

An Epistemic Strategy Logic

1

XIAOWEI HUANG, University of Liverpool
 RON VAN DER MEYDEN, UNSW Sydney

This article presents an extension of temporal epistemic logic with operators that can express quantification over agent strategies. Unlike previous work on alternating temporal epistemic logic, the semantics works with systems whose states explicitly encode the strategy being used by each of the agents. This provides a natural way to express what agents would know were they to be aware of some of the strategies being used by other agents. A number of examples that rely on the ability to express an agent's knowledge about the strategies being used by other agents are presented to motivate the framework, including reasoning about game-theoretic equilibria, knowledge-based programs, and information-theoretic computer security policies. Relationships to several variants of alternating temporal epistemic logic are discussed. The computational complexity of model checking the logic and several of its fragments are also characterized.

2
3
4
5
6
7
8
26
10
11
12

Q1 CCS Concepts:

Q2 Additional Key Words and Phrases:

13

ACM Reference format:

14

Xiaowei Huang and Ron van der Meyden. 2018. An Epistemic Strategy Logic. *ACM Trans. Comput. Logic* 19, 4, Article 26 (September 2018), 45 pages.
<https://doi.org/10.1145/3233769>

15
16
17

1 INTRODUCTION

18

19

In distributed and multi-agent systems, agents typically have a choice of actions to perform and have individual and possibly conflicting goals. This leads agents to act strategically, attempting to select their actions over time so as to guarantee achievement of their goals even in the face of other agents' adversarial behaviour. The choice of actions generally needs to be made on the basis of *imperfect* information concerning the state of the system.

20
21
22
23
24

These concerns have motivated the development of a variety of modal logics that aim to capture aspects of such settings. One of the earliest, dating from the 1980s, was multi-agent *epistemic logic* [23, 44], which introduced modal operators that deal with imperfect information by providing a way to state what agents *know*. Combining such constructs with temporal logic constructs [46] gives *temporal epistemic logics*, which support reasoning about how agents' knowledge

25
26
27
28
29

This article combines results from Reference [31], An epistemic strategy logic, X. Huang and R. van der Meyden, in *Proceedings of the 2nd International Workshop on Strategic Reasoning*, and Reference [33] A temporal logic of strategic knowledge, X. Huang and R. van der Meyden, In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*. It extends these works by including full proofs for all results.

Most of Huang's work was performed when he was at UNSW Sydney, Australia.

Q3

Authors' addresses:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

1529-3785/2018/09-ART26 \$15.00

<https://doi.org/10.1145/3233769>

30 changes over time. Temporal-epistemic logic is an area about which a significant amount is now
 31 understood [19].

32 Logics dealing with reasoning about strategies, which started to be developed in the same period
 33 [43], had a slower initial start but have in recent years become the focus of intense study [2,
 34 30, 45]. *Alternating temporal logic* (ATL) [2], which generalizes branching-time temporal logic to
 35 encompass reasoning about the temporal effects of strategic choices by one group of agents against
 36 all possible responses by their adversaries, has become a popular basis for work in this area.

37 One of the ways in which recent work has extended ATL is to add epistemic operators, yielding
 38 an *alternating temporal epistemic logic*, e.g., ATEL [29]. Many subtle issues arise concerning what
 39 agents know in settings where multiple agents act strategically. In the process of understanding
 40 these issues, there has been a proliferation of epistemic extensions of ATL [34, 35, 39, 55]. Some
 41 of the modal operators introduced in this literature are complex, interweaving ideas about the
 42 knowledge of a group of agents, the strategies available to them, the effect of playing these strate-
 43 gies against strategies available to agents not in the group, and the knowledge that other groups
 44 of agents may have about these effects.

45 Our contribution in this article is to develop a logic that extends the expressive power of previ-
 46 ous work on logics for knowledge and strategies, while at the same time simplifying the syntactic
 47 basis by identifying a small set of primitives that can be composed to represent the more complex
 48 constructs for reasoning about strategies and knowledge from prior literature. We present exam-
 49 ples to show that the logic is useful for a range of applications, including expressing notions of
 50 information flow security (such as strategic notions of noninterference and erasure policies), rea-
 51 soning about implementations of knowledge-based programs, and reasoning about game-theoretic
 52 equilibria. We also conduct a detailed analysis of the complexity of model checking a number of
 53 fragments of the logic. Our semantic framework is able to model a range of semantics for knowl-
 54 edge and strategies including a “perfect recall” interpretation, but since we are interested in model-
 55 checking complexity results at the lower end of the complexity spectrum, we concentrate on an
 56 “imperfect recall” or “observational” semantics of knowledge. (We note that model checking just
 57 ATL, even without knowledge operators, under an imperfect information and perfect recall se-
 58 mantics is already undecidable [16].)

59 At the semantic level, the key way in which our logic extends prior work on alternating
 60 temporal epistemic logic is by treating agents’ strategies as first-class citizens in the semantics,
 61 represented as components of the global state of the system at any moment of time in a run
 62 of the system. This is in contrast to most prior work in the area, where strategies are used to
 63 generate runs of a system, but the runs themselves contain no explicit information about the
 64 specific strategies used by the players to produce them. Our approach provides a referent for the
 65 notion “the strategy being used by player i ,” which cannot be expressed in most prior works on
 66 alternating temporal epistemic logic.

67 We reflect this additional referent at the syntactic level by introducing a syntactic notation $\sigma(i)$,
 68 which refers to the strategy of agent i . Since the strategy of agent i is modelled semantically as a
 69 component of the global state, just like the local state of agent i , we allow this construct to be used
 70 in the same contexts where the agent name i can be used—in particular, in operators for knowledge
 71 (including distributed and common knowledge). An example of what can be expressed with this
 72 extension is $D_{\{i, \sigma(i)\}}\phi$, which says that the truth of ϕ in all possible futures can be deduced from
 73 knowledge of agent i ’s local state plus the strategy being applied by agent i . Intuitively, the con-
 74 struct $D_{\{i, \sigma(i)\}}$ captures what agent i knows when it takes into account the strategy it is running.

75 We show that this extension of temporal epistemic logic gives a logical approach with broad
 76 applicability. In particular, as we show in Section 3.2, temporal epistemic logic extended with the
 77 indices $\sigma(i)$ can express alternating temporal logic constructs (both revocable and irrevocable).

The extension can also express many of the subtly different notions that have been proposed in the literature on alternating temporal epistemic logics. We demonstrate this (in Section 3.3) by results that show how a number of such logics can be translated into our setting. We also present a number of other applications including game-theoretic solution concepts (Section 3.5), issues of concern in computer security (Section 3.6), and reasoning about possible implementations of knowledge-based programs (Section 3.7).

In some applications, however, some richer expressiveness is required. One such application, discussed in Section 3.3, concerns expressing an operator and combining common knowledge and strategic concerns from an extended alternating temporal epistemic logic of Jamroga and van der Hoek [38]. We address this by adding to the logic constructs that can be used to express quantification over strategies. This leads to a logic that, like *strategy logic* [13, 42], supports explicit naming and quantification over strategies. Technically, we achieve this in a slightly more general way: We first generalize temporal epistemic logic to include operators $\exists x$ for quantification over global states x , as well as statements $e_i(x)$ that say that component i in the current global state is the same as component i in the global state denoted by x . Even before the introduction of strategic concerns, this gives a novel extension of temporal epistemic logic in the flavour of *hybrid logic* [4]. (As we show in Section 2, this extension enables the expression of security notions such as *nondeducibility* [53] that cannot be naturally expressed in standard temporal epistemic logics.) We then apply this generalization to a system that includes strategies encoded in the global states and references these using the “strategic” indices $\sigma(i)$. The resulting logic can express that agent i knows what strategy agent j is using, by means of the formula

$$\exists x(e_{\sigma(j)}(x) \wedge K_i e_{\sigma(j)}(x)),$$

in which the first occurrence of $e_{\sigma(j)}(x)$ binds x to a global state in which the strategy of agent j is the same as at the current state, and the remainder of the formula states that every global state considered possible by agent i has the same strategy for agent j . (This cannot be expressed in most alternating temporal epistemic logics, e.g., ATEL [29], since their semantics fails to encode the strategy being run by an agent in the locus of evaluation of formulas.) The framework is able to express the above-mentioned operator from Reference [38], as well as notions of information flow security that quantify over agent strategies, such as *nondeducibility on strategies* [57], which we discuss in Section 3.1.

The main theoretical contribution of the article is a set of results on the complexity of model checking the resulting logic and its fragments. We consider several dimensions: Does the logic have quantifiers, and what is the temporal basis for the logic, branching-time (CTL) or linear time (LTL)? The richest logics in our spectrum turn out to have EXSPACE-complete model-checking problems. However, we identify a number of special cases where model checking is in PSPACE, i.e., no more than the complexity of model checking the temporal logic LTL. One is the fragment where we allow the constructs $\exists x$ and $e_i(x)$ but restrict the temporal operators to be those of the branching-time logic CTL. Another is the fragment in which we do not allow $\exists x$ and $e_i(x)$ but allow strategic indices $\sigma(i)$ in knowledge operators and take the temporal operators from the richer branching-time logic CTL*, which extends the linear time logic LTL.

The structure of the article is as follows. In Section 2, we first develop an extension of temporal epistemic logic that adds the ability to quantify over global states and refer to global state components. We then present a semantic model for the environments in which agents choose their actions. Building on this model, we show how to construct a model for temporal epistemic logic called *strategy space* in which runs build in information about the strategy being used by each of the agents. We then define a spectrum of logics defined over the resulting semantics. These logics are obtained as fragments of the extended temporal epistemic logic, interpreted in strategy

124 space. Section 3 deals with applications of the resulting logics. In particular, we show that the
 125 logics can express reasoning about implementations of knowledge-based programs, many notions
 126 that have been proposed in the area of alternating temporal epistemic logic, game-theoretic so-
 127 lution concepts, and problems from computer security. Next, in Section 4, we provide results on
 128 the complexity of the model-checking problem for the various fragments of the logic, identify-
 129 ing fragments with lower complexity than the general problem. In Section 5, we conclude with a
 130 discussion of related literature.

131 2 AN EXTENDED TEMPORAL EPISTEMIC LOGIC

132 The usual *interpreted systems* semantics for temporal epistemic logic [19] deals with runs, in which
 133 each moment of time is associated with a global state that is composed of a local state for each agent
 134 in the system. We begin by defining the syntax and semantics of an extension of temporal epistemic
 135 logic that adds the ability to quantify over global states and refer to global state components.
 136 This syntax and semantics will be instantiated in what follows by taking some of the global state
 137 components to be the strategies being used by agents.

138 To quantify over global states, we extend temporal epistemic logic with a set of variables Var ,
 139 a quantifier $\exists x$, and a construct $e_i(x)$, where x is a variable. The formula $\exists x.\phi$ says, intuitively,
 140 that there exists in the system a global state x such that ϕ (a formula that may contain uses of the
 141 variable x) holds at the current point. The formula $e_i(x)$ asserts the equality of the local states of
 142 agent i at the current point and in the global state x .

143 Let $Prop$ be a set of atomic propositions and let Ags be a finite set of agent names, excluding the
 144 special name e , which we use to designate the environment in which the agents operate. We write
 145 Ags^+ for the set $\{e\} \cup Ags$. The language $ETLK(Ags, Prop, Var)$ (or just ETLK when the parameters
 146 are obvious) has syntax given by the grammar:

$$\phi \equiv p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid A\phi \mid \circ\phi \mid \phi_1 U \phi_2 \mid \exists x.\phi \mid e_i(x) \mid D_G\phi \mid C_G\phi,$$

147 where $p \in Prop$, $x \in Var$, $i \in Ags^+$, and $G \subseteq Ags^+$. The construct $D_G\phi$ expresses that agents in G
 148 have distributed knowledge of ϕ , i.e., could deduce ϕ if they pooled their information, and $C_G\phi$
 149 says that ϕ is common knowledge to group G . The temporal formulas $\circ\phi$, $\phi_1 U \phi_2$, $A\phi$ have the
 150 same intuitive meanings as in the temporal logic CTL^* [17], i.e., $\circ\phi$ says that ϕ holds at the next
 151 moment of time, $\phi_1 U \phi_2$ says that ϕ_1 holds until ϕ_2 does, and $A\phi$ says that ϕ holds in all possible
 152 evolutions from the present situation.

153 Other operators can be defined in the usual way, e.g., $\phi_1 \wedge \phi_2 = \neg(\neg\phi_1 \vee \neg\phi_2)$, $\diamond\phi = (true U \phi)$,
 154 which says that ϕ holds eventually, $\square\phi = \neg\diamond\neg\phi$, which says that ϕ always holds, $E\phi = \neg A\neg\phi$,
 155 which says that ϕ holds on some path from the current point, and so on. The universal form
 156 $\forall x.\phi = \neg\exists x.\neg\phi$ expresses that ϕ holds for all global states x that occur in the system. For an agent
 157 $i \in Ags^+$, we write $K_i\phi$ for $D_{\{i\}}\phi$ —this expresses that agent i knows the fact ϕ . The notion of
 158 everyone in group G knowing ϕ can then be expressed as $E_G\phi = \bigwedge_{i \in G} K_i\phi$. We write $e_G(x)$ for
 159 $\bigwedge_{i \in G} e_i(x)$. This says that at the current point, the agents in G have the same local state as they
 160 do at the global state named by variable x .

161 We will be interested in a fragment of the logic that restricts the occurrence of the temporal
 162 operators to some simple patterns, in the style of the branching-time temporal logic CTL [15]. We
 163 write $ECTL(Ags, Prop, Var)$ (or just ECTL when the parameters are obvious) for the fragment of
 164 the language $ETLK(Ags, Prop, Var)$ in which the temporal operators occur only in the particular
 165 forms $A\circ\phi$, $E\circ\phi$, $A\phi_1 U \phi_2$, and $E\phi_1 U \phi_2$. In the context of temporal logic, these restrictions reduce
 166 the complexity of model checking from PSPACE to PTIME [15]. It is therefore interesting to study
 167 the impact on complexity of a similar restriction in the context of our additional operators.

The semantics of $\text{ETLK}(\text{Ags}, \text{Prop}, \text{Var})$ builds straightforwardly on the following definitions used in the standard semantics for temporal epistemic logic [19]. Consider a system for a set of agents Ags . A *global state* is an element of the set $\mathcal{G} = \prod_{i \in \text{Ags}^+} L_i$, where L_e is a set of states of the environment and L_i is a set of *local states* for each agent $i \in \text{Ags}$. A *run* is a mapping $r : \mathbb{N} \rightarrow \mathcal{G}$ giving a global state at each moment of time. For $n \leq m$, write $r[n \dots m]$ for the sequence $r(n)r(n+1) \dots r(m)$. We also write $r[n \dots]$ for the infinite sequence $r(n)r(n+1) \dots$. A *point* is a pair (r, m) consisting of a run r and a time $m \in \mathbb{N}$. An *interpreted system* is a pair $\mathcal{I} = (\mathcal{R}, \pi)$, where \mathcal{R} is a set of runs and π is an *interpretation*, mapping each point (r, m) with $r \in \mathcal{R}$ to a subset of Prop . Elements of $\mathcal{R} \times \mathbb{N}$ are called the *points* of \mathcal{I} . For each $i \in \text{Ags}^+$, we write $r_i(m)$ for the corresponding component of $r(m)$ in L_i and then define an equivalence relation on points by $(r, m) \sim_i (r', m')$ if $r_i(m) = r'_i(m')$. We also define $\sim_G^D \equiv \bigcap_{i \in G} \sim_i$, and $\sim_G^E \equiv \bigcup_{i \in G} \sim_i$, and $\sim_G^C \equiv (\bigcup_{i \in G} \sim_i)^*$ for $G \subseteq \text{Ags}$, where $*$ denotes the reflexive transitive closure of a relation. We take \sim_\emptyset^D to be the universal relation on points and (for the sake of preserving monotonicity of these relations in these degenerate cases) take \sim_\emptyset^E and \sim_\emptyset^C to be the identity relation.

To extend this semantic basis for temporal epistemic logic to a semantics for ETLK, we just need to add a construct that interprets variables as global states. A *context* for an interpreted system \mathcal{I} is a mapping Γ from Var to global states occurring in \mathcal{I} , i.e., such that for all $x \in \text{Var}$ there exists a point (r, m) of \mathcal{I} such that $\Gamma(x) = r(m)$. When g is a global state and $x \in \text{Var}$, we write $\Gamma[g/x]$ for the context Γ' with $\Gamma'(x) = g$ and $\Gamma'(y) = \Gamma(y)$ for all variables $y \neq x$. The semantics of the language ETLK is given by a relation $\Gamma, \mathcal{I}, (r, m) \models \phi$, representing that formula ϕ holds at point (r, m) of the interpreted system \mathcal{I} , relative to context Γ . This is defined inductively on the structure of the formula ϕ , as follows:

- $\Gamma, \mathcal{I}, (r, m) \models p$ if $p \in \pi(r, m)$;
- $\Gamma, \mathcal{I}, (r, m) \models \neg\phi$ if not $\Gamma, \mathcal{I}, (r, m) \models \phi$;
- $\Gamma, \mathcal{I}, (r, m) \models \phi \wedge \psi$ if $\Gamma, \mathcal{I}, (r, m) \models \phi$ and $\Gamma, \mathcal{I}, (r, m) \models \psi$;
- $\Gamma, \mathcal{I}, (r, m) \models A\phi$ if $\Gamma, \mathcal{I}, (r', m) \models \phi$ for all $r' \in \mathcal{R}$ with $r[0 \dots m] = r'[0 \dots m]$;
- $\Gamma, \mathcal{I}, (r, m) \models \bigcirc\phi$ if $\Gamma, \mathcal{I}, (r, m+1) \models \phi$;
- $\Gamma, \mathcal{I}, (r, m) \models \phi U \psi$ if there exists $m' \geq m$ such that $\Gamma, \mathcal{I}, (r, m') \models \psi$ and $\Gamma, \mathcal{I}, (r, k) \models \phi$ for all k with $m \leq k < m'$;
- $\Gamma, \mathcal{I}, (r, m) \models \exists x.\phi$ if $\Gamma[r'(m')/x], \mathcal{I}, (r, m) \models \phi$ for some point (r', m') of \mathcal{I} ;
- $\Gamma, \mathcal{I}, (r, m) \models e_i(x)$ if $r_i(m) = \Gamma(x)_i$;
- $\Gamma, \mathcal{I}, (r, m) \models D_G\phi$ if $\Gamma, \mathcal{I}, (r', m') \models \phi$ for all (r', m') such that $(r', m') \sim_G^D (r, m)$;
- $\Gamma, \mathcal{I}, (r, m) \models C_G\phi$ if $\Gamma, \mathcal{I}, (r', m') \models \phi$ for all (r', m') such that $(r', m') \sim_G^C (r, m)$.

