Rational Play and Rational Beliefs under Uncertainty

Nils Bulling and Wojciech Jamroga
Clausthal University of Technology, Germany
{bulling,wjamroga}@in.tu-clausthal.de

ABSTRACT

Alternating-time temporal logic (ATL) is one of the most influential logics for reasoning about agents' abilities. Constructive Strategic Logic (CSL) is a variant of ATL for imperfect information games that allows to express strategic and epistemic properties of coalitions under uncertainty. In this paper, we propose a logic that extends CSL with a notion of plausibility that can be used for reasoning about the outcome of rational behavior (in the game-theoretical sense). Moreover, we show how a particular notion of beliefs can be defined on top of plausibility. The resulting logic, CSLP, turns out to be very expressive.

We show that beliefs satisfy axioms **KD45** in the logic. We also demonstrate how solution concepts for imperfect information games can be characterized and used in CSLP and that the model checking complexity increases only slightly when plausibility and rational beliefs are added.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*; I.2.4 [**Artificial Intelligence**]: Knowledge Representation Formalisms and Methods—*Modal logic*

General Terms

Theory

Keywords

Temporal logic, imperfect information games, knowledge and beliefs

1. INTRODUCTION

Alternating-time temporal logic (ATL) [1] is one of the most influential logics for reasoning about abilities of agents with perfect information. The key constructs are *cooperation modalities* $\langle\langle A\rangle\rangle$ where A is a group of agents. The reading of $\langle\langle A\rangle\rangle\gamma$ is that agents A have a collective strategy to enforce γ. In [7] a variant of ATL for imperfect information scenarios has been proposed. The logic, called *Constructive Strategic Logic* (CSL), unified several attempts to incorporate epistemic concepts into ATL, and solved various problems of these previous attempts. However, it included only strategic and epistemic modalities; in particular, doxastic and rationality concepts were absent.

On the other hand, another extension of ATL, called ATL *with Plausibility* (ATLP), has been proposed in [3] for reasoning about *rational* or *plausible* behavior.[1] That logic allowed to describe and/or impose rationality assumptions on (a subset of) agents, and to reason about the outcome of the play if irrational behavior was disregarded. For example, one might assume that agents are not completely dumb and do not play dominated strategies (in the game-theoretical sense). Such assumptions allow to restrict the vast number of possibilities each agent has to consider.

In this paper we present *Constructive Strategic Logic with Plausibility* (CSLP), a combination of CSL and ATLP where the new language goes far beyond the pure union of both logics. Firstly, the plausibility concept allows us to neatly define the relationship between epistemic and doxastic concepts, in a similar way to [2]. As the basic modality we introduce *weak constructive rational beliefs*: $\mathbb{CW}_A$ (common beliefs), $\mathbb{DW}_A$ (distributed beliefs), and $\mathbb{EW}_A$ (mutual beliefs). The term *constructive* is used in the same sense as in [3], where it referred to an "operational" kind of knowledge that, in order to "know how to play", requires the agents to be able to *identify* and *execute* an appropriate strategy. Like for CSL, the semantics of CSLP is non-standard: formulae are interpreted in *sets of states*. For example, the intuitive reading of $M, Q \models \langle\langle A\rangle\rangle\gamma$ is that agents A have a collective strategy which enforces γ from *each* state in Q. Thanks to the plausibility concept provided by ATLP we can define *knowledge* and *rational beliefs* on top of weak beliefs. We point out that our notion of rational belief is rather specific, and show interesting properties of knowledge, rational belief, and plausibility. In particular, it is shown that knowledge and belief are **KD45** modalities.

We show that CSLP is very expressive, and we demonstrate how solution concepts for imperfect information games can be characterized and used in CSLP. It also turns out that, despite the logic's expressiveness, the model checking complexity does not increase when compared to ATLP, and increases only slightly compared to CSL when plausibility and rational beliefs are added.

1.1 Related Work

Our idea to build beliefs on top of plausibility has been in-

Cite as: Rational Play and Rational Beliefs under Uncertainty, Nils Bulling, Wojciech Jamroga, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 257–264

[1] In this paper we use the terms *rational* and *plausible* interchangeably.

spired by [10, 6]. In [2], we extended CTLK, a straightforward combination of the branching-time logic CTL [4] and standard epistemic logic [5], by a notion of plausibility which in turn was used to define a particular notion of beliefs. Plausibility assumptions were defined in terms of paths in the underlying system. Then, agent's beliefs were given by his knowledge if only plausible paths were considered.

Another source of inspiration is [13, 12], where the semantics of ability was influenced by particular notions of rationality. We generalized these ideas in [3]. Semantically, a subset of strategies (behaviors) was identified as *rational* in the model; a typical formula was $\mathbf{Pl}_B\langle\langle A\rangle\rangle\gamma$ with the following reading: Agents A can enforce γ if agents in B act rationally. We showed how one can use the logic to characterize solution concepts (Nash equilibria, Pareto optimal profiles etc.), and reason about the outcome of rational play.

The current paper is an attempt to integrate the notions of time, knowledge, belief, strategic ability, rationality, and uncertainty in a single logical framework.

2. CSL WITH PLAUSIBILITY

We start with an informal presentation of the idea. Then, we describe the formal syntax and semantics, and we discuss the new operators in more detail.

2.1 Agents, Beliefs, and Rational Play

In the following, let $A \subseteq \mathbb{A}\mathrm{gt}$ be a team of agents where $\mathbb{A}\mathrm{gt}$ denotes the set of all agents. Formulae are interpreted given a model M and a set of states Q. The reading of $M, Q \models \langle\langle A\rangle\rangle\gamma$ is that agents A have a collective strategy which enforces γ from *all* states in Q. $\mathbf{Pl}_A\,\varphi$ assumes that agents in A play plausibly according to some rationality criterion which can be set (resp. refined) by operators (**set-pl** ω) (resp. (**refn-pl** ω)). The set of such *rational agents* is denoted by $\mathbb{R}\mathrm{gt}$. Plausibility terms ω refer to sets of strategy profiles that implement the rationality criteria. Finally, the logic includes operators for *constructive weakly rational belief* (*constructive weak belief*/cwb in short): $\mathbb{CW}_A\varphi$ (agents A have common cwb in φ); $\mathbb{EW}_A\varphi$ (agents A have mutual cwb in φ); and $\mathbb{DW}_A\varphi$ (agents A have distributed cwb in φ). Semantically, the cwb operators yield "epistemic positions" of team A that serve as reference for the semantic evaluation of strategic formulae.

Consider formula $\mathbb{EW}_A\mathbf{Pl}_{\mathbb{A}\mathrm{gt}\setminus A}\,\langle\langle A\rangle\rangle\Box\mathsf{safe}$ (*coalition A has a constructive mutual weak belief that they can keep the system safe forever if the opponents behave rationally*) in model M and set of states Q. Firstly, Q is extended with all states indistinguishable from some state in Q for any agent from A. Let us call the extended set Q'. Now, A have cwb in $\mathbf{Pl}_{\mathbb{A}\mathrm{gt}\setminus A}\,\langle\langle A\rangle\rangle\Box\mathsf{safe}$ iff they have a strategy that maintains safe from all states in Q' assuming that implausible behavior for the agents in $\mathbb{A}\mathrm{gt}\setminus A$ is disregarded.

Later, we will define strongly rational beliefs (resp. knowledge) as a special case of cwb's in which all agents are (resp. no agent is) assumed to play plausibly.

2.2 Syntax

The language of *Constructive Strategic Logic with Plausibility* (CSLP) includes atomic propositions, Boolean connectives, strategic formulae, operators for *constructive weakly rational beliefs*, and operators that handle *plausibility updates.* As we will see, standard/constructive strongly rational beliefs and knowledge can be defined on top of these.

