR: the axioms of Coalition Logic, expressed by each of the operators O_c, O_α, O_β.

II. Additional axiom schemes and rules for O_c:

(O_c1) Monotonicity w.r.t. A:
$\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \rightarrow \langle\langle \mathrm{A} \cup \mathrm{C} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi)$ for any $\mathrm{C} \subseteq \mathsf{Agt}$

(O_c2) Monotonicity w.r.t. B:
$\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \rightarrow \langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \cup \mathrm{C} \rangle \psi)$ for any $\mathrm{C} \subseteq \mathsf{Agt}$

(O_c3) $\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \rightarrow \langle\langle \mathrm{A} \cup \mathrm{B} \rangle\rangle_{\mathsf{c}}((\phi \wedge \psi); \langle \emptyset \rangle \top)$

(O_c4) $\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \emptyset \rangle \psi) \leftrightarrow \langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}((\phi \wedge \psi); \langle \emptyset \rangle \top)$

(O_c5) $\neg \langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\bot; \langle \mathrm{B} \rangle \psi)$

(O_c6) $\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \leftrightarrow \langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \setminus \mathrm{A} \rangle \, \psi)$

(O_c7) $\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \leftrightarrow \langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle (\phi \wedge \psi))$

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

$$\frac{\phi \rightarrow \phi', \ \psi \rightarrow \psi'}{\langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B} \rangle \psi) \rightarrow \langle\langle \mathrm{A} \rangle\rangle_{\mathsf{c}}(\phi'; \langle \mathrm{B} \rangle \psi')}$$

Axiomatic system $\mathsf{Ax}_{\mathsf{ConStR}}$: axioms and rules for O_β

III. Additional axiom schemes and rules for O_β:

(O_β1) Monotonicity w.r.t. B:
$[\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \rangle \psi) \rightarrow [\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \cup \mathrm{C} \rangle \psi)$ for any $\mathrm{C} \subseteq$ Agt.

(O_β2) $[\mathrm{A}]_\beta(\phi; \langle \emptyset \rangle \phi)$

(O_β3) $[\mathrm{A}]_\beta(\bot; \langle \emptyset \rangle \psi)$

(O_β4) $[\emptyset]_\beta(\top; \langle \mathrm{A} \rangle \phi) \rightarrow \neg [\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \rangle \bot)$

(O_β5) $[\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \rangle \psi) \leftrightarrow [\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \setminus \mathrm{A} \rangle \psi)$

(O_β6) $[\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \rangle \psi) \leftrightarrow [\mathrm{A}]_\beta(\phi; \langle \mathrm{B} \rangle (\phi \wedge \psi))$

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

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

Axiomatic system $\mathsf{Ax}_{\mathsf{ConStR}}$: axioms and rules for O_α

IV. Additional axiom schemes and rules for O_α:

All axioms (O_β1) - (O_β6), rewritten for O_α. In addition:

(O_α*) Anti-monotonicity w.r.t. A:
$[\mathrm{A} \cup \mathrm{C}]_\alpha(\phi; \langle \mathrm{B} \rangle \psi) \to [\mathrm{A}]_\alpha(\phi; \langle \mathrm{B} \rangle \psi)$ for any $\mathrm{C} \subseteq$ Agt.

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

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

Axiomatic system $\mathrm{Ax}_{\mathsf{ConStR}}$: interacting axioms

V. Interacting axioms for ConStR:

(ConStR 1) $[\mathrm{A}]_{\alpha}(\phi; \langle \mathrm{B}\rangle\psi) \to [\mathrm{A}]_{\beta}(\phi; \langle \mathrm{B}\rangle\psi)$

(ConStR 2) $[\emptyset]_{\beta}(\top; \langle \mathrm{A}\rangle\phi) \land [\mathrm{A}]_{\beta}(\phi; \langle \mathrm{B}\rangle\psi) \to \langle\langle \mathrm{A}\rangle\rangle_{\mathsf{c}}(\phi; \langle \mathrm{B}\rangle\psi)$

Closing remarks

This work: about one-level, local conditional strategic reasoning.
Next steps in the project:

- Extend ConStR with standard temporal operators, to produce an ATL-like extension enabling *long term* conditional strategic reasoning.

- Add explicitly knowledge, both in the semantics and in the language.

- Formalise agents' higher-order conditional reasoning, where e.g. Alice and Bob know each other's goals and possible actions, and both reason assuming that the other will act in pursuit of her/his goals taking into account that the other will do likewise, etc.

  This calls for iterated conditional strategic reasoning, with fixed-point based formal semantics. Work in (slow) progress.

THE END


Make Logic,
not War!