The definition is standard, except for the constructs $\exists x.\phi$ and $e_i(x)$. The clause for the former says that $\exists x.\phi$ holds at a point (r, m) if there exists a global state $g = r'(m')$ such that ϕ holds at the point (r, m) , provided we interpret x as referring to g . Note that it is required that g is attained at some point (r', m') , so actually occurs in the system \mathcal{I} . The clause for $e_i(x)$ says that this holds at a point (r, m) if the local state of agent i , i.e., $r_i(m)$, is the same as the local state $\Gamma(x)_i$ of agent i at the global state $\Gamma(x)$ that interprets the variable x according to Γ .

We remark that these novel constructs introduce some redundancy, in that the set of epistemic operators D_G could be reduced to the “universal” operator D_\emptyset , since $D_G\phi \equiv \exists x.(e_G(x) \wedge D_\emptyset(e_G(x) \Rightarrow \phi))$. Evidently, given the syntactic complexity of this formulation, D_G remains a useful notation.

Example 2.1. As an example of a property that can be naturally expressed in ESL, but not in most standard temporal epistemic logics (e.g., ESL minus the operators $\exists x$ and $e_i(x)$), consider information flow security properties in the spirit of *nondeducibility* [53]. Suppose that there are

214 two agents Hi and Lo , representing two information security levels High and Low, respectively.
 215 The High level contains secrets that need to be protected from an attacker, represented by the Low
 216 level. Nondeducibility security properties, intuitively, assert that Lo always has no information
 217 about Hi . When the information that needs to be protected is represented in the local state of Hi ,
 218 this means that Lo should always consider all local states of Hi possible. This can be expressed
 219 using the formula

$$\Box(\neg\exists x(K_{Lo}(\neg e_{Hi}(x))).$$

220 Here, $K_{Lo}(\neg e_{Hi}(x))$ expresses that Lo has some information about Hi , because there exists some
 221 local state of Hi that Lo is able to exclude, namely, the local state g_{Hi} where g is the global state
 222 denoted by x . By asserting that it is always the case that there does not exist such a state x whose
 223 Hi -local component Lo is able to exclude, we say that Lo never has information about Hi . Equiva-
 224 lently, pushing the outer negation inwards gives the form $\Box(\forall x(\neg K_{Lo}(\neg e_{Hi}(x))))$ which says that
 225 Lo always considers all local states of Hi to be possible.

226 We remark that the operators $\exists x$ and $e_i(x)$ may be eliminated from the above formula if the
 227 system \mathcal{I} is known and has a sufficiently rich set of atomic propositions that each local state h of
 228 Hi is associated with a conjunction ϕ_h of literals that is true exactly at global states g with $g_{Hi} = h$.
 229 Let L_{Hi} be the set of local states of Hi . This gives the equivalence

$$\exists x(K_{Lo}(\neg e_{Hi}(x))) \equiv \bigvee_{h \in L_{Hi}} K_{Hi}(\neg \phi_h),$$

230 which is valid in \mathcal{I} . However, if the system \mathcal{I} over all systems, then no single formula of the logic
 231 without the operators $\exists x$ and $e_i(x)$ can be equivalent to $\exists x(K_{Lo}(\neg e_{Hi}(x)))$, because a fixed set of
 232 propositions cannot distinguish an arbitrarily large set of states.

233 2.1 Strategic Environments

234 To semantically represent settings in which agents operate by strategically choosing their actions,
 235 we introduce *environments*, a type of transition system that models the available actions and their
 236 effects on the state. This modelling is long established in the literature on reasoning about knowl-
 237 edge [41] and is similar to models used in the tradition of alternating temporal logic [2]. From an
 238 environment and a class of strategies, we construct an instance of the interpreted systems seman-
 239 tics defined in the previous section. One of the innovations in this construction is to introduce new
 240 names that refer to global state components that represent the strategies being used by the agents.

241 An *environment* for agents Ags is a tuple $E = \langle S, I, \{Acts_i\}_{i \in Ags}, \rightarrow, \{O_i\}_{i \in Ags}, \pi \rangle$, where

- 242 (1) S is a set of states,
- 243 (2) I is a subset of S , representing the initial states,
- 244 (3) for each $i \in Ags$, component $Acts_i$ is a nonempty set of actions that may be performed by
 245 agent i ; we define $Acts = \prod_{i \in Ags} Acts_i$ to be the set of joint actions,
- 246 (4) $\rightarrow \subseteq S \times Acts \times S$ is a transition relation, labelled by joint actions,
- 247 (5) for each $i \in Ags$, component $O_i : S \rightarrow L_i$ is an observation function, and
- 248 (6) $\pi : S \rightarrow \mathcal{P}(Prop)$ is a propositional assignment.

249 Here, the range L_i of the observation function O_i is any set, and what will matter in the semantics
 250 is an equivalence relation derived from this function.

251 An environment is said to be finite if all its components, i.e., S , Ags , $Acts_i$, L_i , and $Prop$, are finite.
 252 Intuitively, a joint action $a \in Acts$ represents a choice of action $a_i \in Acts_i$ for each agent $i \in Ags$,
 253 performed simultaneously, and the transition relation resolves this into an effect on the state. We
 254 assume that \rightarrow is serial in the sense that for all $s \in S$ and $a \in Acts$ there exists $t \in S$ such that
 255 $(s, a, t) \in \rightarrow$. We also write $s \xrightarrow{a} t$ for $(s, a, t) \in \rightarrow$.

Example 2.2. We describe an environment for a secure message transmission problem, which models a sender agent HS at a High security level that has a bit of information to be transmitted to a receiver agent HR , also at a High security level, via a channel represented by an agent Lo at the Low security level (e.g., the internet). The transmission is handled by an agent Cr that models cryptography that may be applied to the message before transmission. Thus, we take $Ags = \{HS, Cr, HR, Lo\}$. The environment has the following components:

- The set of states S is the set of assignments to the following variables:
 - s , representing the sender’s secret bit, with value in $\{0, 1\}$
 - k , representing a secret encryption key, with value in $\{0, 1\}$
 - c , representing the unsecured communication channel, with value in $\{0, 1, \perp\}$
 We represent a state in S in the format $\langle s, k, c \rangle$, corresponding to the values of the three variables.
- The set I of initial states is the set $\langle s, k, \perp \rangle$ where $s, k \in \{0, 1\}$. That is, the value of the channel c is initially \perp , representing that no message has yet been sent.
- We associate the following sets of actions with the agents: $Acts_{HS} = Acts_{HR} = Acts_{Lo} = \{\text{skip}\}$ and $Acts_{Cr} = \{c := s \oplus k, c := s \oplus \bar{k}\}$. Thus, agents HS, HR, Lo are inert; they can perform only the action skip , which has no effect on their local states. The only active agent is Cr , which has two actions, each of which encrypts the message bit s using the key k and places the result in the channel c . Encryption is done by computing the exclusive-or \oplus of the message with information from the key. The two actions correspond to taking the information from the key to be either the key bit k itself or its complement \bar{k} . Since agents HS, HR, Lo always perform skip , we may, for brevity, name joint actions using the action names for agent Cr , i.e., if a is one of Cr ’s actions, then we denote a joint action $\langle \text{skip}, a, \text{skip}, \text{skip} \rangle$ in $Acts = Acts_{HS} \times Acts_{Cr} \times Acts_{HR} \times Acts_{Lo}$ as just a .
- The transition relation resolves joint actions denoted as $a \in Acts_{Cr}$ as follows:

$$\langle s, k, c \rangle \xrightarrow{a} \langle s', k', c' \rangle$$

if either a is $c := s \oplus k$ and $s' = s, k' = k$ and $c' = s \oplus k$, or a is $c := s \oplus \bar{k}$ and $s' = s, k' = k$ and $c' = s \oplus \bar{k}$.

- We define the observation functions for each of the agents on states $\langle s, k, c \rangle \in S$ as follows:
 - Agent HS observes just the bit to be transmitted, i.e., $O_{HS}(\langle s, k, c \rangle) = s$.
 - Agent Cr observes both the bit to be transmitted and the value of the encryption key, i.e., $O_{Cr}(\langle s, k, c \rangle) = \langle s, k \rangle$.
 - Agent HR observes the communication channel and the value of the encryption key, i.e., $O_{HR}(\langle s, k, c \rangle) = \langle k, c \rangle$.
 - Agent Lo observes just the communication channel, i.e., $O_{Lo}(\langle s, k, c \rangle) = c$.
- We do not need any propositions in our later uses of this environment, so we take $Prop = \emptyset$ and $\pi : S \rightarrow Prop$ to be the trivial assignment.

A *strategy* for agent $i \in Ags$ in an environment E is a function $\alpha_i : S \rightarrow \mathcal{P}(Acts_i) \setminus \{\emptyset\}$, selecting a nonempty set of actions of the agent at each state.¹ We call these actions *enabled* at the state for agent i . A *group strategy*, or *strategy profile*, for a group G is a tuple $\alpha_G = \langle \alpha_i \rangle_{i \in G}$ where each α_i is a strategy for agent i . A *joint strategy* is a group strategy for the group Ags of all agents. If $\alpha = \langle \alpha_i \rangle_{i \in G}$ is a group strategy for group G , and $H \subseteq G$, then we write $\alpha \upharpoonright H$ for the restriction $\langle \alpha_i \rangle_{i \in H}$ of α to H .

¹More generally, a strategy could be a function of the history, but we focus here on strategies that depend only on the final state.

298 A strategy α_i for agent i is *deterministic* if $\alpha_i(s)$ is a singleton for all s . A strategy α_i for agent i
 299 is *uniform* if for all states s, t , if $O_i(s) = O_i(t)$, then $\alpha_i(s) = \alpha_i(t)$. Intuitively, uniformity captures
 300 the constraint that agents' actions are chosen using no more information than they obtain from
 301 their observations.² A strategy $\alpha_G = \langle \alpha_i \rangle_{i \in G}$ for a group G is *locally uniform (deterministic)* if α_i
 302 is uniform (respectively, deterministic) for each agent $i \in G$.³ Given an environment E , we write
 303 $\Sigma^{det}(E)$ for the set of deterministic joint strategies, $\Sigma^{unif}(E)$ for the set of all locally uniform joint
 304 strategies, and $\Sigma^{unif, det}(E)$ for the set of all deterministic locally uniform joint strategies.

305 *Example 2.3.* We present some joint strategies in the environment of Example 2.2. For agents
 306 $i \in \{HS, HR, Lo\}$, the only available action is skip, so all joint strategies α have $\alpha_i(s) = \{\text{skip}\}$
 307 for all $s \in S$. Thus, each joint strategy α is determined by its component α_{Cr} , the strategy of the
 308 encryption agent.

309 The encryption agent could always choose the action $c := s \oplus k$, giving the strategy α_{Cr}^0 defined
 310 by $\alpha_{Cr}^0(s) = \{c := s \oplus k\}$ for all states s . This strategy is both locally uniform and deterministic.

311 If the encryption agent chooses its action non-deterministically, then we have the strategy α_{Cr}^1
 312 defined by $\alpha_{Cr}^1(s) = \text{Acts}_{Cr}$ for all states s . This strategy is locally uniform but not deterministic.

313 An alternate strategy for the encryption agent is to choose its action based on the values it
 314 observes. Consider the strategy α_{Cr}^2 defined by letting $\alpha_{Cr}^2(\langle s, k, c \rangle)$ be the singleton set $\{c := s \oplus k\}$
 315 if $k = 0$ and the action $\{c := s \oplus \bar{k}\}$ otherwise. This strategy is deterministic. Also, since the value
 316 k is always part of the agent's observation, this strategy is locally uniform.

317 2.2 Strategy Space

318 We now define an interpreted system, called the *strategy space* of an environment, that contains
 319 all the possible runs generated when agents Ags behave by choosing a strategy from some set Σ
 320 of joint strategies in the context of an environment E . To enable reference to the strategy being
 321 used by agent $i \in Ags$, we introduce the notation " $\sigma(i)$ " as a name referring to agent i 's strategy.
 322 For $G \subseteq Ags$, we write $\sigma(G)$ for the set $\{\sigma(i) \mid i \in G\}$.

323 Technically, $\sigma(i)$ will be treated as if it were an agent in the context of temporal epistemic logic,
 324 in the sense that it will be the index of a local state component of the global state. In particular, we
 325 take the value of the local state at index $\sigma(i)$ to be the strategy in use by agent i . We will permit use
 326 of the indices $\sigma(i)$ in epistemic operators. This provides a way to refer, using distributed knowledge
 327 operators D_G , where G contains the strategic indices $\sigma(i)$, to what agents would know should
 328 they take into account not just their own observations but also information about other agents'
 329 strategies. For example, the distributed knowledge operator $D_{\{i, \sigma(i)\}}$ captures the knowledge that
 330 agent i has, taking into account the strategy that it is running. Operator $D_{\{i, \sigma(i), \sigma(j)\}}$ captures what
 331 agent i would know, taking into account its own strategy and the strategy being used by agent j .
 332 Various applications of the usefulness of this expressiveness are given in Section 3.

333 We note, however, that unlike the base agent $i \in Ags$, the index $\sigma(i)$ is not one of the agents
 334 in the environment E , and it is not associated with any actions. The index $\sigma(i)$ exists only in the
 335 interpreted system that we generate from E . (A similar remark applies to the special agent e , which
 336 is also not associated with any actions.) Since the indices $\sigma(i)$ are not agents in the same sense
 337 as agents $i \in Ags$, the reader may prefer to read $D_G \phi$ with $\sigma(i) \in G$ as " ϕ is *deducible* from the

²Recall that we work in this article with agents with imperfect recall. For agents with perfect recall, we would use a notion of uniformity that allows agents choice of action to depend on all their past observations.

³We prefer the term "locally uniform" to just "uniform" in the case of groups, since we could say a strategy α for group G is *globally uniform* if for all states s, t , if $O_i(s) = O_i(t)$ for all $i \in G$, then $\alpha_i(s) = \alpha_i(t)$ for all $i \in G$. While we do not pursue this in the present article, this notion would be interesting in settings where agents share information to collude on their choice of move.

information contained in state components G ” rather than the more standard “it is distributed knowledge to agents G that ϕ .”

Formally, suppose we are given an environment $E = \langle S, I, \{Acts_i\}_{i \in Ags}, \rightarrow, \{O_i\}_{i \in Ags}, \pi \rangle$ for agents Ags , where $O_i : S \rightarrow L_i$ for each $i \in Ags$, and a set $\Sigma \subseteq \prod_{i \in Ags} \Sigma_i$ of joint strategies for the group Ags . We define the *strategy space* interpreted system $\mathcal{I}(E, \Sigma) = (\mathcal{R}, \pi')$ as follows.⁴ The system $\mathcal{I}(E, \Sigma)$ has global states $\mathcal{G} = S \times \prod_{i \in Ags} L_i \times \prod_{i \in Ags} \Sigma_i$. Intuitively, each global state consists of a state of the environment E , a local state for each agent i in E , and a strategy for each agent i . We index the components of this Cartesian product by e , the elements of Ags and the elements of $\sigma(Ags)$, respectively. We take the set of runs \mathcal{R} of $\mathcal{I}(E, \Sigma)$ to be the set of all runs $r : \mathbb{N} \rightarrow \mathcal{G}$ satisfying the following constraints, for all $m \in \mathbb{N}$ and $i \in Ags$:

- (1) $r_e(0) \in I$ and $\langle r_{\sigma(i)}(0) \rangle_{i \in Ags} \in \Sigma$,
- (2) $r_i(m) = O_i(r_e(m))$,
- (3) $(r_e(m), a, r_e(m+1)) \in \rightarrow$ for some joint action $a \in Acts$ such that for all $j \in Ags$ we have $a_j \in \alpha_j(r_e(m))$, where $\alpha_j = r_{\sigma(j)}(m)$, and
- (4) $r_{\sigma(i)}(m+1) = r_{\sigma(i)}(m)$.

The interpretation π' of $\mathcal{I}(E, \Sigma)$ is determined from the interpretation π of E by taking $\pi'(r, m) = \pi(r_e(m))$ for all points (r, m) .

The first constraint on runs says, intuitively, that runs start at an initial state of E , and the initial strategy profile at time 0 is one of the profiles in Σ . The second constraint states that the agent i 's local state at time m is the observation that agent i makes of the state of the environment at time m . The third constraint says that evolution of the state of the environment is determined at each moment of time by agents choosing an action by applying their strategy at that time to the state at that time. The joint action resulting from these individual choices is then resolved into a transition on the state of the environment using the transition relation from E . The final constraint says that agents' strategies are fixed during the course of a run. Intuitively, each agent picks a strategy and then sticks to it.

Our epistemic strategy logic is now just an instantiation of the extended temporal epistemic logic in the strategy space generated by an environment. That is, we start with an environment E and an associated set of strategies Σ and then work with the language $ETLK(Ags \cup \sigma(Ags), Prop, Var)$ in the interpreted system $\mathcal{I}(E, \Sigma)$. (Recall that this notation implicitly includes a local state component e to represent the state of the environment.) We call this instance of the language $ESL(Ags, Prop, Var)$ or just ESL when the parameters are implicit.

Since interpreted systems are always infinite objects, we use environments to give a finite input for the model-checking problem. For an environment E , a set of strategies Σ for E , and a context Γ for $\mathcal{I}(E, \Sigma)$, we write $\Gamma, E, \Sigma \models \phi$ if $\Gamma, \mathcal{I}(E, \Sigma), (r, 0) \models \phi$ for all runs r of $\mathcal{I}(E, \Sigma)$. Often, the formula ϕ will be a sentence, i.e., will have all variables x in the scope of an operator $\exists x$. In this case the statement $\Gamma, E, \Sigma \models \phi$ is independent of Γ and we write simply $E, \Sigma \models \phi$.

We will be interested in a number of fragments of ESL that turn out to have lower complexity. We define $ESL^-(Ags, Prop, Var)$, or just ESL^- , to be the language

$$ECTL(Ags \cup \sigma(Ags), Prop, Var).$$

⁴The construction given here is for an “observational” or “imperfect recall” modelling of knowledge that assumes that an agent reasons, and chooses its next action, on the basis of its current observation only. It is straightforward to give other constructions such as a synchronous perfect recall semantics, where we work with the sequence of observations and actions of the agent instead. Model checking for such a variant would be undecidable, so we do not pursue this here.

377 Another fragment of the language that will be of interest is the language, denoted

$$\text{CTL}^*K(\text{Ags} \cup \sigma(\text{Ags}), \text{Prop}, \text{Var}),$$

378 in which we omit the constructs \exists and $e_i(x)$; this is a standard branching-time temporal epistemic
379 language except that it contains the strategy indices $\sigma(\text{Ags})$.

380 3 APPLICATIONS

381 We now consider a range of applications of the logic ESL and show how it can represent notions
382 from earlier work on alternating temporal epistemic logic. (In a few cases, we prove precise trans-
383 lation results, but due to the large number of operators and distinct semantics underlying these
384 logics in the literature, we just sketch intuitive correspondences in most cases.)