Definition 1 ($\mathcal{L}_{\mathrm{CSLP}}$). *Let $\mathbb{A}\mathrm{gt}$ be a set of agents, Π a set of propositions, and Ω a set of primitive plausibility terms. The logic $\mathcal{L}_{CSLP}(\mathbb{A}\mathrm{gt},\Pi,\Omega)$ is generated by the following grammar:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi\wedge\varphi \mid \langle\langle A\rangle\rangle\bigcirc\varphi \mid \langle\langle A\rangle\rangle\Box\varphi \mid \langle\langle A\rangle\rangle\varphi\,\mathcal{U}\,\varphi \mid \mathbb{CW}_A\varphi \mid \mathbb{EW}_A\varphi \mid \mathbb{DW}_A\varphi \mid \mathbf{Pl}_A\,\varphi \mid (\textbf{\textit{set-pl}}\ \omega)\varphi \mid (\textbf{\textit{refn-pl}}\ \omega)\varphi.$$

The temporal operators $\bigcirc, \Box, \mathcal{U}$ stand for "next", "always", and "until", respectively. We use the standard definitions of $\vee, \rightarrow, \leftrightarrow$, plus the following derived modalities: $\Diamond\varphi \equiv \top\,\mathcal{U}\,\varphi$ (sometime), $\mathsf{Now}\varphi \equiv \varphi\,\mathcal{U}\,\varphi$ (now), $\mathbb{W}_a\varphi \equiv \mathbb{CW}_{\{a\}}\varphi$ (individual cwb), $\mathsf{CW}_A\varphi \equiv \mathbb{CW}_A\langle\langle\emptyset\rangle\rangle\mathsf{Now}\varphi$, $\mathsf{EW}_A\varphi \equiv \mathbb{EW}_A\langle\langle\emptyset\rangle\rangle\mathsf{Now}\varphi$, $\mathsf{DW}_A\varphi \equiv \mathbb{DW}_A\langle\langle\emptyset\rangle\rangle\mathsf{Now}\varphi$ (standard weak belief, wb), $\mathsf{W}_a\varphi \equiv \mathsf{CW}_{\{a\}}\varphi$ (individual wb), $\mathbf{Pl} \equiv \mathbf{Pl}_{\mathbb{A}\mathrm{gt}}$ (reasoning under the assumption that all agents behave plausibly), and $\mathbf{Ph} \equiv \mathbf{Pl}_\emptyset$ (reasoning about outcome of all "physically" possible behaviors). Finally, we define operators for constructive and standard strongly rational belief (csb) as:
$\mathbb{B}\mathsf{el}_a \equiv \mathbb{W}_a\mathbf{Pl}$, $\mathbb{CB}\mathsf{el}_A \equiv \mathbb{CW}_A\mathbf{Pl}$,
$\mathbb{EB}\mathsf{el}_A \equiv \mathbb{EW}_A\mathbf{Pl}$, $\mathbb{DB}\mathsf{el}_A \equiv \mathbb{DW}_A\mathbf{Pl}$,
$\mathsf{Bel}_a \equiv \mathbf{Ph}\,\mathsf{W}_a\mathbf{Pl}$, $\mathsf{CBel}_A \equiv \mathbf{Ph}\,\mathsf{CW}_A\mathbf{Pl}$,
$\mathsf{EBel}_A \equiv \mathbf{Ph}\,\mathsf{EW}_A\mathbf{Pl}$, $\mathsf{DBel}_A \equiv \mathbf{Ph}\,\mathsf{DW}_A\mathbf{Pl}$,
and the constructive and standard knowledge operators as:
$\mathbb{K}_a \equiv \mathbf{Ph}\,\mathbb{W}_a$, $\mathbb{C}_A \equiv \mathbf{Ph}\,\mathbb{CW}_A$, $\mathbb{E}_A \equiv \mathbf{Ph}\,\mathbb{EW}_A$,
$\mathbb{D}_A \equiv \mathbf{Ph}\,\mathbb{DW}_A$, $\mathsf{K}_a \equiv \mathbf{Ph}\,\mathsf{W}_a$, $\mathsf{C}_A \equiv \mathbf{Ph}\,\mathsf{CW}_A$,
$\mathsf{E}_A \equiv \mathbf{Ph}\,\mathsf{EW}_A$, $\mathsf{D}_A \equiv \mathbf{Ph}\,\mathsf{DW}_A$.

We will show in Section 2.4 that these definitions capture the respective notions of knowledge and belief appropriately.

2.3 Semantics

Definition 2 (CEGS). *A* concurrent epistemic game structure *is a tuple $M = \langle\mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o, \sim_1, \ldots, \sim_k\rangle$, with: a nonempty finite set of all agents $\mathbb{A}\mathrm{gt} = \{1, ..., k\}$, a nonempty set of states St, a set of atomic propositions Π, a valuation of propositions $\pi : St \rightarrow 2^\Pi$, and a nonempty finite set of atomic actions Act. $\sim_1, ..., \sim_k \subseteq St\times St$ are epistemic equivalence relations; $q \sim_a q'$ means that, while the system is in state q, agent a cannot determine whether it is in q or q'. Function $d : \mathbb{A}\mathrm{gt}\times St \rightarrow 2^{Act}$ defines nonempty sets of actions available to agents at each state, with $d(a,q) = d(a,q')$ for $q \sim_a q'$. Finally, o is a (deterministic) transition function that assigns the outcome state $q' = o(q, \alpha_1, \ldots, \alpha_k)$ to state q and a tuple of actions $\langle\alpha_1, \ldots, \alpha_k\rangle$, $\alpha_i \in d(i,q)$, that can be executed by $\mathbb{A}\mathrm{gt}$ in q.*

Remark 1. *Relations $\sim^E_A$, $\sim^C_A$ and $\sim^D_A$, used to model group epistemics, are derived from the individual relations of agents from A. First, $\sim^E_A$ is the union of relations $\sim_a$, $a \in A$. Next, $\sim^C_A$ is defined as the transitive closure of $\sim^E_A$. Finally, $\sim^D_A$ is the intersection of all the $\sim_a$, $a \in A$.*

A *strategy* s_a of agent a is a conditional plan that specifies what a is going to do for every possible situation: $s_a : St \rightarrow Act$ such that $s_a(q) \in d(a,q)$. A *collective strategy* s_A for a group of agents A is a tuple of strategies, one per agent from A. Strategy s_a is *uniform* iff $q \sim_a q'$ implies $s_a(q) = s_a(q')$; a collective strategy is uniform iff it consists of only uniform individual strategies. We denote the set of uniform strategies of agent a by Σ_a; the set of uniform collective strategies of team A is given by $\Sigma_A = \times_{a\in A}\Sigma_a$, and the set of all uniform *strategy profiles* by $\Sigma = \Sigma_{\mathbb{A}\mathrm{gt}}$.

Definition 3 (CEGSP, plausibility model). *A* concurrent epistemic game structure with plausibility *is given by*

$M = \langle \mathbb{A}\text{gt}, St, \Pi, \pi, Act, d, o, \sim_1, \ldots, \sim_k, \Upsilon, \mathbb{R}\text{gt}, \Omega, [\![\cdot]\!]\rangle$, *where* $\langle \mathbb{A}\text{gt}, St, \Pi, \pi, Act, d, o, \sim_1, \ldots, \sim_k\rangle$ *is a* CEGS, $\Upsilon \subseteq \Sigma$ *is a set of plausible strategy profiles (called* plausibility set*),* $\mathbb{R}\text{gt} \subseteq \mathbb{A}\text{gt}$ *is a set of* rational agents *(i.e., the agents to whom the plausibility assumption will apply),* Ω *is a set of plausibility terms, and* $[\![\cdot]\!] : \Omega \times 2^{St} \to \Sigma$ *is a* plausibility mapping *that provides denotation of the terms.*[2] *We refer to* $(\Upsilon, \mathbb{R}\text{gt})$ *as the* plausibility model *of* M*. When necessary, we write* X_M *to denote the element* X *of model* M*.*

Note that imposing strategic restrictions on a *subset* $\mathbb{R}\text{gt}$ of agents can be desirable due to several reasons. It might, for example, be the case that only information about the proponents' play is available; hence, assuming plausible behavior of the opponents is neither sensible nor justified. Or, even simpler, a group of (simple minded) agents might be known not to behave rationally.