385 3.1 Variants of Nondeducibility

386 We already mentioned the notion of nondeducibility in Example 2.1, which shows one way that
387 our logic extends the expressiveness of previous work on temporal epistemic logic by allowing
388 quantification over agents' local states to be expressed. We continue discussion of this example
389 here in the context of the environment E of Example 2.2. We also show that our logic can rep-
390 resent a related notion from the security literature called *nondeducibility on strategies* [57] that
391 involves an agent reasoning not just based on its local state but also using knowledge of the strat-
392 egy being employed by another agent. This demonstrates a further dimension in which we can
393 express more than prior work on alternating temporal epistemic logic and shows the value of al-
394 lowing the strategic indices $\sigma(i)$ to occur in epistemic operators. (Our discussion in this section
395 loosely follows examples used in Reference [57] to motivate nondeducibility on strategies.)

396 Consider first the instance

$$\text{NonDed} = \Box(\neg\exists x.(K_{Lo}(\neg e_{HS}(x)))$$

Q4 397 of the formula from Example 2.1, which expresses that the low-level attacker Lo never learns any
398 information about the high-level secret held in the local state of the high sender HS . However, the
399 formula

$$\text{Ded}(G) = \Diamond(\exists x.(D_G(e_{HS}(x)))$$

400 states that group G does eventually learn the value of the secret held by HS . (Note that the for-
401 mula $D_G(e_{HS}(x))$ says that the group G has distributed knowledge that the local state of com-
402 ponent HS is the same as the local state of HS in the global state denoted by variable x . The
403 formulas NonDed and $\text{Ded}(\{Lo\})$ are not opposites, as one might expect from the names. Actu-
404 ally, the negation of NonDed and $\text{Ded}(\{Lo\})$ lead to similar formulas, except that the former has a
405 negation before $e_{HS}(x)$.) Clearly, for cryptography to be effective, we require that the specification
406 $\text{NonDed} \wedge \text{Ded}(\{HR\})$ be satisfied, which expresses that the High receiver HR eventually learns the
407 secret, but that the adversary Lo never has any information about the secret.

408 In what follows, given a joint strategy α , we write $\Sigma(\alpha)$ for the singleton set of strategies $\{\alpha\}$.

409 Suppose first that encryption is always done using the action $c := s \oplus k$, so that the joint strategy
410 is the strategy α^0 from Example 2.3, with $\alpha_{Cr}^0(s) = \{c := s \oplus k\}$ for all states s . Then we work in
411 the interpreted system generated by the set of strategies $\Sigma(\alpha^0) = \{\alpha^0\}$. Note that in $\mathcal{I}(E, \Sigma(\alpha^0))$, it
412 is common knowledge that the strategy being used by Cr is α_{Cr}^0 . The following result shows that
413 in this case, the system satisfies the specification $\text{NonDed} \wedge \text{Ded}(\{HR\})$.

414 PROPOSITION 1. $E, \Sigma(\alpha^0) \models \text{NonDed} \wedge \text{Ded}(\{HR\})$.

415 PROOF. We first show that $E, \Sigma(\alpha^0) \models \text{NonDed}$. Note that, since there is only one joint strategy,
416 and agents' observations are derived from the state of the environment, a run r of $\mathcal{I}(E, \Sigma(\alpha^0))$

is determined by the sequence of states of the environment $r_e[0 \dots] = r_e(0), r_e(1) \dots$. These sequences all have the form

$$\langle s, k, \perp \rangle \langle s, k, s \oplus k \rangle^\infty$$

for some $s, k \in \{0, 1\}$, where t^∞ indicates infinitely many copies of the state t . For each run r of this form, there exists another run r' with $r'_e[0 \dots] = \langle \bar{s}, \bar{k}, \perp \rangle \langle \bar{s}, \bar{k}, \bar{s} \oplus \bar{k} \rangle^\infty$. Now, we have that $(r, n) \sim_{Lo} (r', n)$ for all $n \in \mathbb{N}$, since

$$r_{Lo}(0) = O_{Lo}(\langle s, k, \perp \rangle) = \perp = O_{Lo}(\langle \bar{s}, \bar{k}, \perp \rangle) = r'_{Lo}(0)$$

and

$$\begin{aligned} r_{Lo}(n) &= O_{Lo}(\langle s, k, s \oplus k \rangle) \\ &= s \oplus k \\ &= \bar{s} \oplus \bar{k} \\ &= O_{Lo}(\langle \bar{s}, \bar{k}, \bar{s} \oplus \bar{k} \rangle) \\ &= r'_{Lo}(n) \end{aligned}$$

for $n \geq 1$. Since also $(r, n) \sim_{Lo} (r, n)$, we have that Lo considers both possible values of the local state of agent HS possible, so $\mathcal{I}(E, \Sigma(\alpha^0)), (r, 0) \models \text{NonDed}$.

However, we have $\mathcal{I}(E, \Sigma(\alpha^0)), (r, 0) \models \text{Ded}(\{HR\})$. For, at time 1, we have $r_{HR}(1) = O_{HR}(\langle s, k, s \oplus k \rangle) = \langle k, s \oplus k \rangle$. Let (r', m) be any point with $(r, 1) \sim_{HR} (r', m)$, and let $r'_e[0 \dots] = \langle s', k', \perp \rangle \langle s', k', s' \oplus k' \rangle^\infty$. Then $m \geq 1$ and

$$r'_{HR}(m) = O_{HR}(\langle s', k', s' \oplus k' \rangle) = \langle k', s' \oplus k' \rangle.$$

Thus, from $r_{HR}(1) = r'_{HR}(m)$, we obtain $k = k'$ and $s \oplus k = s' \oplus k'$. Hence also $r'_{HS}(m) = s' = (s' \oplus k') \oplus k' = (s \oplus k) \oplus k = s = r_{HS}(1)$. This shows that $\mathcal{I}(E, \Sigma(\alpha^0)), (r, 1) \models \exists x(K_{HR}(e_{HS}(x)))$, so $\mathcal{I}(E, \Sigma(\alpha^0)), (r, 0) \models \text{Ded}(\{HR\})$. \square

However, not every strategy for the encryption agent similarly satisfies the specification. Consider the joint strategy α^2 from Example 2.3. Here we have that Lo and HR both always learn the value of the secret.

PROPOSITION 2. $E, \Sigma(\alpha^2) \models \text{Ded}(\{HR\}) \wedge \text{Ded}(\{Lo\})$.

PROOF. Strategy α^2 is deterministic. Note that if $k = 0$, then $s \oplus k = s$, and if $k = 1$, then $s \oplus \bar{k} = s$. Thus, the runs of α^2 have sequence of environment states $r_e[0 \dots] = \langle s, k, \perp \rangle \langle s, k, s \rangle^\infty$. As above, since $\Sigma(\alpha^2)$ is a singleton, this sequence determines the run as a whole. Since $O_{Lo}(\langle s, k, s \rangle) = s$ and $O_{HR}(\langle s, k, s \rangle) = \langle k, s \rangle$, both Lo and HR directly observe the value of the secret s in the local state of HS from time 1, so know this value. \square

A corollary of this result is that if we work in a system where all (uniform) strategies for Cr are possible (represented by the set of strategies Σ^{unif}), then while Lo cannot deduce the secret in general, there are encryption strategies for Cr such that, if Lo knew that this strategy is being applied by Cr , then Lo would be able to deduce the secret.

PROPOSITION 3. $E, \Sigma^{unif} \models \text{NonDed}$, but not $E, \Sigma^{unif} \models \neg \text{Ded}(\{Lo, \sigma(Cr)\})$.

PROOF. For $E, \Sigma^{unif} \models \text{NonDed}$, we note that Lo always considers it possible that Cr is running strategy α^0 from above and argue exactly as in Proposition 1. To show that not $E, \Sigma^{unif} \models \neg \text{Ded}(\{Lo, \sigma(Cr)\})$, let r be a run in which Cr runs strategy α^2_{Cr} . Note that if $(r, 0) \sim_{\{Lo, \sigma(Cr)\}} (r', m)$, then $r_{\sigma(Cr)}(0) = r'_{\sigma(Cr)}(m)$, i.e., Cr uses the same strategy in the runs r and r' . Essentially

449 the same argument as applied in Proposition 2 to show that $Ded(\{Lo\})$ holds then shows that
 450 $\mathcal{I}(E, \Sigma^{unif}), (r, 0) \models Ded(\{Lo, \sigma(Cr)\})$. \square

451 By means of a similar example, Wittbold and Johnson [57] argued that nondeducibility is too
 452 weak a notion of security to capture information flow security attacks in which the attacker ex-
 453 ploits a covert channel in a system. Intuitively, it does not take into account that the attacker
 454 may have information about the strategies being used by other agents. One example of how such
 455 knowledge of another agent’s strategy may arise in practice is when the attacker Lo has succeeded
 456 in infiltrating a virus (here represented by the strategy of Cr) into the system being attacked (here
 457 composed of components HS , HR , and Cr , i.e., the High sender, the High receiver, and the encrypt-
 458 tion agent, respectively). When this is the case, a more appropriate modality for the attacker’s
 459 knowledge is the modality $D_{\{Lo, \sigma(Cr)\}}$, which captures what Lo can deduce when it also knows
 460 the strategy $\sigma(Cr)$ being employed by Cr rather than the modality $D_{\{Lo\}}$ used in $Ded(\{Lo\})$. (The
 461 modality $D_{\{Lo, \sigma(Lo), \sigma(Cr)\}}$ that says that Lo also reasons knowing its own strategy would also make
 462 sense in general, though in the model under discussion it is identical to $D_{\{Lo, \sigma(Cr)\}}$, since Lo has
 463 only one action to choose from, so all its uniform strategies are the same.) Wittbold and Johnson’s
 464 notion of *nondeducibility on strategies* (NDS) is a definition of security that takes into account such
 465 reasoning by the attacker. For a two-agent system, composed of Low-level agent Lo and High-level
 466 agent Hi , Wittbold and Johnson define a system to satisfy non-deducibility on strategies if every
 467 Low view is compatible with every High strategy. NDS may be expressed directly in our logic by
 468 the formula⁵

$$D_0 \forall x. (\neg K_{Lo}(\neg e_{\sigma(Hi)}(x))),$$

469 which says that at all points of the system (identifying a Lo view/local state, in particular) for
 470 all global states x (identifying a High strategy, in particular), Lo considers the High strategy in x
 471 to be possible. This notion cannot be expressed in alternating temporal epistemic logics such as
 472 ATEL, discussed below, which do not allow reference to what can be deduced about other agents’
 473 strategies.

474 3.2 Revocable and Irrevocable Strategies in ATL

475 ATL [2] is a generalization of the branching-time temporal logic CTL that can express the capability
 476 of agents’ strategies to bring about temporal effects. We show in this section that ESL is able to
 477 express several variants of ATL. The following section relates various epistemic extensions of ATL
 478 to ESL.

479 The syntax of ATL formulas ϕ is given as follows:

$$\phi \equiv p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid \langle\langle G \rangle\rangle\phi \mid \langle\langle G \rangle\rangle\Box\phi \mid \langle\langle G \rangle\rangle(\phi_1 U \phi_2),$$

480 where $p \in Prop$, $i \in Ags$, and $G \subseteq Ags$. Essentially, each branching construct $A\phi$ of CTL is general-
 481 ized in ATL to an *alternating* construct $\langle\langle G \rangle\rangle\phi$ for a group G of agents, where ϕ is a “prefix temporal”
 482 formula such as $\Box\phi'$, $\Diamond\phi'$, $\Box\phi'$, or $\phi_1 U \phi_2$, as would be used to construct a CTL formula. Intuitively,
 483 $\langle\langle G \rangle\rangle\phi$ says that the group G has a strategy for ensuring that ϕ holds, irrespective of what the other
 484 agents do.

485 The semantics of ATL is given using *concurrent game structures*, which are very similar to en-
 486 vironments as defined above, with the main differences being the following. For each point of
 487 difference, we sketch how to view concurrent game structures as equivalent to environments.

⁵The perfect recall semantics in combination with perfect recall strategies would give the interpretation of this formula that is most adequate for security applications.

- Concurrent game structures lack a set of initial states. It is convenient for technical reasons 488
to treat a concurrent game structure as an environment with all of its states initial. 489
- Concurrent game structures allow that not all actions are available at every state, whereas 490
in environments all actions are always available. In environments, we can treat a choice of 491
a non-enabled action as equivalent to a choice of a default enabled action in the transition 492
relation. 493
- The transition relation in concurrent game structures is deterministic, in the sense that for 494
each state s and joint action a , there exists a unique state t such that $s \xrightarrow{a} t$. Nondetermin- 495
ism in environments can be modelled in concurrent game structures by adding an agent 496
that makes the nondeterministic choice through its actions. 497
- ATL's concurrent game structures do not have a notion of observation. Intuitively, all agents 498
always have perfect information concerning the current state. We may capture this in en- 499
vironments by taking $O_i(s) = s$ for all agents i and states s . 500

Using such correspondences, we can express the ATL semantics in environments E as follows. 501
For reasons discussed below, we generalize the ATL semantics by parameterizing the definition on 502
a set Δ of strategies for groups of agents in the environment E . That is, Δ is a collection of tuples 503
of agent strategies of the form $\langle \alpha_i \rangle_{i \in G}$, with both the strategies α_i and the set G of agents varying. 504
The semantics uses a relation $E, s \models^\Delta \phi$, where $E = \langle S, I, Acts, \rightarrow, \{O_i\}_{i \in Ags}, \pi \rangle$ is an environment 505
and $s \in S$ is a state of E , and ϕ is a formula. 506

For the definition, we need the notion of a path in E : This is a function $\rho : \mathbb{N} \rightarrow S$ such that for all 507
 $k \in \mathbb{N}$ there exists a joint action a with $(\rho(k), a, \rho(k+1)) \in \rightarrow$. A path ρ is *from* a state s if $\rho(0) = s$. 508
A path ρ is *consistent* with a strategy $\alpha = \langle \alpha_i \rangle_{i \in G}$ for a group G if for all $k \in \mathbb{N}$ there exists a joint 509
action a such that $(\rho(k), a, \rho(k+1)) \in \rightarrow$ and $a_i \in \alpha_i(\rho(k))$ for all $i \in G$. It is also convenient to 510
identify the *path formulas* of ATL as formulas of the form $\circ\phi$, $\square\phi$ or $\phi U \psi$ where ϕ and ψ are ATL 511
formulas. 512

The relation $E, s \models^\Delta \phi$, where s is a state of E and ϕ is an ATL formula, is defined by a mutual 513
recursion with the relation $E, \rho \models^\Delta \phi$, where ρ is a path of E and ϕ is a path formula, as follows. 514
Note that if $\langle\langle G \rangle\rangle\phi$ is an ATL formula, then ϕ is a path formula. For evaluation of ATL formulas at 515
a state we have the clauses 516

- $E, s \models^\Delta p$ if $p \in \pi(s)$; 517
- $E, s \models^\Delta \neg\phi$ if not $E, s \models^\Delta \phi$; 518
- $E, s \models^\Delta \phi \wedge \psi$ if $E, s \models^\Delta \phi$ and $E, s \models^\Delta \psi$; 519
- $E, s \models^\Delta \langle\langle G \rangle\rangle\phi$ if there exists a strategy $\alpha_G \in \Delta$ for group G such that for all paths ρ from s 520
that are consistent with α_G , we have $E, \rho \models^\Delta \phi$; 521

and for evaluation of a path formula at a path we have the clauses 522

- $E, \rho \models^\Delta \circ\phi$ if $E, \rho(1) \models^\Delta \phi$; 523
- $E, \rho \models^\Delta \square\phi$ if have $E, \rho(k) \models^\Delta \phi$ for all $k \in \mathbb{N}$; 524
- $E, \rho \models^\Delta \phi U \psi$ if there exists $m \geq 0$ such that $E, \rho(m) \models^\Delta \psi$, and for all $k < m$, we have 525
 $E, \rho(k) \models^\Delta \phi$. 526

The semantics for ATL given in Reference [2] corresponds to the instance of this definition with 527
 Δ equal to the set of perfect recall, perfect information group strategies, but we focus here on the 528
variant where Δ contains just imperfect information strategies. 529

We argue that the ATL construct $\langle\langle G \rangle\rangle\phi$ can be expressed in $CTL^*K(Prop, Ags \cup \sigma(Ags))$ as 530

$$\neg K_e \neg D_{\{e\} \cup \sigma(G)} \phi.$$

531 Intuitively, here the outer operator $\neg K_e \neg$ existentially switches to a point that has the same state
 532 of the environment as the current state (and hence the same local state for all agents in Ags) but
 533 may have different strategies for any of the agents. The inner operator $D_{\{e\} \cup \sigma(G)}$ then fixes both
 534 the state of the environment and the strategies selected by the group G but allows all other agents
 535 to vary their strategy. It quantifies universally over these possibilities. Thus, the formula says that
 536 the group G has a strategy that achieves ϕ from the current state, whatever strategy the other
 537 agents play. (An alternate way to express the formula using the richer expressive power of ESL is
 538 as $\exists x(K_e(e_{\sigma(G)}(x) \Rightarrow \phi))$.)

539 More formally, consider the following translation from ATL to $CTL^*K(Prop, Ags \cup \sigma(Ags))$. For
 540 an ATL formula ϕ , we write ϕ^* for the translation of ϕ , defined inductively on the construction of
 541 ϕ by the following rules:

$$\begin{aligned}
 p^* &= p \\
 (\neg \phi)^* &= \neg \phi^* \\
 (\phi_1 \wedge \phi_2)^* &= \phi_1^* \wedge \phi_2^* \\
 \langle\langle G \rangle\rangle \phi^* &= \neg K_e \neg D_{\{e\} \cup \sigma(G)} \phi^* \\
 (\bigcirc \phi)^* &= \bigcirc \phi^* \\
 (\square \phi)^* &= \square \phi^* \\
 (\phi_1 U \phi_2)^* &= \phi_1^* U \phi_2^*.
 \end{aligned}$$

542 Note that the semantics of the operators using $\langle\langle G \rangle\rangle$ quantifies over runs in which the agents G
 543 run a particular strategy α_G , but there is no constraint on the behaviour of the other agents: These
 544 are not assumed to choose their actions according to any particular strategy. A natural alternative
 545 to the definition above would be to use the clause

546 $E, s \models^\Delta \langle\langle G \rangle\rangle \phi$ if there exists a strategy $\alpha \in \Delta$ for group G such that for all joint strategies
 547 $\beta \in \Delta$ for group Ags with $\beta \upharpoonright G = \alpha$, and all paths ρ from s that are consistent with β , we
 548 have $E, \rho \models^\Delta \phi$.

549 This variant corresponds more directly to the formula $\neg D_e \neg D_{\{e\} \cup \sigma(G)} \square \phi$ than does the ATL
 550 semantics. It is reasonable to take the position that it more naturally captures a concept of interest
 551 in competitive situations where agents are constrained in the strategies they are able to use.

552 In the original semantics of ATL, where perfect information, perfect recall strategies were con-
 553 sidered, the two definitions are equivalent, since for any behaviour of the other agents, there is
 554 a strategy that matches it. However, for the imperfect information, epistemic extension we con-
 555 sider, this does not hold. For example, if all strategies in Δ are deterministic, then the above variant
 556 would not allow paths in which some agent in the complement of G chooses an action a at the
 557 first occurrence of a state s , but some other action b at a later occurrence of s . However, such runs
 558 are allowed in the ATL semantics given above. Since the semantics of ESL assumes that all runs
 559 are generated by all agents running some strategy, we need to make some technical assumptions
 560 on Δ to set up a correspondence with ATL.

561 Define the “random” strategy for agent i to be the strategy $rand_i$ defined by $rand_i(s) = Acts_i$ for
 562 all states $s \in S$. Given a strategy $\alpha = \langle \alpha_i \rangle_{i \in G}$ for a group of agents G in an environment E , define
 563 the *completion* of the strategy to be the joint strategy $comp(\alpha) = \langle \alpha'_i \rangle_{i \in Ags}$ with $\alpha'_i = \alpha_i$ for $i \in G$
 564 and with $\alpha'_i = rand_i$ for all $i \in Ags \setminus G$. Intuitively, this operation completes the group strategy to
 565 a joint strategy for all agents, by adding the “random” strategy for all agents not in G , so that these
 566 agents are completely unconstrained in their behaviour. Given a set of strategies Δ for groups of
 567 agents, we define the set of joint strategies $comp(\Delta) = \{comp(\alpha) \mid \alpha \in \Delta\}$.

A second technicality is needed that results from the way we have used Δ as a parameter in a generalization of the ATL semantics. A constraint on this set is needed to prove our correspondence result. Say that a set Δ of group strategies is *restrictable* if for every $\alpha \in \Delta$ for group of agents G and every group $H \subseteq G$, the restriction $\alpha \upharpoonright H$ of α to agents in H is also in Δ . Say that Δ is *extendable* if for every strategy α for a group H and group $G \supseteq H$, there exists a strategy $\alpha' \in \Delta$ for group G whose restriction $\alpha' \upharpoonright H$ to H is equal to α . Intuitively, restrictability says that group strategies are closed under formation of subgroups, and extensibility says that a group is not able to prevent any other agent from having *some* strategy that they are able to follow at the same time as the group follows its choice of strategy.

The requirement that a set Δ of group strategies be restrictable and extendable is quite mild. For example, if Δ_i is a set of strategies for agent i , for each agent $i \in \text{Ags}$, then the natural set of “Cartesian product strategies”

$$\Delta = \{\langle \alpha_i \rangle_{i \in G} \mid G \subseteq \text{Ags}, \forall i \in G (\alpha_i \in \Delta_i)\}$$

is both restrictable and extendable. In particular, the set of all group strategies, and the set of all locally uniform group strategies, are both restrictable and extendable. Another example of a collection of strategies satisfying this condition is the set of group strategies α in which at most k agents follow a strategy that differs from a designated “correct” strategy σ . Note that this collection is extendable, because an agent always has the option to choose the correct strategy, even if k others have already deviated. This collection models a common assumption in the analysis of fault-tolerant distributed algorithms.

A final technicality relates to the fact that whereas runs of an environment start at an initial state of the environment, and hence an environment may have unreachable states, models in the ATL semantics lack a notion of initial state, and formulas may be evaluated at any state. As already noted above, we resolve this difference by viewing ATL models as environments in which all states are initial (hence reachable).

The following result now captures in a precise way that the ATL semantics can be expressed in our logic as claimed above, provided we allow joint strategies in which some agents run the random strategy.

THEOREM 3.1. *For every environment E in which all states are initial, for every nonempty set of group strategies Δ that is restrictable and extendable, for every state s of E and ATL formula ϕ , we have $E, s \models^\Delta \phi$ iff for all (equivalently, some) points (r, m) of $\mathcal{I}(E, \text{comp}(\Delta))$ with $r_e(m) = s$ we have $\mathcal{I}(E, \text{comp}(\Delta)), (r, m) \models \phi^*$.*

PROOF. For brevity, we write just \mathcal{I} for $\mathcal{I}(E, \text{comp}(\Delta))$. For the claim that the quantifiers “for all” and “some” are interchangeable in the right-hand side, note that formulas of the form ϕ^* are Boolean combinations of atomic propositions and formulas of the form $K_e \psi$, whose semantics at a point (r, m) depends only on $r_e(m)$. This gives the implication from the “some” case to the “for all” case. For the implication from the “for all” case to the “some” case, note that the “for all” case is never trivial, because for all states s of E , there exists a point (r, m) of \mathcal{I} with $r_e(m) = s$. This follows from the fact that all states are initial in E and that the transition relation is serial, so that any group strategy α in Δ is consistent with an infinite path from initial state s . This corresponds to a run r with $r(0) = (s, \text{comp}(\alpha))$.

It therefore suffices to show that $E, s \models^\Delta \phi$ iff for all points (r, m) of \mathcal{I} with $r_e(m) = s$ we have $\mathcal{I}, (r, m) \models \phi^*$. Additionally, for path subformulas ϕ of the form $\bigcirc \psi$, $\square \psi$ and $\psi_1 U \psi_2$ of ATL formulas, we show that for all paths ρ , we have $E, \rho \models^\Delta \phi$ iff for all points (r, m) of \mathcal{I} with $r_e[m \dots] = \rho$ we have $\mathcal{I}, (r, m) \models \phi^*$

612 We proceed by induction on the construction of ϕ . The base case of atomic propositions, as well
 613 as the cases for the Boolean constructs, are trivial. The claim concerning path formulas is also
 614 straightforward from the semantics of the temporal operators and, inductively, the claim concern-
 615 ing state formulas.

616 We consider next the case of $\phi = \langle\langle G \rangle\rangle\psi$. We show that $E, s \models^\Delta \langle\langle G \rangle\rangle\psi$ iff for all points (r, m) of
 617 \mathcal{I} with $r_e(m) = s$ we have $\mathcal{I}, (r, m) \models \neg K_e \neg D_{\{e\} \cup \sigma(G)} \phi^*$.

618 Suppose, first, $E, s \models^\Delta \langle\langle G \rangle\rangle\psi$. Let (r, m) be a point of \mathcal{I} with $r_e(m) = s$. We show that $\mathcal{I}, (r, m) \models$
 619 $\neg K_e \neg D_{\{e\} \cup \sigma(G)} \psi^*$. By the ATL semantics, there exists a strategy $\alpha_G \in \Delta$ for group G such that for
 620 all paths ρ of E from s that are consistent with α_G we have $E, \rho \models^\Delta \psi$. Let $\alpha = \text{comp}(\alpha_G)$ (note that
 621 this is in $\text{comp}(\Delta)$), and (using the fact that all states are initial) let r' be a run of \mathcal{I} with $r'(0) =$
 622 (s, α) . Because $r_e(m) = s = r'_e(0)$, we have $(r, m) \sim_e (r', 0)$, and it suffices to show that $\mathcal{I}, (r', 0) \models$
 623 $D_{\{e\} \cup \sigma(G)} \psi^*$. For this, suppose that (r'', m'') is any point of \mathcal{I} with $(r', 0) \sim_{\{e\} \cup \sigma(G)} (r'', m'')$. We
 624 show that $\mathcal{I}, (r'', m'') \models \psi^*$. Now $r''(m'') = (t, \alpha')$ implies that $\alpha'_i = \alpha_i$ for all $i \in G$. Thus, the path
 625 $\rho = r''_e(m'')r''_e(m'' + 1) \dots$ in E is consistent with α_G , and $\rho(0) = r''_e(m'') = r'_e(0) = s$. It follows
 626 that $E, \rho \models^\Delta \psi$. Using the induction hypothesis, it follows that $\mathcal{I}, (r'', m'') \models \psi^*$. This completes
 627 the argument that $\mathcal{I}, (r', 0) \models D_{\{e\} \cup \sigma(G)} \psi^*$.

628 Conversely, suppose that for all points (r, m) of \mathcal{I} with $r_e(m) = s$ we have $\mathcal{I}, (r, m) \models$
 629 $\neg K_e \neg D_{\{e\} \cup \sigma(G)} \psi^*$. We show that $E, s \models^\Delta \langle\langle G \rangle\rangle\psi$. Using the fact that all states are initial, let r be a run
 630 of \mathcal{I} with $r_e(0) = s$, and hence, $\mathcal{I}, (r, 0) \models \neg K_e \neg D_{\{e\} \cup \sigma(G)} \psi^*$. Then there exists a point (r', m') of \mathcal{I}
 631 such that $r'_e(m') = s$ and $\mathcal{I}, (r', m') \models D_{\{e\} \cup \sigma(G)} \psi^*$. Let $r'(m') = (s, \alpha)$. Then there exists a strategy
 632 $\beta \in \Delta$ for some set of agents G' such that $\alpha = \text{comp}(\beta) \in \text{comp}(\Delta)$. Let $H = G \cap G'$. By restrictabil-
 633 ity, we have $\beta \upharpoonright H \in \Delta$. By extendability, there exists a strategy $\gamma \in \Delta$ for group G such that
 634 $\gamma \upharpoonright H = \beta \upharpoonright H$. Taking $\alpha' = \text{comp}(\beta \upharpoonright H)$, it follows that $\alpha' \in \text{comp}(\Delta)$. Note that $\alpha' \upharpoonright G = \alpha \upharpoonright G$
 635 and $\alpha'_i = \text{rand}_i$ for $i \in \text{Ags} \setminus G$. In particular, $\alpha'_i = \text{rand}_i$ for $i \in G \setminus H$. Thus, any path consistent
 636 with γ is consistent with α' .

637 To prove that $E, s \models^\Delta \langle\langle G \rangle\rangle\psi$, we show that for every path ρ of E from s consistent with γ , we
 638 have $E, \rho \models^\Delta \psi$. For this, let ρ be a path from s consistent with the strategy γ for group G . By the
 639 conclusion of the previous paragraph, ρ is consistent with the joint strategy α' for all agents. Since s
 640 is an initial state of E , there exists a run r'' of \mathcal{I} with $r''(0) = (s, \alpha')$ and $r''_e[0 \dots \infty] = \rho$. Moreover,
 641 $(r', m') \sim_{\{e\} \cup \sigma(G)} (r'', 0)$. Thus, we obtain from $\mathcal{I}, (r', m') \models D_{\{e\} \cup \sigma(G)} \psi^*$ that $\mathcal{I}, (r'', 0) \models \psi^*$. By
 642 the induction hypothesis, we obtain that $E, \rho \models^\Delta \psi$. \square

643 The ESL interpretation unpacks the alternating double quantification in the semantics of $\langle\langle G \rangle\rangle\phi$.
 644 ESL offers the advantage of being able to express notions that are not expressible in ATL. For
 645 example, under assumptions similar to those of Theorem 3.3,

$$\neg D_e \neg ((\neg D_{\{e\} \cup \sigma(\text{Ags})} \neg \diamond p) \wedge (D_{\{e\} \cup \sigma(G)} \Box q))$$

646 says that, from the current state, there is a joint strategy for all agents, such that, some runs of this
 647 joint strategy satisfy $\diamond p$, and group G 's strategy alone suffices to ensure that $\Box q$.

648 There has been discussion in the literature on ATL about whether strategies should be *revocable*
 649 or *irrevocable*. Consider a formula such as

$$\langle\langle A \rangle\rangle \Box (p \wedge \langle\langle A, B \rangle\rangle \diamond q).$$

650 This says that A has a strategy that ensures that it is always the case both that p holds and that A and
 651 B together have a strategy that ensures that eventually q . Under the ATL semantics, the strategy
 652 of A used to satisfy the inner formula $\langle\langle A, B \rangle\rangle \diamond q$ is allowed to be different from the strategy of A
 653 referred to by the outer operator. That is, to satisfy the inner formula, A is allowed to *revoke* the
 654 strategy selected by the outer operator.

This aspect of the ATL semantics has been questioned [1], and it has been proposed that the semantics of the formula $\langle\langle G \rangle\rangle\phi$ should be defined so that it fixes the strategies of agents in the group G and does not allow these to be varied in interpreting operators in the formula ϕ . In such a semantics, the strategy choices are *irrevocable*. Using our framework, both revocable and irrevocable interpretations of the formula can be represented. We show this with two formulas that are almost identical, with the point of difference indicated by use of bold type. The interpretation allowing strategy revocation would be captured by translating both operators as described above, yielding the formula

$$\neg D_e \neg D_{\{e, \sigma(A)\}} (\Box p \wedge \neg \mathbf{D}_e \neg D_{\{e, \sigma(A), \sigma(B)\}} \Diamond q).$$

Note that here the outer operator prefix $\neg D_e \neg D_{\{e, \sigma(A)\}}$ selects a strategy for A and plays it against all strategies of the other agents, and because the operator \mathbf{D}_e allows all agent's strategies to vary, the inner operator prefix $\neg \mathbf{D}_e \neg D_{\{e, \sigma(A), \sigma(B)\}}$ drops the selected strategy of A and selects a fresh strategy for A and B together to play against all strategies of other agents. However, we can force the strategy of agent A to remain fixed in the inner choice of strategies by means of the formula

$$\neg D_e \neg D_{\{e, \sigma(A)\}} (\Box p \wedge \neg \mathbf{D}_{\{e, \sigma(A)\}} \neg D_{\{e, \sigma(A), \sigma(B)\}} \Diamond q).$$

Note that the inner operator $\mathbf{D}_{\{e, \sigma(A)\}}$ varies all agent's strategies, except that of A . Evidently, at any point in a nested formula, our approach gives us the freedom to choose which players' strategies we wish to vary and which to fix.

A logic with revocable strategies is presented in Brihaye et al. [7], which considers the extension of ATL with strategy context, or ATL_{sc} . Formulas are evaluated with respect to a context that is a group strategy γ_G for some group G . The logic has modalities $\cdot)H\langle\cdot\phi$, and $\langle\cdot H\cdot\rangle\phi$. Intuitively, $\cdot)H\langle\cdot\phi$ reduces the context group G to $G \setminus H$ by restricting γ_G to $G \setminus H$. The modality $\langle\cdot H\cdot\rangle\phi$ selects a new group strategy γ_H for group H and constructs the new context $\gamma_H \circ \gamma_G$ for group $G \cup H$ in which agents i in H play $\gamma_H(i)$, and agents i in $G \setminus H$ play $\sigma_G(i)$. The formula ϕ is then evaluated with respect to context $\gamma_H \circ \gamma_G$ in all runs in which $G \cup H$ plays $\gamma_H \circ \gamma_G$ against an arbitrary behaviour of all other agents.

Evaluation of formulas commences with respect to the empty context, so each subformula is evaluated with respect to a context for a group G that can be determined from the operators on the path from the root to that subformula. This means that to represent a formula ϕ of ATL_{sc} , we need to translate it with respect to a group G ; we write the translation as ϕ^G . Roughly, with respect to a context for group G , the formula $\langle\cdot H\cdot\rangle\phi$ can then be expressed with our logic as

$$(\langle\cdot H\cdot\rangle\phi)^G = \neg D_{\{e, \sigma(G \setminus H)\}} \neg D_{\{e, \sigma(G \cup H)\}} \phi^{G \cup H},$$

and the formula $\cdot)H\langle\cdot\phi$ can be expressed as

$$(\cdot)H\langle\cdot\phi)^G = \phi^{G \setminus H}.$$

However, we note that the semantics in Reference [7] is based on perfect recall. This explains that the complexity of model checking ATL_{sc} is non-elementary, while the complexity of model checking our logic ESL is EXPSPACE-complete (Theorem 4.1 and Theorem 4.2).

Q5

Another work by van der Hoek, Jamroga, and Wooldridge [28] introduces constants that refer to strategies and adds to ATL a new (counterfactual) modality $C_i(c, \phi)$, with the intended reading "if it were the case that agent i committed to the strategy denoted by c , then ϕ ." (The meaning of c is bound in the semantic context, and the logic does not allow quantification over c .) The formula ϕ here is not permitted to contain further references to agent i strategies. To interpret the formula $C_i(c, \phi)$ in an environment E , the environment is first updated to a new environment E' by removing all transitions that are inconsistent with agent i running the strategy referred to by c , and then the formula ϕ is evaluated in E' . In ESL, the assertion that i is running a particular strategy

696 can be made by the formula $e_{\sigma(i)}(x)$, where x is taken to denote a global state in which the local
 697 component $\sigma(i)$ denotes the strategy denoted by c . The formula $C_i(c, \phi)$ can then be expressed in
 698 our framework as

$$D_{\{e\} \cup \sigma(\text{Ags} \setminus \{i\})}(e_{\sigma(i)}(x) \Rightarrow \phi^{+\sigma(i)}),$$

699 where in the translation $\phi^{+\sigma(i)}$ of ϕ we ensure that there is no further deviation from the strategy of
 700 agent i by adding $\sigma(i)$ to the group of every knowledge operator occurring later in the translation.
 701 We remark that because it deletes information from the transition relation, strategy choices made
 702 by the construct $C_i(c, \phi)$ are irrevocable, whereas our logic is richer in that it allows revocation of
 703 the corresponding choices.

704 3.3 Connections to Variants of ATEL

705 Alternating temporal epistemic logic (ATEL) adds epistemic operators to ATL [29]. As a number
 706 of subtleties arise in the formulation of such logics, several variants of ATEL have since been
 707 developed. In this section, we consider a number of such variants and argue that our framework
 708 is able to express the main strategic concepts from these variants. We begin by recalling ATEL as
 709 defined in Reference [29].

710 The syntax of ATEL is given as follows:

$$\phi \equiv p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid \langle\langle G \rangle\rangle\phi \mid \langle\langle G \rangle\rangle\Box\phi \mid \langle\langle G \rangle\rangle(\phi_1 U \phi_2) \mid K_i\phi \mid D_G\phi \mid C_G\phi,$$

711 where $p \in Prop$, $i \in Ags$, and $G \subseteq Ags$. This just adds the operators K_i , D_G , and C_G to the syntax
 712 for ATL given above. As usual, we may define $E_G\phi$ as $\bigwedge_{i \in G} K_i\phi$. The intuitive meaning of the
 713 constructs is as in CTL*K above, with, additionally, $\langle\langle G \rangle\rangle\phi$ having the intuitive reading that group
 714 G has a strategy for assuring that ϕ holds.

715 The relation $E, s \models^\Delta \phi$ is extended from ATL to ATEL by adding the following clauses to the
 716 inductive definition:

- 717 • $E, s \models^\Delta K_i\phi$ if $E, t \models^\Delta \phi$ for all $t \in S$ with $t \sim_i s$;
- 718 • $E, s \models^\Delta D_G\phi$ if $E, t \models^\Delta \phi$, for all $t \in S$ with $(s, t) \in \bigcap_{i \in G} \sim_i$;
- 719 • $E, s \models^\Delta C_G\phi$ if $E, t \models^\Delta \phi$ for all $t \in S$ with $(s, t) \in (\bigcup_{i \in G} \sim_i)^*$;

720 where we define, for each $i \in Ags$, the equivalence relation \sim_i on states S by $s \sim_i t$ if and only if
 721 $O_i(s) = O_i(t)$.