Consider now formula $\langle\!\langle A\rangle\!\rangle\gamma$: The team A looks for a strategy that brings about γ, but the members of the team who are also in $\mathbb{R}\text{gt}$ can only choose plausible strategies. The same applies to A's opponents that are contained in $\mathbb{R}\text{gt}$.

Definition 4 (Plausibility of strategies). *Let* $s_A|_B$ *be the* $(A \cap B)$*'s substrategy of* s_A*, and* $\Upsilon|_B = \{s_B \in \Sigma_B \mid \exists s \in \Upsilon\ s|_B = s_B\}$*. We say that* s_A *is* plausible *iff* $\mathbb{R}\text{gt}$*'s substrategy in* s_A *is part of some strategy profile in* Υ*, i.e., if* $s_A|_{A\cap\mathbb{R}\text{gt}} \in \Upsilon|_{A\cap\mathbb{R}\text{gt}}$*.*

By Σ^* *we denote the set of all plausible strategy profiles in the model. That is,* $\Sigma^* = \{s \in \Sigma \mid s|_{\mathbb{R}\text{gt}} \in \Upsilon|_{\mathbb{R}\text{gt}}\}$*. Note that* s_A *is plausible iff* $s_A \in \Sigma^*|_A$*.*

A *path* $\lambda = q_0 q_1 \cdots \in St^\omega$ is an infinite sequence of states such that there is a transition between each q_i, q_{i+1}. By $\lambda[i] = q_i$ we denote the i-th state of λ. Λ denotes all paths in the model, and $\Lambda(q)$ the set of all paths starting in q.

Definition 5 (Plausible outcome paths). *The* plausible outcome*,* $out(q, s_A)$*, of strategy* s_A *from state* q *is defined as the set of paths (starting from* q*) which can occur when only plausible strategy profiles can be played and agents in* A *follow* s_A*; that is,* $out(q, s_A) = \{\lambda \in \Lambda(q) \mid \exists t \in \Sigma^*\ t|_A = s_A \text{ and } out(q,t) = \{\lambda\}\}$

Now we define the notion of *formula* φ *being satisfied* by a (non-empty) set of states Q in model M, written $M, Q \models \varphi$. We will also write $M, q \models \varphi$ as a shorthand for $M, \{q\} \models \varphi$. Note that it is the latter notion of satisfaction (in single states) that we are ultimately interested in – but it is defined in terms of the (more general) satisfaction in sets of states. Let $\text{img}(q, \mathcal{R})$ be the image of state q with respect to binary relation $\mathcal{R}$, i.e., the set of all states q' such that $q\mathcal{R}q'$. Moreover, we use $out(Q, s_A)$ as a shorthand for $\bigcup_{q\in Q} out(q, s_A)$, and $\text{img}(Q, \mathcal{R})$ as a shorthand for $\bigcup_{q\in Q}\text{img}(q, \mathcal{R})$. The semantics is given through the following clauses.

$M, Q \models p$ iff $p \in \pi(q)$ for every $q \in Q$;

$M, Q \models \neg\varphi$ iff $M, Q \not\models \varphi$;

$M, Q \models \varphi \wedge \psi$ iff $M, Q \models \varphi$ and $M, Q \models \psi$;

$M, Q \models \langle\!\langle A\rangle\!\rangle \bigcirc \varphi$ iff there exists $s_A \in \Sigma^*|_A$ such that, for every $\lambda \in out(Q, s_A)$, we have that $M, \{\lambda[1]\} \models \varphi$;

$M, Q \models \langle\!\langle A\rangle\!\rangle \Box \varphi$ iff there exists $s_A \in \Sigma^*|_A$ such that, for every $\lambda \in out(Q, s_A)$ and $i \geq 0$, we have $M, \{\lambda[i]\} \models \varphi$;

$M, Q \models \langle\!\langle A\rangle\!\rangle \varphi \mathcal{U} \psi$ iff there exists $s_A \in \Sigma^*|_A$ such that, for every $\lambda \in out(Q, s_A)$, there is an $i \geq 0$ for which $M, \{\lambda[i]\} \models \psi$ and $M, \{\lambda[j]\} \models \varphi$ for every $0 \leq j < i$.

$M, Q \models \hat{\mathcal{K}}\mathbb{W}_A\varphi$ iff $M, \text{img}(Q, \sim_A^{\mathcal{K}}) \models \varphi$ (where $\hat{\mathcal{K}} = \mathbb{C}, \mathbb{E}, \mathbb{D}$ and $\mathcal{K} = C, E, D$, respectively).

$M, Q \models \mathbf{Pl}_A\, \varphi$ iff $M', Q \models \varphi$, where the new model M' is equal to M but the new set $\mathbb{R}\text{gt}_{M'}$ of rational agents in M' is set to A.

$M, Q \models (\mathbf{set\text{-}pl}\ \omega)\varphi$ iff $M', Q \models \varphi$ where M' is equal to M with $\Upsilon_{M'}$ set to $[\![\omega]\!]_M^Q$.

$M, Q \models (\mathbf{refn\text{-}pl}\ \omega)\varphi$ iff $M', Q \models \varphi$ where M' is equal to M with $\Upsilon_{M'}$ set to $\Upsilon_M \cap [\![\omega]\!]_M^Q$.

Like in CSL, we use two notions of validity, *weak* and *strong*, depending on whether formulae are evaluated with respect to single states or sets of states.

Definition 6 (Validity). *We say that* φ *is* valid *if* $M, q \models \varphi$ *for all* CEGSP*'s* M *with plausibility model* $(\Sigma, \emptyset)$ *(i.e. all strategies are assumed to be plausible and no agent plays plausibly yet) and all states* $q \in St_M$*.*

In addition to that, we say that φ *is* strongly valid *if* $M, Q \models \varphi$ *for* all CEGSP*'s* M *and all sets of states* $Q \subseteq St_M$*.*

Note that strong validity is interpreted in *all* models and not only in those with plausibility model $(\Sigma, \emptyset)$. This stronger notion is necessary for interchangeability of (sub)formulae. The following results are straightforward.

Proposition 2. *Strong validity implies validity.*

Proposition 3. *If* $\varphi_1 \leftrightarrow \varphi_2$ *is strongly valid, and* ψ' *is obtained from* ψ *through replacing an occurrence of* φ_1 *by* φ_2*, then* $M, Q \models \psi$ *iff* $M, Q \models \psi'$*.*

We also say that φ is *satisfiable* if $M, q \models \varphi$ for some CEGSP with plausibility model $(\Sigma, \emptyset)$.

2.4 Interpretation of Derived Operators

In this section we motivate the logic's epistemic and doxastic operators. We especially show that the syntactic definitions for the derived knowledge and belief operators have an intuitive semantics.

2.4.1 Knowledge

The concept behind knowledge is very simple: It is about everything which is "physically" possible, i.e., *all* behaviors are taken into account (not only the plausible ones). In particular this means that, once a knowledge operator occurs, the set of rational agents in the plausibility model becomes void, indicating that *no* agent is assumed to play rationally.

2.4.2 Weakly and Strongly Rational Beliefs

Constructive weak beliefs (cwb) ("common belief", "distributed belief", and "mutual belief") are primitive operators in our logic. All other belief/knowledge operators are derived from cwb and plausibility. In this section, we mainly discuss individual knowledge and beliefs, but the analysis extends to collective attitudes in a straightforward way.

Let us for example consider the individual cwb operator $\mathbb{W}_a\varphi$, with the following reading: Agent a has *constructive*

[2]In this section, the denotation of such terms is fixed; in Section 4 we present a more flexible version.

weak belief in φ iff φ holds in all states that a considers possible, where all agents behave according to the currently specified plausibility model (Υ, A). That is, agents in A are assumed to play as specified in Υ. It is important to note that *weakly rational* beliefs restrict *only* the behavior of the agents specified in the current plausibility model (i.e. A). This is the difference between weak and strong beliefs – the latter assume plausible behavior of *all* the agents. This is why we call such beliefs *strongly rational*, as it restricts the behavior of the system in a more rigorous way due to stronger rationality assumptions.