722 The specific version of ATEL defined in Reference [29] is obtained from the above defini-
 723 tions by taking $\Delta = \{\sigma_G \mid G \subseteq Ags, \sigma_G \text{ a deterministic } G\text{-strategy in } E\}$. That is, following the def-
 724 initions for ATL, this version works with arbitrary deterministic group strategies, in which an
 725 agent selects its action as if it had full information of the state. This aspect of the definition
 726 has been criticized by Jonker [39] and (in the case of ATL without epistemic operators) by
 727 Schobbens [50], who argue that this choice is not in the spirit of the epistemic extension, in
 728 which observations are intended precisely to represent that agents do not have full informa-
 729 tion of the state. They propose that the definition instead be based on the set $\Delta = \{\sigma_G \mid G \subseteq$
 730 $Ags, \sigma_G \text{ a locally uniform deterministic } G\text{-strategy in } E\}$. This ensures that in choosing an action,
 731 agents are able to use only the information available in their observations.

732 We concur that the use of locally uniform strategies is the more appropriate choice, but in either
 733 event, we now argue that our approach using strategy space is able to express everything that can
 734 be expressed in ATEL. We may extend the translation into our logic given above from ATL to ATEL
 735 by adding the following rules:

$$(K_i\phi)^* = K_i\phi^* \quad (D_G\phi)^* = D_G\phi^* \quad (C_G\phi)^* = C_G\phi^*.$$

To obtain a correspondence with ATEL, which does not have a notion of initial states, we again work with environments in which all states are initial. The following result shows that Theorem 3.1 extends from ATL to the logic ATEL.

THEOREM 3.2. *For every environment E in which all states are initial, for every nonempty set of group strategies Δ that is restrictable, for every state s of E and ATEL formula ϕ , we have $E, s \models^\Delta \phi$ iff for all (equivalently, some) points (r, m) of $\mathcal{I}(E, \text{comp}(\Delta))$ with $r_e(m) = s$ we have $\mathcal{I}(E, \text{comp}(\Delta)), (r, m) \models \phi^*$.*

PROOF. The proof extends the proof of Theorem 3.1. The argument for the equivalence of the universal and existential quantifications in the right-hand side of the “iff” continues to apply, even though the translation now contains formulas of the form $K_i\phi$, because, by construction in Section 2.2, $r_e(m) = r'_e(m')$ implies $r_i(m) = r'_i(m')$. The remainder of the proof extends the inductive argument.

Consider $\phi = K_i\psi$. Then $(K_i\psi)^* = K_i\psi^*$. We suppose first that $E, s \models^\Delta \phi$ and show that for all points (r, m) of \mathcal{I} with $r_e(m) = s$ we have $\mathcal{I}, (r, m) \models \phi^*$, i.e., $\mathcal{I}, (r, m) \models K_i\psi^*$. Let (r, m) be a point of \mathcal{I} with $r_e(m) = s$. We need to show that for all points (r', m') of \mathcal{I} with $(r, m) \sim_i (r', m')$ we have $\mathcal{I}, (r', m') \models \psi^*$. But if $(r, m) \sim_i (r', m')$ then $r'_e(m') \sim_i r_e(m) = s$ in E . Thus, from $E, s \models^\Delta K_i\psi$ it follows that $E, r'_e(m') \models^\Delta \psi$. By the induction hypothesis, we obtain that $\mathcal{I}, (r', m') \models \psi^*$, as required.

Conversely, suppose that for all points (r, m) of \mathcal{I} with $r_e(m) = s$ we have $\mathcal{I}, (r, m) \models K_i\psi^*$. We show that $E, s \models^\Delta K_i\psi$. Let t be any state of E with $s \sim_i t$. We have to show $E, t \models^\Delta \psi$. First, since s is an initial state of E , there exists a run r of \mathcal{I} with $r_e(0) = s$, and joint strategy equal to any strategy in $\text{comp}(\Delta)$, so we take $m = 0$, and we have $\mathcal{I}, (r, m) \models K_i\psi^*$. Then for all points (r', m') of \mathcal{I} with $r'_e(m') = t$, we have $(r, 0) \sim_i (r', m')$, from which it follows that $\mathcal{I}, (r', m') \models \psi^*$. By the induction hypothesis, we have $E, t \models \psi$, as required. This completes the proof for the case of $\phi = K_i\psi$. The argument for the distributed and common knowledge operators is similar, and left to the reader. \square

We remark that our translation maps ATEL into $\text{CTLK}(\text{Ags} \cup \sigma(\text{Ags}), \text{Prop})$, the fragment that we show in Theorem 4.5 below to have PSPACE-complete model-checking complexity. This strongly suggests that this fragment has a strictly stronger expressive power than ATEL, since the complexity of model-checking ATEL, assuming uniform strategies, is known to be P^{NP} -complete. (The class P^{NP} consists of problems solvable by PTIME computations with access to an NP oracle.) For ATEL, model checking can be done with a polynomial time (with respect to the size of formula) computation with access to an oracle that is in NP with respect to both the number of states and the number of joint actions. In particular, Reference [50] proves this upper bound, and Reference [37] proves a matching lower bound.

Similar translation results can be given for other alternating temporal epistemic logics from the literature. We sketch a few of these translations here.

Jamroga and van der Hoek [38] discuss issue of *de dicto* and *de re* interpretations of ATEL formulas. They consider the formula $K_i\langle\langle i \rangle\rangle\phi$. (Note that here ϕ is a path formula). The ATEL semantics states that for an environment E and a state s , we have $E, s \models K_i\langle\langle i \rangle\rangle\phi$ when in every state t consistent with agent i 's knowledge, some strategy for agent i , depending on t , is guaranteed to satisfy ϕ . This is consistent with there being no *single* strategy for agent i that agent i knows will work to achieve ϕ in all such states t . To express that a single strategy is known to guarantee ϕ , they formulate a general construct $\langle\langle G \rangle\rangle_{\mathcal{K}(H)}^\bullet\phi$ that says, effectively, that there is a strategy for a group G that another group H knows (for notion of group knowledge \mathcal{K}) to achieve goal ϕ . (Here again, ϕ is a path formula.) The notion of group knowledge \mathcal{K} could be E for everyone knows, D for

781 distributed knowledge, or C for common knowledge. More precisely⁶,

$E, s \models^\Delta \langle\langle G \rangle\rangle_{\mathcal{K}(H)}^\bullet \phi$ if there exists a locally uniform group strategy $\alpha \in \Delta$ for group G such that for all states t with $s \sim_H^\mathcal{K} t$, and for all paths ρ from t that are consistent with α , we have that $E, \rho \models^\Delta \phi$.

782 Here, $\sim_H^\mathcal{K}$ is the appropriate epistemic indistinguishability relation on states of E . The particular
783 case $\langle\langle G \rangle\rangle_{E(G)}^\bullet \phi$ is also proposed as the semantics for the ATL construct $\langle\langle G \rangle\rangle \phi$ in References [35, 39,
784 50].

785 The construct $\langle\langle G \rangle\rangle_{D(H)}^\bullet \phi$ can be represented in the $\text{CTLK}(Ags \cup \sigma(Ags), Prop)$ fragment of ESL
786 as

$$\neg K_e \neg D_{H \cup \sigma(G)} \phi.$$

787 Intuitively, here the first modal operator $\neg K_e \neg$ switches the strategy of all the agents while main-
788 taining the state s , thereby selecting a strategy α for group G in particular, and the next operator
789 $D_{H \cup \sigma(G)}$ verifies that the group H knows that the strategy being used by group G guarantees ϕ .
790 Similarly, $\langle\langle G \rangle\rangle_{E(H)}^\bullet \phi$ can be represented as

$$\neg K_e \neg \bigwedge_{i \in H} D_{\{i\} \cup \sigma(G)} \phi.$$

791 The precise statement and proof of these correspondences is similar to that in Theorem 3.3.

792 In the case of the construct $\langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$, the definition involves the common knowledge that a
793 group H of agents would have if they knew a particular strategy being used by another group
794 G . By analogy with the above cases, one might expect this to be expressible using the for-
795 mula $\neg K_e \neg C_{H \cup \sigma(G)} \phi$. However, this does not give the intended meaning. Note that the seman-
796 tics of the formula $C_{H \cup \sigma(G)} \phi$ quantifies over points (r', m') reachable through chains $(r, m) =$
797 $(r_0, m_0) \sim_{i_1} (r_1, m_1) \sim_{i_2} \dots \sim_{i_n} (r_n, m_n) = (r', m')$, where each i_j is in the set $H \cup \sigma(G)$. But this
798 loses the connection to common knowledge of group H and fails to fix the strategy of group G .
799 Instead, what we would need to capture is chains of the form $(r_0, m_0) \sim_{\{i_1\} \cup \sigma(H)} (r_1, m_1) \sim_{\{i_2\} \cup \sigma(H)}$
800 $\dots \sim_{\{i_n\} \cup \sigma(H)} (r_n, m_n) = (r', m')$, where each i_j is in the set G . For this, it appears we need to be
801 able to express the greatest fixpoint X of the equation $X \equiv \bigwedge_{i \in G} D_{\{i\} \cup \sigma(H)}(X \wedge \phi)$. The language
802 $\text{CTLK}(Ags \cup \sigma(Ags), Prop)$ does not include fixpoint operators and it does not seem that this fix-
803 point is expressible. Indeed, the construct $\langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$ does not appear to be expressible using the
804 fragment $\text{CTLK}(Ags \cup \sigma(Ags), Prop)$.

805 However, common knowledge of group H about the effects of a fixed strategy of group G can
806 be expressed with ESL in a natural way by the formula

$$C_H(e_{\sigma(G)}(x) \Rightarrow \phi),$$

807 which says that it is common knowledge to the group H that ϕ holds if the group G is running
808 the strategy profile captured by the variable x . Using this idea, the construct $\langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$ can be
809 represented with ESL as

$$\exists x. C_H(e_{\sigma(G)}(x) \Rightarrow \phi).$$

810 The following result states this claim precisely.

811 **THEOREM 3.3.** *Let E be an environment in which all states are initial, and let Δ be a restrictable and*
812 *extendable set of group strategies in E . Let $\mathcal{I} = \mathcal{I}(E, \text{comp}(\Delta))$. Assume that ϕ is a path formula and*

⁶As above, we have generalized the definition to be relative to a set of group strategies Δ . The strategies used in Reference [38] are imperfect information, perfect recall strategies; we formulate the definition here with imperfect information, imperfect recall strategies.

that ϕ^* is an ESL formula without free variables, such that for every path ρ of E , we have $E, \rho \models^\Delta \phi$ 813
iff for all (equivalently, some) points (r, m) of \mathcal{I} with $r_e[m \dots] = \rho$ we have $\mathcal{I}, (r, m) \models \phi^*$. 814

Then for all states s of E , we have $E, s \models^\Delta \langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$ iff for all (equivalently, some) points (r, m) of 815
 \mathcal{I} with $r_e(m) = s$ we have $\mathcal{I}, (r, m) \models \exists x. C_H(e_{\sigma(G)}(x) \Rightarrow \phi^*)$. 816

PROOF. The argument for the equivalence between the universal and existential versions of the 817
right hand side of the iff is similar to that in Theorem 3.1. 818

Suppose, first, that $E, s \models^\Delta \langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$. Let (r, m) be a point of \mathcal{I} with $r_e(m) = s$. We need to prove 819
that $\mathcal{I}, (r, m) \models \exists x. C_H(e_{\sigma(G)}(x) \Rightarrow \phi^*)$. From $E, s \models^\Delta \langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$ it follows that there exists a strat- 820
egy $\alpha \in \Delta$ for group G , such that for all states t with $s \sim_H^C t$ and paths ρ from t consistent with α , we 821
have $E, \rho \models^\Delta \phi$. Let r' be any run with $r'_{\sigma(G)}(0) = \alpha$, and define Γ to be a context with $\Gamma(x) = r(0)$. 822
To prove $\mathcal{I}, (r, m) \models \exists x. C_H(e_{\sigma(G)}(x) \Rightarrow \phi^*)$, we show that $\Gamma, \mathcal{I}, (r, m) \models C_H(e_{\sigma(G)}(x) \Rightarrow \phi^*)$. For 823
this, suppose that $(r, m) = (r^0, m^0) \sim_{i_1} (r^1, m^1) \sim_{i_2} \dots \sim_{i_k} (r^k, m^k)$, where $i_j \in H$ for $j = 1 \dots k$, 824
and assume that $\Gamma, \mathcal{I}, (r^k, m^k) \models e_{\sigma(G)}(x)$. We need to show that $\Gamma, \mathcal{I}, (r^k, m^k) \models \phi^*$. 825

Note that we have $s = r(m) = r_e^0(m^0) \sim_{i_1} \dots \sim_{i_k} r_e^k(m^k)$. Since $\Gamma, \mathcal{I}, (r^k, m^k) \models e_{\sigma(G)}(x)$, we 826
have that $r_{\sigma(G)}^k(m^k) = \Gamma(x)_{\sigma(G)} = \alpha$. Thus, the sequence $\rho = r_e^k[m^k \dots]$ is a path of E con- 827
sistent with the group strategy α . It follows that $E, \rho \models^\Delta \phi$. By assumption, this means that 828
 $\Gamma, \mathcal{I}, (r^k, m^k) \models \phi^*$. 829

Conversely, let (r, m) be a point of \mathcal{I} with $r_e(m) = s$ and $\mathcal{I}, (r, m) \models \exists x. C_H(e_{\sigma(G)}(x) \Rightarrow \phi^*)$, 830
witnessed by $\Gamma, \mathcal{I}, (r, m) \models C_H(e_{\sigma(G)}(x) \Rightarrow \phi^*)$. Note that $\Gamma(x)_{\sigma(Ags)} = comp(\beta)$, where $\beta \in \Delta$ is 831
a group strategy for some group G' . For agents $i \in G \setminus G'$, we have that $\Gamma(x)_{\sigma(i)}$ is the random 832
strategy $rand_i$. It follows that any path consistent with $\Gamma(x)_{\sigma(G \cap G')}$ is also consistent with $\Gamma(x)_{\sigma(G)}$. 833
Let $\alpha \in \Delta$ be any group strategy for group G with $\alpha \upharpoonright (G \cap G') = \Gamma(x)_{\sigma(G \cap G')}$. Such a strategy 834
exists by the fact that Δ is restrictable and extendable: We may take α to be an extension of $\beta \upharpoonright$ 835
 $(G \cap G')$. Then we have that any path consistent with α is consistent with $\Gamma(x)_{\sigma(G)}$. 836

We show $E, s \models^\Delta \langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$, with α as the witnessing strategy for group G . For this, let $s =$ 837
 $s_0 \sim_{i_1} s_1 \sim_{i_2} \dots \sim_{i_k} s_k$, where $i_j \in H$ for $j = 1 \dots k$, and let ρ be a path from s_k consistent with 838
 α . We show $E, \rho \models^\Delta \phi$. By the observation above, ρ is also consistent with $\Gamma(x)_{\sigma(G)}$. Let r^k be a 839
run with $r_e^k[0 \dots] = \rho$, and $r_{\sigma(G)}^k(0) = \Gamma(x)_{\sigma(G)}$. (We can take $r_{\sigma(Ags)}^k = comp(\beta \upharpoonright (G \cap G'))$, which 840
is in $comp(\Delta)$.) Then $\Gamma, \mathcal{I}, (r, 0) \models e_{\sigma(G)}(x)$. Moreover, for each $j = 1 \dots k - 1$, let r^j be any run 841
with $r_e^j(0) = s_j$. Then $(r, m) = (r^0, m^0) \sim_{i_1} (r^1, 0) \sim_{i_2} \dots \sim_{i_k} (r^k, 0)$. It follows from $\Gamma, \mathcal{I}, (r, m) \models$ 842
 $C_H(e_{\sigma(G)}(x) \Rightarrow \phi^*)$ that $\Gamma, \mathcal{I}, (r^k, 0) \models \phi^*$, and in fact $\mathcal{I}, (r^k, 0) \models \phi^*$, since ϕ^* has no free variables. 843
Since $r_e^k[0 \dots] = \rho$, by assumption, we have $E, \rho \models \phi$. This proves $E, s \models^\Delta \langle\langle G \rangle\rangle_{C(H)}^\bullet \phi$. \square 844

The above equivalences give a reduction of the complex operators of Reference [38] that makes 845
their epistemic content more explicit by expressing this using standard epistemic operators. 846

An alternate approach to decomposing the operators $\langle\langle G \rangle\rangle_{\mathcal{K}(H)}^\bullet$ is proposed in Reference [35]. By 847
comparison with our standard approach to the semantics of the epistemic operators, this proposal 848
uses “constructive knowledge” operators that require a nonstandard semantics in which formulas 849
are evaluated at sets of states rather than at individual states. Evaluation at single world q is treated 850
as equivalent to evaluation at the set $\{q\}$. For each standard (group) epistemic operator $\mathcal{K} = E, D, C$, 851
there is a constructive version $\hat{\mathcal{K}} = \mathbb{E}, \mathbb{D}, \mathbb{C}$. Atomic propositions p are evaluated at sets of states 852
 Q by 853

$$E, Q \models p \text{ if for all states } q \in Q \text{ we have } E, q \models p.$$

854 (As above, we define the semantics on environments E rather than ATEL models.) For the con-
855 structive epistemic operators,

$$E, Q \models \hat{\mathcal{K}}_G \phi \text{ if } E, \{q' \in Q \mid \exists q \in Q (q \sim_G^{\mathcal{K}} q')\} \models \phi$$

856 and for the ATL operator $\langle\langle G \rangle\rangle \phi$ we have

857 $E, Q \models \langle\langle G \rangle\rangle \phi$ if there exists a strategy α for group G such that ϕ holds in all runs
858 starting in a state in Q in which group G plays the strategy α .

859 Note that the ATEL formula $\mathcal{K}_G \langle\langle G \rangle\rangle \phi$ says that at each world considered possible (in the appro-
860 priate sense for \mathcal{K}) by group G , there exists a (possibly different) strategy for G that achieves ϕ .
861 By contrast, $\hat{\mathcal{K}}_G \langle\langle G \rangle\rangle \phi$ says that there exists a *single* strategy for G that achieves ϕ from each world
862 considered possible (in the appropriate sense for \mathcal{K}) by G .

863 This logic is shown in Reference [36] to have a normal form, in which every subformula starting
864 with a constructive knowledge operator $\hat{\mathcal{K}}_G^1$ is of the form $\hat{\mathcal{K}}_{G_1}^1 \dots \hat{\mathcal{K}}_{G_n}^n \phi$, where ϕ starts with a
865 strategy modality and each $\mathcal{K}^i \in \{E, D, C\}$. Such a normal form subformula, evaluated at a single
866 state, can be represented in ESL as

$$\exists x. \mathcal{K}_{G_1}^1 \dots \mathcal{K}_{G_n}^n (e_{\sigma(H)}(x) \Rightarrow \phi).$$

867 Precise formulation and proof of the claim are similar to the proofs above and left to the reader.

868 3.4 Strategy Logic

869 Chatterjee et al's strategy logic [13], which we call CHP-SL, following the convention in Refer-
870 ence [42], is an extension of ATL* for two-player games. Let x, y be two variables ranging over
871 Player 1 and Player 2's strategies. The logic allows these variables to be quantified: If ϕ is a formula,
872 then $\exists x. \phi$ and $\forall x. \phi$ are formulas. Additionally the effects of a particular combination of player
873 strategies can be expressed using the formula $\phi(x, y)$, which says that ϕ holds if player 1 plays
874 strategy x and player 2 plays strategy y . Thus, the ATL* formula $\langle\langle 1 \rangle\rangle \phi$ can be expressed in CHP-SL
875 with $(\exists x)(\forall y)\phi(x, y)$.

876 Strategy logic (SL) [42] generalises CHP-SL, with the syntax as follows:

$$\phi \equiv p \mid \neg \phi \mid \phi \wedge \phi \mid \phi \circ \phi \mid \phi U \phi \mid \langle\langle v \rangle\rangle \phi \mid [[v]] \phi \mid (i, v) \phi,$$

877 where $v \in \text{Var}_{\text{SL}}$ such that Var_{SL} is a set of strategy variables, and $i \in \text{Ags}$ is an agent. Intuitively,
878 $\langle\langle v \rangle\rangle \phi$ says that there exists a strategy v such that ϕ , formula $[[v]] \phi$ says that ϕ holds for all strategies
879 v , and $(i, v) \phi$ says that ϕ holds if agent i plays strategy v . A formula is a *sentence* if every occurrence
880 of (a, x) is within the scope of an occurrence of $\langle\langle x \rangle\rangle$ or $[[x]]$, and every temporal subformula $\phi \circ \phi$ or
881 $\phi U \phi$ occurs within the context of some binding (i, x) , for every agent i . The ATL* formula $\langle\langle 1 \rangle\rangle \phi$
882 can be expressed in SL as $\langle\langle x \rangle\rangle [[y]] (1, x) (2, y) \phi$.

883 Let Str be a set of agent strategies, and $\chi : \text{Ags} \cup \text{Var}_{\text{SL}} \rightarrow \text{Str}$ be a partial mapping from agents
884 and variables to the set of strategies. Then, the semantics can be formulated with respect to our
885 environments E as follows⁷:

- 886 • $E, \chi, (r, m) \models \langle\langle v \rangle\rangle \phi$ iff there exists a strategy $\sigma \in \text{Str}$ such that $E, \chi[v \mapsto \sigma], s \models \phi$;
887 • $E, \chi, (r, m) \models [[v]] \phi$ iff for all strategies $\sigma \in \text{Str}$ it holds that $E, \chi[v \mapsto \sigma], s \models \phi$;
888 • $E, \chi, (r, m) \models (i, v) \phi$ iff $E, \chi[i \mapsto \chi(v)], (r', m) \models \phi$ for all runs r' where $r(m) = r'(m)$ and
889 r' is a run consistent with $\chi[i \mapsto \chi(v)]$ from time m .

⁷We make some simplifications; the authors of Reference [42] distinguish between path and state formulas.

Atomic, Boolean, and temporal formulas are handled as usual. We remark that because (1) the transition relation is assumed in SL to be deterministic, i.e., \rightarrow can be written as a function of type $S \times Acts \rightarrow S$, and (2) temporal operators in a sentence appear only in contexts where every agent is bound to a strategy, the final binding (i, v) before temporal operators are evaluated in fact quantifies over just a single run.

Given an SL formula ϕ , we let $V(\phi)$ be the set of variables in the operators $\langle\langle \cdot \rangle\rangle$ or $[[\cdot]]$. SL allows the assignment of a strategy to multiple agents, e.g., in formula $\langle\langle v \rangle\rangle((i, v)\phi_1 \wedge (j, v)\phi_2)$ the agents i and j have the same strategy represented in the variable v . For this to make sense in an imperfect information system, without allowing implausible bindings or artificially complex interpretations of quantification, all agents need to have the same actions and the same observations. This does not match the setting of our framework particularly well. We remark that CHP-SL does not allow this expressivity, as players 1 and 2 are associated with their dedicated strategy variables x and y , respectively.

In the following, we consider the fragment of SL in which every variable v_i is uniquely associated with an agent $i \in Ags$, so that v_i occurs only in bindings $(j, v_i)\psi$ with $j = i$. Then we can translate a SL formula ϕ into an ESL formula ϕ^* as follows:

$$\begin{aligned}
p^* &= p \\
(\neg\phi)^* &= \neg\phi^* \\
(\phi_1 \wedge \phi_2)^* &= \phi_1^* \wedge \phi_2^* \\
(\circ\phi)^* &= \circ\phi^* \\
(\phi_1 U \phi_2)^* &= \phi_1^* U \phi_2^* \\
(\langle\langle v_i \rangle\rangle\phi)^* &= \exists v_i \phi^* \\
([[v_i]]\phi)^* &= \forall v_i \phi^* \\
((i, v_i)\phi)^* &= D_{e \cup \sigma(Ags \setminus \{i\})}(e_{\sigma(i)}(v_i) \Rightarrow \phi^*).
\end{aligned}$$

Intuitively, to decide if $\langle\langle v_i \rangle\rangle\phi$, we need to determine the existence of a strategy v_i with respect to which the formula ϕ is satisfied. In the ESL translation, v_i refers to a global state rather than a strategy, but the only component of this global state that is used in the remainder of the evaluation is the component $\sigma(i)$, which picks out a strategy for agent i and similarly for $[[v_i]]\phi$. To decide if $(i, v_i)\phi$, we need to satisfy ϕ on (all) runs where agent i 's strategy is switched to that represented in v_i . The translation handles this using the operator $D_{e \cup \sigma(Ags \setminus \{i\})}$, which refers to points in which the state of the environment and the strategies of all agents are fixed, while the strategy of agent i is allowed to vary. The assertion $e_{\sigma(i)}(v_i)$ checks that the strategy of agent i is in fact switched to that represented in the global state v_i .

Similarly, CHP-SL formulas can be translated into ESL formulas as follows:

$$\begin{aligned}
(\exists x.\phi)^* &= \exists x \phi^* \\
(\forall x.\phi)^* &= \forall x \phi^* \\
(\phi(x, y))^* &= D_e(e_{\sigma(1)}(x) \wedge e_{\sigma(2)}(y) \Rightarrow \phi^*).
\end{aligned}$$

Finally, we remark that both the CHP-SL semantics in Reference [13] and the SL semantics in Reference [42] are for perfect recall. Since we have formulated ESL for imperfect recall, we leave the above translations as indicative rather than attempting a formal proof.

919 3.5 Game-theoretic Solution Concepts

920 It has been shown for a number of logics for strategic reasoning that they are expressive enough
 921 to state a variety of game-theoretic solution concepts, e.g., References [13, 28] show that Nash
 922 Equilibrium is expressible. We now sketch the main ideas required to show that the fragment
 923 $\text{CTLK}(Ags \cup \sigma(Ags) \cup \{e\}, Prop)$ of our framework also has this expressive power. We assume two
 924 players $Ags = \{0, 1\}$ in a normal form perfect information game and assume that these agents play
 925 a deterministic strategy. The results in this section can be easily generalized to multiple players
 926 and extensive form games.

927 Given a game \mathcal{G} we construct an environment $E_{\mathcal{G}}$ that represents the game. Each player has
 928 a set of actions that correspond to the moves that the player can make. We assume that $E_{\mathcal{G}}$ is
 929 constructed to model the game so that play happens in the first step from a unique initial state and
 930 that subsequent transitions do not change the state. We let agents have perfect information in $E_{\mathcal{G}}$,
 931 i.e., we define the observation of agent i in state s by $O_i(s) = s$. (Consequently, although we use
 932 uniform strategies $\Sigma^{unif, det}$ below, the uniformity constraint is vacuous in these environments.)

933 We write $-i$ to denote the adversary of player i . Let u_i for $i \in \{0, 1\}$ be a variable denoting the
 934 utility gained by player i when play is finished. Let V_i be the set of possible values for u_i , and let
 935 $V = V_0 \cup V_1$. We work with the following atomic propositions. Atomic proposition $u \leq v$, where
 936 $u, v \in V$, expresses the ordering on utilities. Atomic proposition $u_i = v$, where $i \in \{0, 1\}$ and $v \in V_i$,
 937 expresses that player i 's utility has value v . We use formula

$$U_i(v) = \circ(u_i = v)$$

938 to express that value v is player i 's utility once play finishes.

939 **Nash equilibrium (NE)** is a solution concept that states that no player can gain by unilaterally
 940 changing their strategy. We may write

$$BR_i(v) = U_i(v) \wedge K_{\sigma(-i)} \bigwedge_{v' \in V_i} (U_i(v') \Rightarrow v' \leq v)$$

941 to express that, given the current strategy $\sigma(-i)$ of the adversary of i , the value v attained by player
 942 i 's current strategy is the best possible utility attainable by player i , i.e., the present strategy of
 943 player i is a best response to the adversary. Thus,

$$BR_i = \bigvee_{v \in V_i} BR_i(v)$$

944 says that player i is playing a best-response to the adversary's strategy. The following statement
 945 then expresses the existence of a (pure) Nash equilibrium for the game \mathcal{G} :

$$E_{\mathcal{G}}, \Sigma^{unif, det}(E_{\mathcal{G}}) \models \neg D_{\emptyset} \neg (BR_0 \wedge BR_1).$$

946 That is, in a Nash equilibrium, each player is playing a best response to the other's strategy.

947 **Perfect cooperative equilibrium (PCE)** is a solution concept intended to overcome deficiencies
 948 of Nash equilibrium for explaining cooperative behaviour [26]. It says that each player does at
 949 least as well as she would if the other player were best-responding. The following formula:

$$BU_i(v) = D_{\emptyset} \left(\bigwedge_{v' \in V_i} ((BR_{-i} \wedge U_i(v')) \Rightarrow v' \leq v) \right)$$

950 states that v is as good as any utility that i can obtain if the adversary always best-responds to
 951 whatever i plays. Thus,

$$BU_i = \bigvee_{v \in V_i} (U_i(v) \wedge BU_i(v))$$

says that i is currently getting a utility as good as the best utility that i can obtain if the adversary is a best-responder. Now, the following formula expresses the existence of perfect cooperative equilibrium for the game \mathcal{G} :


$$E_{\mathcal{G}}, \Sigma^{unif, det}(E_{\mathcal{G}}) \models \neg D_0 \neg (BU_0 \wedge BU_1).$$

That is, in a PCE, no player has an incentive to change their strategy, on the assumption that the adversary will best-respond to any change.

3.6 Computer Security Example: Erasure Policies

Formal definitions of computer security frequently involve reference to the strategies available to the players, and to agents' reasoning based on these strategies. In this section, we sketch an example that illustrates how our framework might be applied in this context.

Consider the scenario depicted in the following diagram:



A customer C can purchase items at a web merchant M . Payment is handled by a trusted payment processor P (this could be a service or device), which interacts with the customer, merchant, and a bank B to securely process the payment. (To keep the example simple, we suppose that the customer and merchant use the same bank). One of the guarantees provided by the payment processor is to protect the customer from attacks on the customer's credit card by the merchant: The specification for the protocol that runs the transaction requires that the merchant should not obtain the customer's credit card number. In fact, the specification for the payment processor is that after the transaction has been successfully completed, the payment processor should *erase* the credit card data to ensure that even the payment processor's state does not contain information about the customer's credit card number. The purpose of this constraint is to protect the customer against subsequent attacks by an attacker A , who may be able to use vulnerabilities in the payment processor's software to obtain access to the payment processor's state.

We sketch how one might use our framework to express the specification. To capture reasoning about all possible behaviours of the agents, and what they can deduce from knowledge of those behaviours, we work in $\mathcal{I}^{unif}(E)$ for a suitably defined environment E . To simplify matters, we take $\text{Ags} = \{C, M, P, A\}$. We exclude the strategy of the bank from consideration: This amounts to assuming that the bank has no actions and is trusted to run a fixed protocol. We similarly assume that the payment processor P has no actions, but to talk about what information is encoded in the payment processor's local state, we do allow that this agent has observations. The customer C may have actions such as entering the credit card number in a web form, pressing a button to submit the form to the payment processor, and pressing a button to approve or cancel the transaction. The customer observes variable cc , which records the credit card number drawn from a set CCN , and Boolean variable $done$, which records whether the transaction is complete (which could mean either committed or aborted).

987 We assume that the attacker A has some set of exploit actions, as well as some innocuous actions
 988 (e.g., setting a local variable or performing *skip*). The effect of the exploit actions is to exploit a
 989 vulnerability in the payment processor's software and copy parts of the local state of the payment
 990 processor to variables that are observable by the attacker. We include in the environment state a
 991 Boolean variable `exploited`, which records whether the attacker has executed an exploit action
 992 at some time in the past. The merchant M may have actions such as sending cost information to
 993 the payment processor and acknowledging a receipt certifying that payment has been approved
 994 by the bank (we suppose this receipt is transmitted from the bank to the merchant via the payment
 995 processor).

996 We may then capture the statement that the system is *potentially* vulnerable to an attack that
 997 exploits an erasure flaw in the implementation of the payment processor, by the following formula:

$$\neg D_{\emptyset} \neg \left(\text{done} \wedge \bigvee_{x \in \text{CCN}} K_P(\text{cc} \neq x) \right),$$

998 where $\text{cc} \neq x$ is an atomic proposition for each $x \in \text{CCN}$, with the obvious meaning that the cus-
 999 tomer's credit card number is not x . This says that there exist behaviours of the agents, which can
 1000 (at least on some points in some runs) leave the payment processor in a state where the customer
 1001 has received confirmation that the transaction is done, but in which the payment processor's local
 1002 state somehow still encodes *some* information about the customer's credit card number. This
 1003 encoding could be direct (e.g., by having a variable `customer_cc` that still stores the credit card
 1004 number) or indirect (e.g., by the local state including both a symmetric encryption key K and an
 1005 encrypted version of the credit card number, `enc_customer_cc`, with value $\text{Encrypt}_K(\text{cc})$ that was
 1006 used for secure transmission to the bank). Note that for a breach of security, it is only required
 1007 that the information suffices to *rule out* some credit card number (so that, e.g., knowing the first
 1008 digit of the number would constitute a vulnerability)

1009 The vulnerability captured by this formula is only potential, because it does not necessarily
 1010 follow that the attacker is able to obtain the credit card information. Whether this is possible can
 1011 be checked using the formula

$$\neg D_{\emptyset} \neg \left(\text{done} \wedge \neg \text{exploited} \wedge E \diamond \bigvee_{x \in \text{CCN}} D_{\{A, \sigma(A)\}}(\text{cc} \neq x) \right),$$

1012 which says that it is possible for the attacker to obtain information about the credit card number
 1013 even after the transaction is done. (To focus on erasure flaws, we deliberately wish to exclude here
 1014 the possibility that the attack occurs during the processing of the transaction.) Note that here we
 1015 assume that the attacker knows his own strategy when making deductions from the information
 1016 obtained in the attack. This is necessary, because the attacker can typically write his own local
 1017 variables, so it needs to be able to distinguish between a value it wrote itself and a value it copied
 1018 from the payment processor.

1019 However, even this formula may not be sufficiently strong. Suppose that the payment processor
 1020 implements erasure by writing, to its variable `customer_cc`, a random value. Then, even if the
 1021 attacker obtains a copy of this value, and it happens to be equal to the customer's actual credit
 1022 card number, the attacker would not have any knowledge about the credit card number, since, as
 1023 far as the attacker knows, it could be looking at a randomly assigned number. However, there may
 1024 still be vulnerabilities in the system. Suppose that the implementation of the payment processor
 1025 operates so that the customer's credit card data is not erased by randomization until the merchant
 1026 has acknowledged the receipt of payment from the bank, but to avoid annoying the customer with
 1027 a hanging transaction, the customer is advised that the transaction is approved (setting `done true`)

if the merchant does not respond within a certain time limit. It is still the case that on observing the copied value of *customer_cc*, the attacker would not be able to deduce that this is the customer's credit card number, since it might be the result of erasure in the case that the merchant responded promptly. However, if the attacker knows that the merchant has not acknowledged the receipt, then the attacker can deduce that the value is not due to erasure. One way in which the attacker might know that the merchant has not acknowledged receipt is that the attacker is in collusion with the merchant, who has agreed to omit sending the required acknowledgement messages.

This type of attack can be captured by replacing the term $D_{\{A, \sigma(A)\}}(\text{cc} \neq x)$ by $D_{\{A, \sigma(A), \sigma(M)\}}(\text{cc} \neq x)$, capturing that the attacker reasons using knowledge of both its own strategy as well as the strategy of the merchant or even $D_{\{A, \sigma(A), \sigma(M), M\}}(\text{cc} \neq x)$ for a collusion in which the merchant shares information observed. Similarly, to focus on erasure flaws in the implementation of the payment gateway, independently of the attackers capability, we would replace the term $K_P(\text{cc} \neq x)$ above by $D_{\{P, \sigma(M)\}}(\text{cc} \neq x)$.

We remark that in the case of the attacker's knowledge, it would be appropriate to work with a perfect recall semantics of knowledge, but when using knowledge operators to express information in the payment gateway's state for purposes of reasoning about erasure policy, the more appropriate semantics of knowledge is imperfect recall.

This example illustrates some of the subtleties that arise in the setting of reasoning about security and the way that our framework helps to represent them. Erasure policies have previously been studied in the computer security literature, beginning with Reference [14], though generally without consideration of strategic behaviour by the adversary.

3.7 Reasoning about Knowledge-based Programs

Knowledge-based programs [20] are a form of specification of a multi-agent system in the form of a program structure that describes how an agent's actions are related to its knowledge. They have been shown to be a useful abstraction for several areas of application, including the development of optimal protocols for distributed systems [20], robot motion planning [6], and game-theoretic reasoning [24].

Knowledge-based programs cannot be directly executed, since there is a circularity in their semantics: Which actions are performed depends on what the agents know, which in turn depends on which actions the agents perform. The circularity is not vicious and can be resolved by means of a fixed point semantics, but it means that a knowledge-based program may have multiple distinct implementations (or none), and the problem of reasoning about these implementations is quite subtle. In this section, we show that our framework can capture reasoning about the set of possible implementations of a knowledge-based program.