Using rationality assumptions to define beliefs makes them rather specific. They differ from most "standard" concepts of belief in two main respects. Firstly, our notion of beliefs is focused on *behavior* and *abilities* of agents. When no action is considered, all epistemic and doxastic notions coincide.

PROPOSITION 4. *Let φ be a propositional formula. Then, $\mathbb{W}_a\varphi \leftrightarrow \mathbb{B}\mathsf{el}_a\varphi \leftrightarrow \mathbb{K}_a\varphi$ is strongly valid.*

Secondly, rational beliefs are about *restricting the expected behavior* due to rationality assumptions: Irrational behaviors are simply disregarded. To strengthen this important point consider the following statements:

(i) *Ann (a) knows how Bill (b) can commit suicide* (which can be formalized as $\mathbb{K}_a\langle\langle b\rangle\rangle\Diamond\mathsf{suicide}$);

(ii) *Ann constructively believes that Bill can commit suicide* (which we tentatively formalize as $\mathbb{B}\mathsf{el}_a\langle\langle b\rangle\rangle\Diamond\mathsf{suicide}$).

In the usual treatment of beliefs, statement (i) should imply statement (ii), but this does not apply to *rational* beliefs. That is because, typically, beliefs and knowledge are both about "hard facts". Thus, if a knows some fact to be true, she should also include it in her belief base. On the other hand, our reading of $\mathbb{B}\mathsf{el}_a\langle\langle b\rangle\rangle\Diamond\mathsf{suicide}$ is given as follows: If all agents are constrained to act rationally then Ann knows a strategy for Bill by which he can commit suicide. However, it is natural to assume that no rational entity would commit suicide.[3] Hence, Bill's ability to commit suicide is out of question if we assume him to act rationally. Such an irrational behavior is just unthinkable and thus disregarded by Ann! While she knows how Bob can commit suicide in general, she has no *plausible* recipe for Bob to do that.

A similar analysis can be conducted for standard (i.e., non-constructive) beliefs. Consider the following variants of (i) and (ii):

(i') *Ann knows that Bill has some way of committing suicide* ($\mathsf{K}_a\langle\langle b\rangle\rangle\Diamond\mathsf{suicide}$);

(ii') *Ann believes, taking only rational behavior of all agents into account (in particular of Bill), that Bill has the ability to commit suicide* ($\mathsf{Bel}_a\langle\langle b\rangle\rangle\Diamond\mathsf{suicide}$).

Like before, (i') does not imply (ii'). While Ann knows that Bill "physically" has some way of killing himself, by assuming him to be rational she disregards the possibility. Bob's assumed rationality constrains his choices in Ann's view. This shows that in our logic knowing φ does not imply rational beliefs in φ. We will justify the intuition on a more concrete example.

[3]This assumption is given in the plausibility model; it can be any assumption the designer would like to impose on the agents.

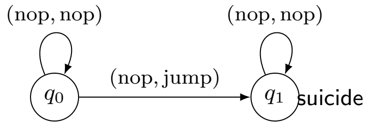

Figure 1: Simple CEGSP.

EXAMPLE 1. *There are two agents 1 (Ann) and 2 (Bill). Agent 2 has the ability to jump from a building and commit suicide. However, agent 1 disregards this possibility and considers it rational that 2 will not jump. The corresponding* CEGSP *is shown in Figure 1 where all different states are distinguishable from each other; the set of plausible strategy profiles consists of the single profile s in which both agents play action nop, i.e., they do nothing (in particular, we want to impose that Bill does not jump). Hence, we have $M, q_0 \models \mathsf{K}_1\langle\langle 2\rangle\rangle\bigcirc$ suicide but $M, q_0 \not\models \mathsf{Bel}_1\langle\langle 2\rangle\rangle\bigcirc$ suicide.*

The following result, in line with [2], is immediate:

THEOREM 5. *In general, standard (resp. constructive) knowledge does not imply standard (resp. constructive) rational belief. That is, formulae $\mathbb{K}_a\varphi\wedge\neg\mathbb{B}\mathsf{el}_a\varphi$, $\mathbb{K}_a\varphi\wedge\neg\mathbb{W}_a\varphi$, $\mathsf{K}_a\varphi\wedge\neg\mathsf{Bel}_a\varphi$, $\mathsf{K}_a\varphi\wedge\neg\mathsf{W}_a\varphi$ are satisfiable.*

2.4.3 Non-Constructive Knowledge and Beliefs

In this section, we have a closer look at the standard (non-constructive) epistemic and doxastic operators. We mainly focus on strong beliefs; the cases for knowledge and weak beliefs are given analogously.

The non-constructive versions of distributed, common, and everybody belief are based on a specific construction involving the "until" operator. For example, the non-constructive belief of agent a in φ, $\mathsf{Bel}_a\varphi$, is defined as a's *constructive* belief in the ability of the empty coalition to enforce φ until φ. In [7] it was already shown that this definition captures the right notion; we recall the intuition here.

The cooperation modality $\langle\langle\emptyset\rangle\rangle$ ensures that the state formula φ is evaluated *independently* in each indistinguishable state in Q (thus getting rid of its constructive flavour). However, a cooperation modality must be followed directly by a path formula, and φ is a state formula. The trick is to use $\varphi\,\mathcal{U}\,\varphi$ instead, which ensures that φ is true in the initial state of the path. Thus, a believes in φ iff $\mathbf{Pl}\,\varphi$ is independently true in every indistinguishable state. The following proposition (analogous to [7, Theorem 46]) states that all non-constructive operators match their intended intuitions.

PROPOSITION 6. *Let M be a* CEGSP, *$q \in St_M$, and φ be a* CSLP *formula. Then the following holds, where $\mathcal{K} = C, E, D$, respectively:*

1. *$M, Q \models \mathcal{K}\mathsf{W}_A\varphi$ iff $\Upsilon_M \neq \emptyset$ and $M, q \models \varphi$ for all $q \in \mathrm{img}(Q, \sim_A^{\mathcal{K}})$;*
2. *$M, Q \models \mathcal{K}\mathsf{Bel}_A\varphi$ iff $M, q \models \mathbf{Pl}\,\varphi$ for all $q \in \mathrm{img}(Q, \sim_A^{\mathcal{K}})$;*
3. *$M, Q \models \mathcal{K}_A\varphi$ iff $M, q \models \mathbf{Ph}\,\varphi$ for all $q \in \mathrm{img}(Q, \sim_A^{\mathcal{K}})$.*

3. PROPERTIES OF CSLP

In this section, we examine the relationship between plausibility, knowledge and beliefs, and discuss the standard axioms about epistemic and doxastic concepts.

3.1 Plausibility, Knowledge and Beliefs

Firstly, we observe that knowledge is commutative with **Ph** and belief with **Pl**, which is a technically important property.

PROPOSITION 7. *Let φ be a* CSLP *formula. Then, we have that* $\mathbf{Ph}\,\mathbb{K}_a\varphi \leftrightarrow \mathbb{K}_a\mathbf{Ph}\,\varphi$ *and* $\mathbf{Pl}\,\mathbb{B}\mathsf{el}_a\varphi \leftrightarrow \mathbb{B}\mathsf{el}_a\mathbf{Pl}$ *are strongly valid.*

From the definition of knowledge an belief it follows that a sequence of such operators collapses to the final operator in the sequence.