We consider joint knowledge-based programs P (as defined by Reference [20]) where for each agent i we have a knowledge-based program

$$P_i = \mathbf{do} \phi_1^i \rightarrow a_1^i \square \dots \square \phi_{n_i}^i \rightarrow a_{n_i}^i \mathbf{od},$$

where each ϕ_j^i is a formula of $\text{CTL}^*K(\text{Ags}, \text{Prop})$ of the form $K_i\psi$, and each a_i appears just once. The formulas ϕ_j^i are called the *guards* of the knowledge-based program.⁸ Intuitively, this program says to repeat forever the following operation: Nondeterministically execute one of the actions a_j^i such that the corresponding guard ϕ_j^i is true. To ensure that it is always the case that at least one action

⁸The guards in Reference [20] are allowed to be Boolean combinations of formulas $K_i\psi$ and propositions p local to the agent: Since for such propositions $p \Leftrightarrow K_i p$, and the operator K_i satisfies positive and negative introspection, our form for the guards is equally general. They do not require that a_i appears just once, but the program can always be put into this form by aggregating clauses for a_i into one and taking the disjunction of the guards.

1068 enabled, we assume that $\phi_1^i \vee \dots \vee \phi_{n_i}^i$ is a valid formula; this can always be ensured by taking
 1069 the last condition $\phi_{n_i}^i$ to be the “otherwise” condition $K_i \neg(\phi_1^i \vee \dots \vee \phi_{n_i-1}^i)$, which is equivalent
 1070 to $\neg(\phi_1^i \vee \dots \vee \phi_{n_i-1}^i)$ by introspection. In general, the guards in a knowledge-based program may
 1071 contain common knowledge operators C_G , but we assume for technical reasons (explained below)
 1072 that no ϕ_j^i contains such an operator.

1073 We present a formulation of semantics for knowledge-based programs that refactors the def-
 1074 initions of Reference [20], following the approach of Reference [41], which uses the notion of
 1075 environment defined above rather than the original notion of *context*. A *potential implementation*
 1076 of a knowledge-based program P in an environment E is a joint strategy α in E . (Recall that we
 1077 use “joint strategy” to refer to a group strategy for the group of all agents.) Given a potential im-
 1078 plementation α in E , we can construct the interpreted system $\mathcal{I}_\alpha = \mathcal{I}(E, \{\alpha\})$, which captures the
 1079 possible runs of E when the agents choose their actions according to the single possible joint strat-
 1080 egy α . Given this interpreted system, we can now interpret the epistemic guards in P . Say that a
 1081 state s of E is α -reachable if there is a point (r, m) of \mathcal{I}_α with $r_e(m) = s$. We note that for a formula
 1082 $K_i \phi$, and a point (r, m) of \mathcal{I}_α , the statement $\mathcal{I}_\alpha, (r, m) \models K_i \phi$ depends only on the state $r_e(m)$ of the
 1083 environment at (r, m) . Recall that $r_e(m)$ determines $r_i(m)$ for $i \in \text{Ags}$. For an α -reachable state s of
 1084 E , it therefore makes sense to define satisfaction of $K_i \phi$ at s rather than at a point by $\mathcal{I}_\alpha, s \models K_i \phi$ if
 1085 $\mathcal{I}_\alpha, (r, m) \models K_i \phi$ for all (r, m) with $r_e(m) = s$. We define a joint strategy α to be an *implementation*
 1086 of P in E if for all α -reachable states s of E and agents i , we have

$$\alpha_i(s) = \{a_j^i \mid 1 \leq j \leq n_i, \mathcal{I}_\alpha, s \models \phi_j^i\}.$$

1087 Intuitively, the right-hand side of this equation is the set of actions that are enabled at s by P_i
 1088 when the tests for knowledge are interpreted using the system obtained by running the strategy α
 1089 itself. The condition states that the strategy is an implementation if it enables precisely this set of
 1090 actions at every reachable state. It is easily checked that a strategy α_i satisfying the above equation
 1091 is uniform.

1092 We now show that our framework for strategic reasoning can express the same content as a
 1093 knowledge-based program by means of a formula and that this enables the framework to be used
 1094 for reasoning about knowledge-based program implementations. In general, implementations α of
 1095 a knowledge-based program P can be hard to find, and there may be one, many, or no implementa-
 1096 tions of a given knowledge-based program. We therefore work in strategy space $\mathcal{I}(E, \Sigma^{unif})$, which
 1097 contains all candidate implementations, and develop a formula $\mathbf{imp}(P)$ such that for a given run
 1098 r , the formula $\mathbf{imp}(P)$ holds at a point of r iff the joint strategy encoded in r is an implementation
 1099 of P in E .

1100 We need one constraint on the environment. Say that an environment E is *action recording* if
 1101 for all agents i , for each $a \in \text{Acts}_i$ there exists an atomic proposition $\text{did}_i(a)$ such that for $s \in I$
 1102 we have $\text{did}_i(a) \notin \pi(s)$ and for all states s, t and joint actions a such that $(s, a, t) \in \rightarrow$, we have
 1103 $\text{did}_i(b) \in \pi(t)$ iff $b = a_i$. Intuitively, this means that we can determine from a non-initial state the
 1104 joint action that was executed in reaching that state. It is easily seen that any environment can be
 1105 made action recording just by adding a component to the states that records the latest joint action.

1106 We can now express knowledge-based program implementations as follows. The main issue that
 1107 we need to deal with is that the semantics of knowledge formulas in knowledge-based programs
 1108 is given with respect to a system \mathcal{I}_α , in which it is *common knowledge* that the joint strategy in
 1109 use is α . In general, strategies are not common knowledge in the strategy space $\mathcal{I}(E, \Sigma^{unif})$ within
 1110 which we wish to reason about knowledge-based program implementations. We handle this by
 1111 means of a transformation of formulas.

1112 For a formula ϕ of $\text{CTL}^*K(\text{Ags}, \text{Prop})$, not containing common knowledge operators, write ϕ^s for
 1113 the formula of ESL (in fact, of $\text{CTL}^*K(\text{Ags} \cup \sigma(\text{Ags}), \text{Prop})$) obtained from the following recursively

defined transformation:

1114

$$\begin{aligned}
p^\$ &= p \\
(\neg\phi)^\$ &= \neg\phi^\$ \\
(\phi_1 \wedge \phi_2)^\$ &= \phi_1^\$ \wedge \phi_2^\$ \\
(D_G\phi)^\$ &= D_{G \cup \sigma(Ags)} \phi^\$ \\
(A\phi)^\$ &= A\phi^\$ \\
(\circ\phi)^\$ &= \circ\phi^\$ \\
(\phi_1 U \phi_2)^\$ &= (\phi_1^\$ U \phi_2^\$).
\end{aligned}$$

Intuitively, this substitution says that knowledge operators in ϕ are to be interpreted as if it is known that the current joint strategy is being played. In the case of an operator D_G , which includes the special case $K_i = D_{\{i\}}$, the translation handles this by adding $\sigma(Ags)$ to the set of agents that are kept fixed when moving through the indistinguishability relation.

Let

$$\mathbf{imp}(P) = D_{\sigma(Ags)} \left(\bigwedge_{i \in Ags, j=1 \dots n_i} ((\phi_j^i)^\$ \Leftrightarrow E \circ did_i(a_j^i)) \right).$$

Intuitively, this formula says that the current joint strategy gives an implementation of the knowledge-based program P . More precisely, we have the following:

PROPOSITION 4. *Suppose that P is a knowledge-based program in which the guards do not contain common knowledge operators. Let α be a locally uniform joint strategy in E and let r be a run of $\mathcal{I}(E, \Sigma^{unif}(E))$, in which the agents are running joint strategy α , i.e., $r(0) = (s, \alpha)$ for some state s . Let $m \in \mathbb{N}$. Then*

$$\mathcal{I}(E, \Sigma^{unif}(E)), (r, m) \models \mathbf{imp}(P)$$

iff the strategy α is an implementation of knowledge-based program P in E .

PROOF. For brevity, we write just \mathcal{I} for $\mathcal{I}(E, \Sigma^{unif}(E))$. First, we claim that for a formula ϕ not containing common knowledge operators, we have $\mathcal{I}, (r, m) \models \phi^\$$ iff $\mathcal{I}_\alpha, (r, m) \models \phi$, where $r(m) = (s, \alpha)$. The proof is by induction on the construction of ϕ . The base case of atomic propositions, and the cases for Boolean and linear temporal operators are straightforward.

Consider the case $\phi = A\psi$, where we have $(A\psi)^\$ = A(\psi^\$)$. Observe that if r and r' are runs of \mathcal{I} , with $r[0 \dots m] = r'[0 \dots m]$, then r and r' encode the same strategy $\alpha = r_{\sigma(Ags)}(0)$. Now $\mathcal{I}, (r, m) \models A(\psi^\$)$ iff $\mathcal{I}, (r', m) \models \psi^\$$ for all runs r' of \mathcal{I} with $r[0 \dots m] = r'[0 \dots m]$. By the observation, this is equivalent to $\mathcal{I}, (r', m) \models \psi^\$$ for all runs r' of \mathcal{I}_α with $r[0 \dots m] = r'[0 \dots m]$. By induction, the latter is equivalent to $\mathcal{I}_\alpha, (r', m) \models \psi$ for all runs r' of \mathcal{I}_α with $r[0 \dots m] = r'[0 \dots m]$, i.e., to $\mathcal{I}_\alpha, (r, m) \models A\psi$. Hence $\mathcal{I}, (r, m) \models (A\psi)^\$$ iff $\mathcal{I}_\alpha, (r, m) \models A\psi$.

Finally, consider the case $\phi = D_G\psi$, where we have $(D_G\psi)^\$ = D_{G \cup \sigma(Ags)}(\psi^\$)$. Observe that if (r, m) and (r', m') are points of \mathcal{I} with $(r, m) \sim_{G \cup \sigma(Ags)} (r', m')$, then r and r' encode the same strategy $\alpha = r_{\sigma(Ags)}(0)$ and $(r, m) \sim_G (r', m')$. Conversely, if (r, m) and (r', m') are points of \mathcal{I}_α , i.e., both encode joint strategy α , then $(r, m) \sim_G (r', m')$ implies $(r, m) \sim_{G \cup \sigma(Ags)} (r', m')$. Now $\mathcal{I}, (r, m) \models D_{G \cup \sigma(Ags)}(\psi^\$)$ iff $\mathcal{I}, (r', m') \models \psi^\$$ for all points (r', m') of \mathcal{I} with $(r, m) \sim_{G \cup \sigma(Ags)} (r', m')$. By the observation, this is equivalent to $\mathcal{I}, (r', m') \models \psi^\$$ for all points (r', m') of \mathcal{I}_α with $(r, m) \sim_G (r', m')$. By induction, this is equivalent to $\mathcal{I}_\alpha, (r', m') \models \psi$ for all points (r', m') of \mathcal{I}_α with $(r, m) \sim_G (r', m')$, i.e., to $\mathcal{I}_\alpha, (r, m) \models D_G\psi$.

This completes the proof of the claim. Next, note that, for a point (r, m) with $r(m) = (s, \alpha)$, for action α_j^i of agent i , we have $\mathcal{I}, (r, m) \models E \circ did_i(a_j^i)$ iff $a_j^i \in \alpha_i(s)$.

1147 Suppose that α is an implementation of P in E , and let (r, m) be a point of \mathcal{I} with $r(m) = (s, \alpha)$,
 1148 We show that $\mathcal{I}, (r, m) \models \mathbf{imp}(P)$. For this, we let (r', m') be a point with $(r', m') \sim_{\sigma(Ags)} (r, m)$
 1149 and show that $\mathcal{I}, (r', m') \models \bigwedge_{i \in Ags, j=1 \dots n_i} ((\phi_j^i)^{\mathfrak{S}} \Leftrightarrow E \circ did_i(a_j^i))$. From $(r', m') \sim_{\sigma(Ags)} (r, m)$, it fol-
 1150 lows that $r'(m') = (t, \alpha)$ for some state t of E . Thus, from what was noted above, $\mathcal{I}, (r', m') \models$
 1151 $E \circ did_i(a_j^i)$ iff $a_j^i \in \alpha_i(t)$. Since α is an implementation of P in E , this holds iff $\mathcal{I}_\alpha, (r', m') \models \phi_j^i$. By
 1152 the claim proved above, $\mathcal{I}_\alpha, (r', m') \models \phi_j^i$ is equivalent to $\mathcal{I}, (r', m') \models (\phi_j^i)^{\mathfrak{S}}$. Thus, we have that
 1153 $\mathcal{I}, (r', m') \models (\phi_j^i)^{\mathfrak{S}} \Leftrightarrow E \circ did_i(a_j^i)$. It follows that $\mathcal{I}, (r, m) \models \mathbf{imp}(P)$.

1154 Conversely, suppose that $\mathcal{I}, (r, m) \models \mathbf{imp}(P)$, and let $r(m) = (s, \alpha)$. We show that α is an imple-
 1155 mentation of P in E . Let t be any α -reachable state, with, in particular, (r', m') a point of \mathcal{I}_α with
 1156 $r'(m') = (t, \alpha)$. We need to show that for all agents i , we have

$$\alpha_i(t) = \{a_j^i \mid 1 \leq j \leq n_i, \mathcal{I}_\alpha, t \models \phi_j^i\},$$

1157 i.e., that for all i, j we have $a_j^i \in \alpha_i(t)$ iff $\mathcal{I}_\alpha, t \models \phi_j^i$. Note that $(r, m) \sim_{\sigma(Ags)} (r', m')$, so we have
 1158 that

$$\mathcal{I}, (r', m') \models \bigwedge_{i \in Ags, j=1 \dots n_i} ((\phi_j^i)^{\mathfrak{S}} \Leftrightarrow E \circ did_i(a_j^i)).$$

1159 As in the previous paragraph, $a_j^i \in \alpha_i(t)$ iff $\mathcal{I}, (r', m') \models E \circ did_i(a_j^i)$, which is equivalent to
 1160 $\mathcal{I}, (r', m') \models (\phi_j^i)^{\mathfrak{S}}$, and by the claim proved above, equivalent to $\mathcal{I}_\alpha, (r', m') \models \phi_j^i$, i.e., $\mathcal{I}_\alpha, t \models \phi_j^i$.
 1161 Thus, $a_j^i \in \alpha_i(t)$ iff $\mathcal{I}_\alpha, t \models \phi_j^i$, for all i, j , which is what we needed to prove. \square

1162 In particular, as a consequence of this result, it follows that several properties of knowledge-
 1163 based programs (that do not make use of common knowledge operators) can be expressed in the
 1164 system $\mathcal{I}(E, \Sigma^{unif}(E))$:

1165 (1) The statement that there exists an implementation of P in E can be expressed by

$$\mathcal{I}(E, \Sigma^{unif}(E)) \models \neg D_0 \neg \mathbf{imp}(P).$$

1166 (2) The statement that all implementations of P in E guarantee that formula ϕ of
 1167 CTL*K($Ags, Prop$) (which may contain knowledge operators) holds at all times can be
 1168 expressed by

$$\mathcal{I}(E, \Sigma^{unif}(E)) \models D_0(\mathbf{imp}(P) \Rightarrow \phi^{\mathfrak{S}}).$$

1169 We remark that as a consequence of these encodings and Theorem 4.5 (in Section 4 below)
 1170 that CTL*K($Ags \cup \sigma(Ags), Prop$) model checking in strategy space is in PSPACE, we obtain the
 1171 following result:

1172 COROLLARY 1. *The following are in PSPACE:*

- 1173 (1) *Given a finite environment E and a knowledge-based program P , determine if P has an im-*
 1174 *plementation in E .*
 1175 (2) *Given a finite environment E and a knowledge-based program P and a CTL*K($Ags, Prop$)*
 1176 *formula ϕ , determine if $\mathcal{I}_\alpha \models \phi$ for all implementations α of P in E .*

1177 For testing existence (part 1 of Corollary 1), this result was known [20], but the result on verifi-
 1178 cation (part 2 of Corollary 1) has not previously been noted (though it could also have been shown
 1179 using the techniques in Reference [20].)

1180 One might expect that Proposition 4 can be extended to knowledge-based programs in which
 1181 formulas may contain common knowledge operators, simply by adding the condition

$$(C_G \phi)^{\mathfrak{S}} = C_{G \cup \sigma(Ags)} \phi^{\mathfrak{S}}$$

to the transformation of formulas. However, this does not work, because the interpretation of $C_G\phi$ 1182
 in a subsystem \mathcal{I}_α is based on chains of points $(r_0, m_0) \sim_{i_1} (r_1, m_1) \sim_{i_2} \dots \sim_{i_k} (r_k, m_k)$, such that r_j 1183
 is a run of joint strategy α for all $j = 1 \dots k$. By contrast, the semantics of $C_{G \cup \sigma(Ags)}\phi^s$ in \mathcal{I} involves 1184
 chains of points that are not required to preserve the joint strategy: Rather, each step preserves 1185
 the local state of one of the agents in G or the strategy of one of the agents. Neither does it work 1186
 to use the translation 1187

$$(C_G\phi)^s = \exists x(e_{\sigma(Ags)}(x) \wedge C_G(e_{\sigma(Ags)}(x) \Rightarrow \phi^s)),$$

since the operator C_G similarly does not preserve the joint strategy, and it is not enough to test 1188
 only at the end of the chain that the joint strategy has been preserved. 1189

It is not clear that the translation we require for common knowledge is expressible in ESL. What 1190
 would work is to generalize the common knowledge operator to the form $C_X\phi$, where X is a set 1191
 of sets of agents (instead of a set of agents) and to define the semantics of this more general form 1192
 as the greatest fixpoint of equation 1193

$$C_X\phi = \bigwedge_{G \in X} D_G(\phi \wedge C_X\phi).$$

We could then use the translation 1194

$$(C_G\phi)^s = C_{\{\{i\} \cup \sigma(Ags) \mid i \in G\}}\phi^s.$$

Here the semantics involves chains of points in which we preserve the joint strategy and one of 1195
 the agents in G . While this is an interesting extension, that we consider worthy of study, we do not 1196
 pursue this as an ad hoc extension here, leaving it for future consideration in a broader context, 1197
 such as a logic that extends ESL by mu-calculus operators. 1198

4 MODEL CHECKING 1199

Model checking is the problem of computing whether a formula of a logic holds in a given model. 1200
 We now consider the problem of model-checking ESL and various of its fragments. 1201

The model-checking problem is to determine whether $\Gamma, E, \Sigma \models \phi$ for a finite-state environment 1202
 E , a set Σ of strategies and a context Γ , where ϕ is an ESL formula. 1203

For purposes of results concerning the complexity of model checking, we need a measure of 1204
 the size of a finite environment. Conventionally, the size of a model is taken to be the length of 1205
 a string that lists its components, and, typically, this is polynomial in the number of states of the 1206
 model. We note that in the case of environments, the set of labels *Acts* of the transition relation is 1207
 an n -fold Cartesian product, where $n = \text{Ags}$, so (if the number of agents is a variable in the class 1208
 of environments we consider) the size of the transition relation may be exponential in the number 1209
 of agents.⁹ 1210

However, there is a more severe issue with respect to the parameter Σ of the model-checking 1211
 problem. A strategy for a single agent is a mapping from states to sets of actions of the agent. 1212
 Hence the number of strategies we may need to list to describe Σ explicitly could be exponential 1213

⁹For certain classes of environments, we could address this by allowing that the transition relation \rightarrow is presented in some notation with the property that (1) given states s, t and a joint action a , the representation of \rightarrow has size polynomial in the size of $|S|$ and $|Acts|$, and (2) determining whether $s \xrightarrow{a} t$ is in PTIME given s, a, t and the representation of \rightarrow . One example of a presentation format with this property is the class of *turn-based environments*, where at each state s , there exists an agent i such that if $s \xrightarrow{a} t$ for a joint action a , then for all joint actions b with $a_i = b_i$ we have $s \xrightarrow{b} t$. That is, the set of states reachable in a single transition from s depends only on the action performed by agent i . In this case, the transition relation can be presented more succinctly as a subset of $S \times (\cup_{i \in \text{Ags}} A_i) \times S$. While it would be interesting to consider the effect of such optimized representations on our complexity results, we do not pursue this here.