PROPOSITION 8. *Let $a \in \mathbb{A}\mathrm{gt}$, φ be a* CSLP *formula, and X, Y be sequences of belief/knowledge operators; i.e. $X, Y \in \{\mathbb{B}\mathsf{el}_a, \mathbb{K}_a\}^*$. Then the following formulae are strongly valid:*

(i) $X\mathbb{B}\mathsf{el}_a\varphi \leftrightarrow Y\mathbb{B}\mathsf{el}_a\varphi$ *(ii)* $X\mathbb{K}_a\varphi \leftrightarrow Y\mathbb{K}_a\varphi$

In particular, we have that the following formulae are strongly valid: (1) $\mathbb{K}_a\mathbb{B}\mathsf{el}_a\varphi \leftrightarrow \mathbb{B}\mathsf{el}_a\varphi$: Agent a knows that he believes φ iff he believes φ; and (2) $\mathbb{B}\mathsf{el}_a\mathbb{K}_a\varphi \leftrightarrow \mathbb{K}_a\varphi$: Agent a believes that he knows φ iff he knows φ.

PROPOSITION 9. *Let the premises be as in Proposition 8. Then, the following formulae are* not *valid: (i)* $X\mathbb{B}\mathsf{el}_a\varphi \leftrightarrow Y\mathbb{K}_a\varphi$*; (ii)* $\mathbb{B}\mathsf{el}_a\varphi \to \mathbb{B}\mathsf{el}_a\mathbb{K}_a\varphi$*; (iii)* $\mathbb{B}\mathsf{el}_a\varphi \to \mathbb{K}_a\varphi$*.*

Proposition 9 says in particular that (ii) an agent who has rational belief in φ does not necessarily believe that he also knows φ; and (iii) an agent who believes in φ does not necessarily know φ. Indeed, both formulae should not hold in a logics of knowledge and belief.

Our definitions of epistemic and doxastic operators from Section 2.2 strongly suggest that the underlying concepts are related. Let us consider formula $\mathbb{K}_a\mathbf{Pl}_B\varphi$: Agent a has constructive knowledge in φ if agents in B behave rationally. This sounds quite similar to beliefs which is formally shown below.

PROPOSITION 10. $\mathbf{Pl}_A\mathbb{K}_a\mathbf{Pl}_A\varphi \leftrightarrow \mathbf{Pl}_A\mathbb{W}_a\varphi$ *is strongly valid. We also have that* $\mathbb{K}_a\varphi \leftrightarrow \mathbb{W}_a\varphi$ *is valid (but not strongly valid).*

Finally, we conclude that rational beliefs and knowledge can also be defined in terms of each other.

THEOREM 11. $\mathbb{B}\mathsf{el}_a\varphi \leftrightarrow \mathbb{K}_a\mathbf{Pl}\,\varphi$ *and* $\mathbb{K}_a\varphi \leftrightarrow \mathbb{B}\mathsf{el}_a\mathbf{Ph}\,\varphi$ *are strongly valid.*

That is, believing in φ is knowing that φ plausibly holds, and knowing that φ is believing that φ is the case in all physically possible plays.

3.2 Axiomatic Properties

In this section we review the well-known **KDT45** axioms. For modality O these axioms are given as follows:

($\mathbf{K}_O$) $O(\varphi \to \psi) \to (O\varphi \to O\psi)$ ($\mathbf{D}_O$) $O\varphi \to \neg O\neg\varphi$
($\mathbf{T}_O$) $O\varphi \to \varphi$ ($\mathbf{4}_O$) $O\varphi \to OO\varphi$
($\mathbf{5}_O$) $\neg O\varphi \to O\neg O\varphi$

We say, for instance, that O is an **K4** modality if axioms $\mathbf{K}_O$ and $\mathbf{4}_O$ are *strongly valid*. The following result is obtained in a way analogous to [7, Theorem 37].

THEOREM 12 (WEAK BELIEFS: **KD45**). W_a *(standard weak beliefs) and* $\mathbb{W}_a$ *(constructive weak beliefs) are* ***KD45*** *modalities. Axiom* ***T*** *is not valid for both notions of weak beliefs.*

REMARK 13. *Despite the similarities to [2], axiom* ***D*** *was not strongly valid for beliefs in* CTLKP *because the belief operator directly referred to plausible paths. Hence, if the set of paths was empty some formulae were trivially true (*$\mathsf{Bel}\varphi$*) and others are trivially false (*$\neg\mathsf{Bel}\varphi$*). In* CSLP *the notions of belief and plausibility are more modular.*

As knowledge and strong beliefs are special kinds of weak beliefs, both operators have to satisfy the same axioms as the weak belief operator. It just remains to check whether axiom **T** holds for knowledge or strong beliefs. However, for the same reason as in pure CSL this axiom does usually not hold; we refer to [7] for a rigorous discussion of this issue – including ways how axiom **T** can be restored for knowledge. The problem that **T** is not true for knowledge (what is usually assumed to be a sensible requirement) is due to the definition of negation in the non-standard semantics defined wrt sets of states.

THEOREM 14 (STRONG BELIEFS: **KD45**). *Standard strong beliefs* Bel_a *and constructive strong beliefs* $\mathbb{B}\mathsf{el}_a$ *are* ***KD45*** *modalities. Axiom* ***T*** *is not valid for both notions of beliefs.*

THEOREM 15 (KNOWLEDGE: **KD45**). *Standard knowledge* K_a *and constructive knowledge* $\mathbb{K}_a$ *are* ***KD45*** *modalities. Axiom* ***T*** *is not valid for both notions of knowledge.*

Note that if we consider a formula φ which does not contain any constructive operators then the following holds.

THEOREM 16. *Let $\mathcal{L}$ consist of all* CSLP *formulae that contain no constructive operators. Then:*

1. K_a *is a* ***KD45*** *modality in $\mathcal{L}$. Axiom* $\mathbf{T}_{\mathsf{K}_a}$ *is valid (but not strongly valid), and* $\mathbf{Ph}\,(\mathsf{K}_a\varphi \to \varphi)$ *is strongly valid in $\mathcal{L}$.*
2. Bel_a *is a* ***KD45*** *modality and* $\mathbf{Pl}\,(\mathsf{Bel}_a\varphi \to \varphi)$ *is strongly valid in $\mathcal{L}$.*

We observe that the validities $\mathbf{Ph}\,(\mathsf{K}_a\varphi \to \varphi)$ and $\mathbf{Pl}\,(\mathsf{Bel}_a\varphi \to \varphi)$ are very similar to the truth axiom **T**.

3.3 Relationship to Existing Logics

In this section, we compare CSLP with several relevant logics and show their formal relationships. To this end, we define the notion of *embedding*. *Logic L_1 embeds logic L_2* iff there is a translation tr of L_2 formulae into formulae of L_1, and a transformation TR of L_2 models into models of L_1, such that $M, q \models_{L_2} \varphi$ iff $TR(M), q \models_{L_1} tr(\varphi)$ for every pointed model M, q and formula φ of L_2.

The following theorem is straightforward from the definition of the logic.

THEOREM 17. CSLP *embeds* ATL*,* ATLP*, and* CSL*.*

It is easy to see that W_a is even a **KDT45** modality for a sublanguage of CSLP and that this sublanguage can embed standard epistemic propositional logic.

PROPOSITION 18. CSLP *embeds standard epistemic propositional logic.*

The following result is not that obvious but follows from Proposition 18 and [3, Proposition 5].

PROPOSITION 19. CSLP *embeds* CTLKP *in the class of epistemic Kripke structures.*

REMARK 20. *In [7] and [3] it was shown that* CSL *and* ATLP *embed several other logics, e.g.,* ATEL *[11],* ATLI *[8], and* GLP *[13]. Due to Theorem 17 all these logics are also embeddable in* CSLP.

4. FLEXIBLE SPECIFICATIONS

In [3] we showed that ATLP can be used to reason about temporal properties of rational play. In particular it was shown that the logic allows to characterize game theoretic solution concepts of perfect information games [9]. These characterizations were then used to describe agents rational behavior and impose the resulting rationality constraints on them. Here we show that CSLP can be used for the same purpose in the more general case of *imperfect information games* (IIG). A natural question is how solution concepts for both game-types differ?

Actually, they do not differ much. For instance, a Nash equilibrium is a strategy profile from which no agent can deviate to obtain a better payoff, for both the perfect and imperfect information case. However, only *uniform strategies* are considered for IIG. Moreover, we require the agent to *know/identify* a strategy successful in *all* states indistinguishable for him.

Before we present how solution concepts can be described in Section 4.2 we need to pave the way for it: CSLP is not yet expressive enough to *describe* strategies in the object language, only predefined plausibility terms are available.

4.1 Nesting Formulae in CSLP

In this section we present $\mathcal{L}^1_{\text{CSLP}}$ which extends $\mathcal{L}_{\text{CSLP}}$ so that plausibility terms are constructed from $\mathcal{L}_{\text{CSLP}}$ formulae.[4] In the following we proceed in an analogous way to [3]. The *extended plausibility terms* of $\mathcal{L}^1_{\text{CSLP}}$ have a structure similar to $\sigma_1.D(\sigma_1)$. Such a term *selects* all strategy profiles s_1 (referred to by the *strategic variable* σ_1) that satisfy a property D which depends on a given model, set of states, and σ_1. Let us be more precise about the structure of such properties. We allow them to be quantified $\mathcal{L}_{\text{CSLP}}$ formulae, e.g., $D(\sigma_1,\dots,\sigma_n) = \forall\sigma_2\exists\sigma_3\dots\forall\sigma_n\varphi(\sigma_1,\dots,\sigma_n)$, where the quantification takes places over strategy profiles which can be used inside φ in the same way as basic plausibility terms would be used. The variable σ_1 takes on a specific role; it *collects* the "good" strategy profiles.

Before we formally define the language we need one more notation. Solution concepts often require to combine strategies or focus on substrategies. For example, given a term ω_{NE} (describing Nash equilibria) and a term ω_{PO} (describing Pareto optimal strategies) we can use $\langle\omega_{\text{NE}},\omega_{\text{PO}}\rangle$ to refer to all profiles in which agent 1 plays his part of an Nash equilibrium and agent 2 plays a Pareto optimal strategy. Likewise, $\omega_{\text{NE}}[1]$ refers to the strategy profiles in which 1's substrategy is a part of some Nash equilibrium.

Formally, given a non-empty set X we say that y is a *strategic combination of* X if it is generated by the following grammar: $y ::= x \mid \langle y,\dots,y\rangle \mid y[A]$ where $x\in X$, $\langle y,\dots,y\rangle$ is a vector of length $|\mathbb{A}\text{gt}|$, and $A\subseteq\mathbb{A}\text{gt}$. The set of *strategic combinations* over X is defined by $\mathcal{T}(X)$. It is easy to see that operator $\mathcal{T}$ is idempotent ($\mathcal{T}(X)=\mathcal{T}(\mathcal{T}(X))$). Below, we define the language $\mathcal{L}^1_{\text{CSLP}}$.

DEFINITION 7 ($\mathcal{L}^1_{\text{CSLP}}$). *Let* $\mathbb{A}\text{gt}$ *be a set of agents,* Π *a set of propositions, and* $\mathcal{V}ars$ *a set of* strategic variables *(with typical element* σ*). The logic* $\mathcal{L}^1_{CSLP}(\mathbb{A}\text{gt},\Pi,\mathcal{V}ars)$ *is defined as* $\mathcal{L}_{CSLP}(\mathbb{A}\text{gt},\Pi,\mathcal{T}(\Omega_1))$ *where* Ω_1 *is given by*

$$\{\sigma_1.(Q_2\sigma_2)\dots(Q_n\sigma_n)\varphi \mid n\in\mathbb{N},\forall i\ (1\le i\le n\Rightarrow \sigma_i\in\mathcal{V}ars,\ Q_i\in\{\forall,\exists\},\varphi\in\mathcal{L}_{CSLP}(\mathbb{A}\text{gt},\Pi,\mathcal{T}(\{\sigma_1,\dots,\sigma_n\}))))\ \}.$$

The semantics of $\mathcal{L}^1_{\text{CSLP}}$ formulae is analogously defined as for the base language but instead of the basic plausibility mapping $[\![\cdot]\!]$, the *extended plausibility mapping* $\widehat{[\![\cdot]\!]}_M$ is used, defined as follows:

1. If $\omega\in\Omega$ then $\widehat{[\![\omega]\!]}^Q_M = [\![\omega]\!]^Q_M$;
2. If $\omega=\omega'[A]$ then $\widehat{[\![\omega]\!]}^Q_M = \{s\in\Sigma \mid \exists s'\in\widehat{[\![\omega']\!]}^Q_M\ s|_A = s'|_A\}$;
3. If $\omega = \langle\omega_1,\dots\omega_k\rangle$ then $\widehat{[\![\omega]\!]}^Q_M = \{s\in\Sigma \mid \exists t_1\in\widehat{[\![\omega_1]\!]}^Q_M,\dots,\exists t_k\in\widehat{[\![\omega_k]\!]}^Q_M\forall i=1,...,k\ s|_{a_i}=t_i|_{a_i})\}$;
4. If $\omega=\sigma_1.(Q_2\sigma_2)\dots(Q_n\sigma_n)\varphi$ then $\widehat{[\![\omega]\!]}^Q_M = \{s_1\in\Sigma\mid Q_2s_2\in\Sigma,\dots,Q_ns_n\in\Sigma\ \ (M^{s_1,\dots,s_n},q\models\varphi)\}$, where $M^{s_1,\dots,s_n}$ is equal to M except that we fix $\Upsilon_{M^{s_1,\dots,s_n}}=\Sigma$, $\Omega_{M^{s_1,\dots,s_n}}=\Omega_M\cup\{\sigma_1,\dots,\sigma_n\}$, $[\![\sigma_i]\!]^Q_{M^{s_1,\dots,s_n}}=\{s_i\}$, and $[\![\omega]\!]^Q_{M^{s_1,\dots,s_n}}=[\![\omega]\!]^Q_M$ for all $\omega\neq\sigma_i$, $1\le i\le n$, and $Q\subseteq St_M$. That is, the denotation of σ_i in $M^{s_1,\dots,s_n}$ is set to strategy profile s_i.

An example $\mathcal{L}^1_{\text{CSLP}}$ formula is
(**set-pl** $\sigma.\langle\!\langle\emptyset\rangle\!\rangle\Box(\mathbf{Ph}\,\langle\!\langle\mathbb{A}\text{gt}\rangle\!\rangle\bigcirc\mathsf{alive}\rightarrow(\mathbf{set\text{-}pl}\ \sigma)\mathbf{Pl}\,\langle\!\langle\emptyset\rangle\!\rangle\bigcirc\mathsf{alive}))$ $\neg\mathsf{Bel}_a\langle\!\langle b\rangle\!\rangle\Diamond\mathsf{suicide}$: Assuming that rational agents avoid death whenever they can, it is not rational of Ann to believe that Bob can commit suicide.

REMARK 21. *The nestings can be increased step by step which results in a hierarchy of logics,* $\mathcal{L}^k_{CSLP}$ $(k=1,2,\dots)$ *as in [3].*

4.2 Solution Concepts under Uncertainty

In this section we characterize solution concepts for imperfect information games in $\mathcal{L}^1_{\text{CSLP}}$. Before we do that, however, we need some way to *evaluate* different strategies. In game theory real values (payoffs) or preference relations are used to define the outcome of a given strategy. Here, we follow the approach from [3] which equips agents with *winning criteria* $\vec{\eta}=\langle\eta_1,\dots,\eta_k\rangle$ (one per agent) where $k=|\mathbb{A}\text{gt}|$. Each criterion η_a of agent a is a temporal formula. Intuitively, a given strategy profile is successful for an agent a iff the winning criterion is fulfilled on *all* resulting paths starting from *any* indistinguishable state given the strategy profile. This requirement is motivated by the fact that an agent does not know whether the system is in q or q' provided that q and q' are indistinguishable for him. So, he should play a strategy which is "good" in both states to ensure success.

DEFINITION 8 (FROM CEGSP TO NF GAME). *Let* M *be a* CEGSP*,* $q\in St_M$*, and* $\vec{\eta}$ *be a vector of winning criteria.*

We define $\mathcal{N}(M,\vec{\eta},q)$*, the* normal form game associated with M, $\vec{\eta}$, and q, *as the normal form game* $\langle\mathbb{A}\text{gt},\mathcal{S}_1,\dots,\mathcal{S}_k,\mu\rangle$*, where the set* $\mathcal{S}_a$ *of* a*'s strategies is given by* Σ_a *(*a*'s uniform*

[4] In order to give a brief presentation we do not allow "basic" plausibility terms anymore.

strategies) for each $a \in \mathbb{A}\mathrm{gt}$, and the payoff function is defined as follows:

$$\mu_a(a_1,\dots,a_k) = \begin{cases} 1 & \textit{if } M,\lambda \models \eta_a \\ & \textit{for all } \lambda \in out(img(q,\sim_a),\langle a_1,\dots,a_k\rangle), \\ 0 & \textit{else} \end{cases}$$

To give a clear meaning to solution concepts in a CEGSP, we relate them to the associated normal form game. The first solution concept we will define is a *best-response strategy* for IIG. Given a strategy profile $s_{-a} := (s_1,\dots,s_{a-1},s_{a+1},\dots,s_k)$ where $k = |\mathbb{A}\mathrm{gt}|$ a strategy s_a is said to be a *best response* to s_{-a} if there is no better strategy for agent a given s_{-a}. Now, s is a *best response profile* wrt a if $s|_a$ is a best response against $s|_{\mathbb{A}\mathrm{gt}\setminus\{a\}}$. According to [3] σ is a best response profile for perfect information games wrt a and $\vec{\gamma}$ in M,q if $M,q \models (\textbf{set-pl } \sigma\mathbb{A}\mathrm{gt}\setminus\{a\}])\textbf{Pl}\,(\langle\!\langle a\rangle\!\rangle\eta_a \rightarrow (\textbf{set-pl } \sigma)\langle\!\langle\emptyset\rangle\!\rangle\eta_a)$. It is read as follows: If agent a has any strategy to enforce η_a against $\sigma[\mathbb{A}\mathrm{gt}\setminus\{a\}]$ then his strategy given in σ should enforce η_a as well.

What do we have to modify to make it suitable for imperfect information games? Firstly, we have to ensure that the strategy σ is uniform, and indeed only uniform strategies are taken into account in the semantics of CSLP. Secondly, since the agent might not be aware of the real state of the system the described strategy should have its desired characteristics in every indistinguishable state. The agent should be able to *identify* the strategy; the key motivation behind CSL. For this purpose CSLP provides the constructive belief operators; recall that $\mathbb{W}_a\langle\!\langle a\rangle\!\rangle$ means that a has a single strategy successful in all indistinguishable states. To ensure this second point we just have to couple strategic operators with constructive operators. So we obtain the following description of a best response strategy for IIG:

$$BR_a^{\vec{\eta}}(\sigma) \equiv$$
$$(\textbf{set-pl } \sigma[\mathbb{A}\mathrm{gt}\setminus\{a\}])\textbf{Pl}\,(\mathbb{W}_a\langle\!\langle a\rangle\!\rangle\eta_a \rightarrow (\textbf{set-pl } \sigma)\mathbb{W}_a\langle\!\langle\emptyset\rangle\!\rangle\eta_a).$$

Other solution concepts characterized in [3] can be adapted to IIG's following the same scheme, e.g.:

Nash equilibrium (NE): $NE^{\vec{\eta}}(\sigma) \equiv \bigwedge_{i\in\mathbb{A}\mathrm{gt}} BR_i^{\vec{\eta}}(\sigma)$;
Subgame perfect NE: $SPN^{\vec{\eta}}(\sigma) \equiv \mathbb{E}\mathbb{W}_{\mathbb{A}\mathrm{gt}}\langle\!\langle\emptyset\rangle\!\rangle\Box\ NE^{\vec{\eta}}(\sigma)$;
Pareto optimal strategy (PO):

$$PO^{\vec{\eta}}(\sigma) \equiv \forall\sigma'\ \textbf{Pl}\Big(\bigwedge_{a\in\mathbb{A}\mathrm{gt}} ((\textbf{set-pl } \sigma')\mathbb{W}_a\langle\!\langle\emptyset\rangle\!\rangle\eta_a \rightarrow (\textbf{set-pl } \sigma)\mathbb{W}_a\langle\!\langle\emptyset\rangle\!\rangle\eta_a) \vee \bigvee_{a\in\mathbb{A}\mathrm{gt}} ((\textbf{set-pl } \sigma)\mathbb{W}_a\langle\!\langle\emptyset\rangle\!\rangle\eta_a \wedge \neg(\textbf{set-pl } \sigma')\mathbb{W}_a\langle\!\langle\emptyset\rangle\!\rangle\eta_a\Big).$$

The following result shows that these concepts match the underlying intuitions.

THEOREM 22. *Let M be a CEGSP, $q \in St_M$, $\vec{\eta}$ a vector of winning criteria, and $\mathcal{N} := \mathcal{N}(M,\overrightarrow{\eta},q)$. Then:*

1. *The set of* NE *strategies in $\mathcal{N}$ is given by $[\![\sigma.\widehat{NE^{\vec{\eta}}(\sigma)}]\!]_M^{\{q\}}$*
2. *The set of* PO *strategies in $\mathcal{N}$ is given by $[\![\sigma.\widehat{PO^{\vec{\eta}}(\sigma)}]\!]_M^{\{q\}}$*
3. *Let Q' collect the states that any agent from A considers possible, i.e., $\mathrm{img}(\{q\},\sim_{\mathbb{A}\mathrm{gt}}^E)$), plus all states reachable from them by (a sequence of) temporal transitions. Then, $[\![\sigma.\widehat{SPN^{\vec{\eta}}(\sigma)}]\!]_M^{\{q\}}$ is equal to $\bigcap_{q'\in Q'} [\![\sigma.\widehat{NE^{\vec{\eta}}(\sigma)}]\!]_M^{\{q'\}}$.*

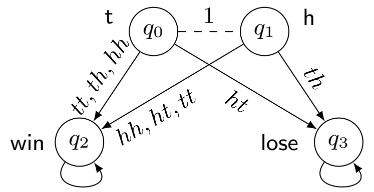

1\2	h	t
h	1, 1	0, 1
t	0, 0	1, 1

Figure 2: Simple CEGSP**.**

EXAMPLE 2. *Consider the* CEGSP *given in Figure 2. There are two agents, 1 and 2, and a coin which initially shows tail (q_0) or head (q_1); agent 1 cannot distinguish between them. Now, both agents win if 1 guesses the right side of the coin or if both agents agree on one side (regardless of whether it is the right one). For instance, the tuple th denotes that 1 says tail and 2 head. Moreover, we assume that both agents have the winning criterion $\bigcirc$win. The associated NF game wrt q_0 is also given in Figure 2. Now we have that $[\![\sigma.\widehat{NE^{\vec{\eta}}(\sigma)}]\!]_M^{\{q\}} = \{hh, tt\}$: Only if both agents agree on the same side, winning is guaranteed.*

5. MODEL CHECKING RATIONAL PLAY UNDER IMPERFECT INFORMATION

In this section we discuss the model checking complexity of $\mathcal{L}_{\mathrm{CSLP}}$ and $\mathcal{L}^1_{\mathrm{CSLP}}$. Given a formula φ, a model M, and a set of states $Q \subseteq St_M$ the associated *model checking* problem is to determine whether $M,Q \models \varphi$ holds or not. In the following we use l to refer to the length of φ and m to denote the number of transitions in M. We only consider a restricted class of models in which the check for plausibility of a strategy profile can be done in polynomial time (wrt l and m) by a non-deterministic Turing machine. In order to conduct a sensible analysis such an assumption is necessary. To this end, we adapt an important notion from [3].

DEFINITION 9 (WELL-BEHAVED CEGSP). *A* CEGSP *M is called* well-behaved *if, and only if, (1) $\Upsilon_M = \Sigma$: all the strategy profiles are plausible in M; and (2) there is an algorithm which determines whether $s \in [\![\omega]\!]_M^Q$ for every set $Q \subseteq St_M$, strategy profile $s \in \Sigma$, and plausibility term $\omega \in \Omega$ in nondeterministic polynomial time wrt the length of ω and the number of transitions in M.*

We begin by reviewing the existing results for CSL and ATLP separately. The complexity results for CSLP follow in a natural way. In [7] it was shown that CSL model checking is $\mathbf{\Delta_2^P}$-complete,[5] the hard cases being formulae $\langle\!\langle A\rangle\!\rangle\Box\varphi$ and $\langle\!\langle A\rangle\!\rangle\varphi_1\,\mathcal{U}\,\varphi_2$. The formulae require existence of a single uniform strategy which is successful in *all* states of Q. In the algorithm from [7], the strategy is guessed by the oracle and then verified in polynomial time (see further). Nested cooperation modalities are model-checked recursively (bottom-up) which puts the algorithm indeed in $\mathbf{\Delta_2^P}$.

We also recall from [3] that ATLP model checking is $\mathbf{\Delta_3^P} = \mathbf{P^{NP^{NP}}}$-complete. The algorithm for checking the hard cases ($\langle\!\langle A\rangle\!\rangle\Box\varphi$ and $\langle\!\langle A\rangle\!\rangle\varphi_1\,\mathcal{U}\,\varphi_2$) is similar: Firstly, a *plausible* strategy of A is guessed (first **NP**-oracle call) and verified

[5] $\mathbf{\Delta_2^P} = \mathbf{P^{NP}}$ is the class of problems that can be solved in polynomial time by a deterministic Turing machine that makes adaptive calls to an **NP** oracle.

against all *plausible* strategies of the opponents (second **NP**-oracle call, the "worst" response of the opponents is guessed). Note that, as soon as the relevant strategy (or strategy profile) s is fixed, the remaining verification can be done in deterministic polynomial time: it is enough to "trim" the model by deleting all transitions which cannot occur when the agents follow s, and model check a CTL formula in the trimmed model (which can be done in polynomial time [4]).

For $\mathcal{L}_{\text{CSLP}}$, we essentially use the ATLP model checking algorithm from [3] with an additional check for uniformity of strategies. This does not influence the complexity. We obtain the following result (we refer to [3, 7] for details).

THEOREM 23. *Model checking $\mathcal{L}_{CSLP}$ in the class of well-behaved* CEGSP*'s is $\mathbf{\Delta_3^P}$-complete with respect to l and m.*

PROOF (SKETCH). The hardness follows from the fact that $\mathcal{L}_{\text{ATLP}}$ is $\mathbf{\Delta_3^P}$-complete and can be embedded in $\mathcal{L}_{\text{CSLP}}$ (cf. Proposition 17). For the inclusion in $\mathbf{\Delta_3^P}$, we sketch the algorithm for $M, Q \models \langle\langle A\rangle\rangle\Box\varphi$: (1) Model-check φ recursively for each $q \in St$, and label the states for which $M, q \models \varphi$ with a new proposition p; (2) Guess a "good" plausible uniform strategy s_A; (3) Guess a "bad" uniform plausible strategy profile t such that $t|_A = s_A$; and (4) Return true if $Q \subseteq mcheck_{\text{CTL}}(M', \mathsf{A}\bigcirc\mathsf{p})$ and false otherwise, where M' is the trimmed model of M wrt profile t. □

In the previous section we showed how CSLP can be used to characterize incomplete information solution concepts. However, for this reason we had to extend the language. An obvious question arises: How much does the complexity increase? The answer is quite appealing: The increase depends on how much extra-expressiveness we actually use; and in any case, we get some expressiveness for free! This can be shown analogously to [3]; here, we just give a brief summary. The model checking complexity can be completely characterized in the number of quantifier *alternations* used in the extended plausibility terms. If we have no quantifiers at all, the resulting sublanguage is no more costly to verify than the base version. Note that the quantifier-free sublanguage of $\mathcal{L}^1_{\text{CSLP}}$ is already sufficient to "plug in" important solution concepts (e.g., Nash equilibria). For each additional quantifier alternation (starting with a universal quantifier) the complexity is pushed one level up in the polynomial hierarchy. For a more detailed discussion, cf. [3].

THEOREM 24. *Let $\mathcal{L} \subseteq \mathcal{L}^1_{CSLP}$ such that each sequence of quantifiers starting with an universal one in any extended plausibility term has at most i quantifier alternations. Then, model checking formulae of $\mathcal{L}$ in the class of well-behaved* CEGSP*'s is in $\mathbf{\Delta^P_{3+i}}$ with respect to l and m.*

PROOF (SKETCH). The extension of the base algorithm discussed above is done in an analogous way to [3]. For each quantifier alternation one has to guess a new strategy. But the first existential quantified strategic variables can be guessed together with the proponents and opponents strategies; thus, no more oracle levels need to be added. □

6. CONCLUSIONS

In this paper, we propose a logic which relates epistemic and doxastic concepts in a specific way; more importantly, it allows to reason about the outcome of rational play in imperfect information games. In the logic, called CTLKP, beliefs are defined on top of the primitive notions of plausibility and indistinguishability. We analyze the relationship between beliefs, knowledge, and rationality, and prove in particular that rational beliefs form a **KD45** modality. CSLP embeds both CTLKP and CSL; thus, the combination of knowledge, rationality, and strategic action turns out to be strictly more expressive than each of the subsets.

Moreover, we show how some important solution concepts can be characterized and used for reasoning about imperfect information scenarios. Finally, we prove that the model checking problem for the basic variant of CSLP is $\mathbf{\Delta_3^P}$-complete. That is, the complexity of model checking is only slightly higher than for CSL, and no worse than for ATLP.

Acknowledgements. Wojciech Jamroga acknowledges support of the Polish development project no. O R000024 04.

7. REFERENCES

[1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.

[2] N. Bulling and W. Jamroga. Agents, beliefs and plausible behaviour in a temporal setting. In *Proceedings of AAMAS'07*, pages 570–577, 2007.

[3] N. Bulling, W. Jamroga, and J. Dix. Reasoning about temporal properties of rational play. *Annals of Mathematics and Artificial Intelligence*, 2009.

[4] E. A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 995–1072. Elsevier Science Publishers, 1990.

[5] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* MIT Press: Cambridge, MA, 1995.

[6] N. Friedman and J. Halpern. A knowledge-based framework for belief change, Part I: Foundations. In *Proceedings of TARK*, pages 44–64, 1994.

[7] W. Jamroga and T. Ågotnes. Constructive knowledge: What agents can achieve under incomplete information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.

[8] W. Jamroga, W. van der Hoek, and M. Wooldridge. Intentions and strategies in game-like scenarios. In *Proceedings of EPIA 2005*, volume 3808 of *LNAI*, pages 512–523. Springer Verlag, 2005.

[9] M. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

[10] K. Su, A. Sattar, G. Governatori, and Q. Chen. A computationally grounded logic of knowledge, belief and certainty. In *Proceedings of AAMAS'05*, pages 149–156. ACM Press, 2005.

[11] W. van der Hoek and M. Wooldridge. Cooperation, knowledge and time: Alternating-time Temporal Epistemic Logic and its applications. *Studia Logica*, 75(1):125–157, 2003.

[12] S. van Otterloo and G. Jonker. On Epistemic Temporal Strategic Logic. *Electronic Notes in Theoretical Computer Science*, XX:35–45, 2004. Proceedings of LCMAS'04.

[13] S. van Otterloo, W. van der Hoek, and M. Wooldridge. Preferences in game logics. In *Proceedings of AAMAS-04*, pages 152–159, 2004.