1214 in the number of *states* of the environment, even in the case of a single agent. To address this issue,
 1215 we abstract the strategy set Σ to a parameterized class such that for each environment E , the set
 1216 $\Sigma(E)$ is a set of strategies for E . When C is a complexity class, we say that the parameterized class
 1217 Σ can be presented in C , if the problem of determining, given an environment E and a joint strategy
 1218 α for E , whether $\alpha \in \Sigma(E)$, is in complexity class C . For example, the class Σ of all strategies for E
 1219 can be PTIME-presented, as can Σ^{unif} , Σ^{det} , and $\Sigma^{unif, det}$.

1220 We first consider the complexity of model checking the full language ESL. The following result
 1221 gives an upper bound of EXPSPACE for this problem.

1222 **THEOREM 4.1.** *Let Σ be a parameterized class of strategies that can be presented in EXPSPACE. The*
 1223 *complexity of deciding, given an environment E , an ESL formula ϕ and a context Γ for $\mathcal{I}(E, \Sigma(E))$,*
 1224 *defined on the free variables of ϕ , whether $\Gamma, E, \Sigma(E) \models \phi$, is in EXPSPACE.*

1225 **PROOF.** The problem can be reduced to that of model checking the temporal epistemic logic
 1226 CTL*K obtained by omitting the constructs \exists and $e_i(x)$ from the language ESL. This is known to
 1227 be PSPACE-complete.¹⁰ The reduction involves an exponential blowup of size of both the formula
 1228 and the environment, so we obtain an EXPSPACE upper bound.

1229 Model checking for temporal epistemic logic takes as input a formula and a structure that is like
 1230 an environment, except that its transitions are not based on a set of actions for the agents. More
 1231 precisely, an *epistemic transition system* for a set of agents Ags is a tuple $\mathcal{E} = \langle S, I, \rightarrow, \{O_i\}_{i \in Ags}, \pi \rangle$,
 1232 where S is a set of states, $I \subseteq S$ is the set of initial states, $\rightarrow \subseteq S \times S$ is a state transition relation,
 1233 for each $i \in Ags$, component $O_i : S \rightarrow L_i$ is a function giving an observation in some set L_i for the
 1234 agent i at each state, and $\pi : S \rightarrow \mathcal{P}(Prop)$ is a propositional assignment. A *run* of \mathcal{E} is a sequence
 1235 $r : \mathbb{N} \rightarrow S$ such that $r(0) \in I$ and $r(k) \rightarrow r(k+1)$ for all $k \in \mathbb{N}$. To ensure that every partial run can
 1236 be completed to a run, we assume that the transition relation is *serial*, i.e., that for all states s there
 1237 exists a state t such that $s \rightarrow t$.

1238 Given an epistemic transition system \mathcal{E} , we define an interpreted system $\mathcal{I}(\mathcal{E}) = (\mathcal{R}, \pi')$ as
 1239 follows. For a run $r : \mathbb{N} \rightarrow S$ of \mathcal{E} , define the lifted run $\hat{r} : \mathbb{N} \rightarrow S \times \prod_{i \in Ags} L_i$ (here $L_e = S$), by
 1240 $\hat{r}_e(m) = r(m)$ and $\hat{r}_i(m) = O_i(r(m))$ for $i \in Ags$. Then we take \mathcal{R} to be the set of lifted runs \hat{r} with
 1241 r a run of \mathcal{E} . The assignment π' is given by $\pi'(r, m) = \pi(r(m))$. The model-checking problem for
 1242 temporal epistemic logic CTL*K is to decide, given an epistemic transition system \mathcal{E} and a formula
 1243 $\phi \in \text{CTL}^*K$, whether $\mathcal{I}(\mathcal{E}), (r, 0) \models \phi$ for all runs r of $\mathcal{I}(\mathcal{E})$.

1244 We now show how to reduce ESL model checking to CTL*K model checking. Given an environ-
 1245 ment $E = \langle S, I, Acts, \rightarrow, \{O_i\}_{i \in Ags}, \pi \rangle$ for ESL($Ags, Prop, Var$), we first introduce a set of new propo-
 1246 sitions $Prop^* = \{p_{(s, \alpha)} \mid s \in S, \alpha \in \Sigma(E)\}$, which will be interpreted at global states of the generated
 1247 interpreted system. Each proposition $p_{(s, \alpha)}$ will be true only at the global state (s, α) . These propo-
 1248 sitions will help to eliminate the constructs $e_i(x)$ and $\exists x$. We then define the epistemic transition
 1249 system $\mathcal{E} = \langle S^*, I^*, \rightarrow^*, \{O_i^*\}_{i \in Ags}, \pi^* \rangle$ for the language $\text{CTL}^*K(Ags \cup \sigma(Ags), Prop \cup Prop^*, Var)$,
 1250 in which the propositions have been extended by the set $Prop^*$, as follows:

- 1251 (1) $S^* = \{(s, \alpha) \in S \times \Sigma(E) \mid s \text{ is reachable in } E \text{ using } \alpha\}$,
- 1252 (2) $I^* = I \times \Sigma(E)$,
- 1253 (3) $(s, \alpha) \rightarrow^* (t, \beta)$ iff $s \xrightarrow{a} t$ (in E) for some joint action a and $\beta = \alpha$,
- 1254 (4) $O_i^*(s, \alpha) = O_i(s)$ and $O_{\sigma(i)}^*(s, \alpha) = \alpha_i$, for $i \in Ags$,
- 1255 (5) $\pi^*(s, \alpha) = \pi(s) \cup \{p_{(s, \alpha)}\}$.

¹⁰The result is stated explicitly in Reference [18], but techniques sufficient for a proof (involving guessing a labelling of states by knowledge subformulas in order to reduce the problem to LTL model checking and also verifying the guess by LTL model checking) were already present in Reference [56]. The branching operator A can be treated as a knowledge operator for purposes of the proof.

We can treat the states $(s, \alpha) \in S^*$ as tuples indexed by $\text{Ags} \cup \sigma(\text{Ags}) \cup \{e\}$ by taking $(s, \alpha)_i =$ 1256
 $O_i(s)$ and $(s, \alpha)_{\sigma(i)} = \alpha_i$ for $i \in \text{Ags}$, and $(s, \alpha)_e = s$. 1257

Note that a joint strategy for an environment E can be represented in space $\Sigma_{i \in \text{Ags}} |S| \times |\text{Acts}_i|$, 1258
 and the number of strategies is exponential in the space requirement. Thus, the size of \mathcal{E} is 1259
 $O(2^{\text{poly}(|E|)})$. Note also that the construction of \mathcal{E} can be done in EXPSPACE so long as verify- 1260
 ing whether an individual strategy α is in $\Sigma(E)$ can be done in EXPSPACE. 1261

We also need a transformation of the formula. Given a formula ϕ of ESL and a context Γ for E , 1262
 we define a formula ϕ^Γ , inductively, by 1263

- (1) $p^\Gamma = p$, for $p \in \text{Prop}$, 1264
- (2) $e_i(x)^\Gamma = \bigvee \{p_g \mid g \in S^*, g_i = \Gamma(x)_i\}$ 1265
- (3) $(\neg\phi)^\Gamma = \neg\phi^\Gamma$, $(\phi_1 \wedge \phi_2)^\Gamma = \phi_1^\Gamma \wedge \phi_2^\Gamma$, 1266
- (4) $(\circ\phi)^\Gamma = \circ(\phi^\Gamma)$, $(\phi_1 U \phi_2)^\Gamma = (\phi_1^\Gamma) U (\phi_2^\Gamma)$, $(A\phi)^\Gamma = A(\phi^\Gamma)$ 1267
- (5) $(D_G\phi)^\Gamma = D_G\phi^\Gamma$, $(C_G\phi)^\Gamma = C_G\phi^\Gamma$, 1268
- (6) $\exists x(\phi)^\Gamma = \bigvee \{\phi^{\Gamma[g/x]} \mid g \in S^*\}$. 1269

Plainly, the size of ϕ^Γ is $O(2^{\text{poly}(|E|, |\phi|)})$, and this formula is in $\text{CTL}^*K(\text{Ags} \cup \sigma(\text{Ags}), \text{Prop} \cup$ 1270
 $\text{Prop}^*)$. A straightforward inductive argument based on the semantics shows that 1271

$\Gamma, E, \Sigma(E) \models \phi$ iff $\mathcal{I}(\mathcal{E}) \models \phi^\Gamma$. It therefore follows from the fact that model checking CTL^*K with 1272
 respect to the observational semantics for knowledge is in PSPACE that ESL model checking is in 1273
 EXPSPACE. \square 1274

The following result shows that a restricted version of the model-checking problem, where we 1275
 consider systems with just one agent and uniform deterministic strategies is already EXPSPACE 1276
 hard. 1277

THEOREM 4.2. *The problem of deciding, given an environment E for a single agent, and an ESL* 1278
sentence ϕ , whether $E, \Sigma^{\text{unif}, \text{det}}(E) \models \phi$, is EXPSPACE-hard. 1279

PROOF. We show how polynomial size inputs to the problem can simulate exponential space 1280
 deterministic Turing machine computations. Let $T = \langle Q, q_0, q_f, q_r, A_I, A_T, \delta \rangle$ be a one-tape Turing 1281
 machine solving an EXPSPACE-complete problem, with states Q , initial state q_0 , final (accepting) 1282
 state q_f , final (rejecting) state q_r , input alphabet A_I , tape alphabet $A_T \supseteq A_I$, and transition function 1283
 $\delta : Q \times A_T \rightarrow Q \times A_T \times \{L, R\}$. We assume that T runs in space $2^{p(n)} - 2$ for a polynomial $p(n)$, and 1284
 that the transition relation is defined so that the machine idles in its final state q_f on accepting 1285
 and idles in state q_r on rejecting. The tape alphabet A_T is assumed to contain the blank symbol \perp . 1286

Define $C_{T, Q} = A_T \cup (A_T \times Q)$ to be the set of ‘‘cell-symbols’’ of T . We may represent a configura- 1287
 tion of T as a finite sequence over the set $C_{T, Q}$, containing exactly one element (x, q) of $A_T \times Q$, 1288
 representing a cell containing symbol x where the machine’s head is positioned, with the machine 1289
 in state q . For technical reasons, we pad configurations with a blank symbol to the left and right (so 1290
 configurations take space $2^{p(n)}$), so that the initial configuration has the head at the second tape 1291
 cell and, without loss of generality, assume that the machine is designed so that it never moves 1292
 the head to the initial or final padding blank. This means that the transition function δ can also be 1293
 represented as a set of tuples $\Delta \subseteq C_{T, Q}^4$, such that $(a, b, c, d) \in \Delta$ iff, whenever the machine is in 1294
 a configuration with a, b, c at cells at positions $k - 1, k, k + 1$, respectively, the next configuration 1295
 has d at the cell at position k . 1296

Given the TM T and a number $N = p(n)$ (for some polynomial p) we construct an environment 1297
 $E_{T, N}$ such that for every input word w , with $|w| = n$, there exists a sentence ϕ_w of size polyno- 1298
 mial in n such that $E_{T, N}, \Sigma^{\text{unif}}(E_{T, N}) \models \phi_w$ iff T accepts w . The idea of the simulation, depicted 1299
 in Figure 1, is to represent a computation of the Turing machine, using space 2^N , by representing 1300

We take the set of states of the environment to be

1322

$$S = \{s_x \mid x \in C_{T,Q}\} \cup \{c_0, c_1\} \cup \{(t_i, j) \mid i = 0 \dots N-1, j \in \{0, 1\}\}.$$

The set of initial states of the environment is defined to be $I = \{s_\perp\} \cup \{(t_i, j) \mid i = 0 \dots N-1, j \in \{0, 1\}\}$. We define the assignment π so that $\pi(s_a) = \{a\}$ for $a \in C_{T,Q}$, $\pi(c_0) = \emptyset$, $\pi(c_1) = \{c\}$ and $\pi((t_i, 0)) = \{t_i\}$ and $\pi((t_i, 1)) = \{t_i\} \cup \{c\}$.

1323

1324

1325

We take the set of actions of the single agent to be the set $\{a_0, a_1\}$. The transition relation \rightarrow is defined so that for the only transitions are

1326

1327

$$\begin{aligned} s_x &\xrightarrow{a_k} c_i \\ c_i &\xrightarrow{a_k} c_j \\ c_i &\xrightarrow{a_k} s_x \\ (t_m, j) &\xrightarrow{a_k} (t_m, k) \end{aligned}$$

for $x \in C_{T,Q}$ and $i, j, k \in \{0, 1\}$ and $m \in \{0 \dots N-1\}$. Intuitively, this forces the runs starting at state s_\perp to alternate between selecting a symbol from $C_{T,Q}$ and a sequence of bits $\{0, 1\}$ for the counter. Note that for every sequence ρ in $\perp \cdot \{c_0, c_1\}^+ \cdot (C_{T,Q} \cdot \{c_0, c_1\}^+)^{\omega}$, and for every strategy α for the single agent, there exists a run r with $r_{\sigma(1)} = \alpha$ and $r_e[0 \dots] = \rho$. For each $i = 0, \dots, N-1$, the states of the form (t_i, j) for $j \in \{0, 1\}$ form an isolated component in the transition relation and are used to ensure that there is a sufficiently rich set of strategy choices for strategies to encode counter values.

1328

1329

1330

1331

1332

1333

1334

The length of the counter sequence segments of a run generated by this transition system can vary within the run, but we can use a formula of length $O(N)$ to state that these segments always have length N wherever they appear in the run; let ϕ_{clock}^N be the formula

1335

1336

1337

$$\Box \left(\alpha_{T,Q} \Rightarrow \left(\bigcirc^{N+1} \alpha_{T,Q} \wedge \bigwedge_{i=1 \dots N} \bigcirc^i \neg \alpha_{T,Q} \right) \right),$$

where we write $\alpha_{T,Q}$ for $\bigvee_{x \in C_{T,Q}} x$. By definition of the transition relation, this formula holds on a run starting in state s_\perp just when it consists of states of the form s_x alternating with sequences of states of the form c_i of length exactly N .

1338

1339

1340

The transition system generates arbitrary such sequences of states c_i of length N , intuitively constituting a guess for the correct counter value. Note that a temporal formula of length $O(N^2)$ can say that these guesses for the counter values are correct, in that the counter values encoded along the run are $0, 1, 2, \dots, 2^N - 1, 0, 1, 2, \dots, 2^N - 1$ (etc.). Specifically, this is achieved by the following formula ϕ_{count}^N :

1341

1342

1343

1344

1345

$$\phi_{zero} \wedge \Box \left(\alpha_{T,Q} \Rightarrow \left(\begin{aligned} &(\phi_{max} \Rightarrow \bigcirc^{N+1}(\phi_{zero})) \wedge \\ &\bigwedge_{i=1 \dots N} ((\bigcirc c \wedge \dots \wedge \bigcirc^{i-1} c \wedge \bigcirc^i \neg c) \Rightarrow \\ &\bigcirc^{N+1}(\bigcirc \neg c \wedge \dots \wedge \bigcirc^{i-1} \neg c \wedge \bigcirc^i c) \\ &\wedge \bigwedge_{j=i+1 \dots N} ((\bigcirc^j c) \Leftrightarrow (\bigcirc^{j+N+1} c)) \end{aligned} \right) \right),$$

where $\phi_{zero} = \bigwedge_{i=1 \dots N} \bigcirc^i \neg c$ and $\phi_{max} = \bigwedge_{i=1 \dots N} \bigcirc^i c$. Intuitively, the first line of the inner formula handles the steps from $2^N - 1$ to 0, and the remainder of the inner formula uses the fact that, in binary, $x01^i + 1 = x10^i$. (Recall that in the run, low order bits are represented to the left.)

1346

1347

1348

The following formula ϕ_{init}^w then says that the run is initialized with word $w = a_1 \dots a_n$

1349

$$\perp \wedge \bigcirc^{N+1}((q_0, a_1) \wedge \bigcirc^{N+1}(a_2 \wedge \bigcirc^{N+1}(\dots \bigcirc^{N+1}(a_n \wedge \bigcirc((\alpha_{T,Q} \Rightarrow (\perp \wedge \neg \phi_{zero}))U(\alpha_{T,Q} \wedge \phi_{zero})))) \dots)),$$

1350 where \perp is the blank symbol. This formula has size $O(N \cdot |w|) = O(p(n) \cdot n)$. Intuitively, the for-
 1351 mula says that the sequence of symbols w is followed by a sequence of \perp symbols until the first
 1352 time that the counter has value zero (this corresponds to the start of the second configuration).

1353 We now need a formula that expresses that whenever we consider two consecutive configura-
 1354 tions C, C' encoded in a run, C' is derived from C by a single step of the TM T . The padding blanks
 1355 are easily handled by the following formula ϕ_{pad} :

$$\Box((\alpha_{T,Q} \wedge (\phi_{zero} \vee \phi_{max})) \Rightarrow \perp).$$

1356 For the remaining cell positions, we need to express that for each cell position $k = 1 \dots 2^N - 2$,
 1357 the cell value at position k in C' is determined from the cell value at positions $k - 1, k, k + 1$ in C
 1358 according to the transition relation encoding Δ . This means that we need to be able to identify the
 1359 corresponding positions k in C and C' . To capture the counter value at a given position in the run,
 1360 we represent counter values using a strategy for the single agent, as follows.

1361 We define the observation function O_1 for the single agent in $E_{T,N}$, so that observation
 1362 $O_1((t_i, j)) = i$ for $i = 0 \dots N - 1$. (The values of the observation function on other states are
 1363 not used in the encoding, and can be defined arbitrarily.) The number with binary represen-
 1364 tation $B = b_{N-1} \dots b_0$ can then be represented by the strategy α_B such that $\alpha_B(t_i, j) = a_{b_i}$, for
 1365 $i = 0 \dots N - 1$ and $j \in \{0, 1\}$ and $\alpha_B(s) = a_0$ for all other states s . (Note that this strategy is uni-
 1366 form, and, conversely, for any uniform strategy α there exists a unique binary number $b_{N-1} \dots b_0$
 1367 such that $\alpha_B(t_i, j) = a_{b_i}$, for $i = 0 \dots N - 1$ and $j \in \{0, 1\}$.) Comparing this representation with the
 1368 encoding of numbers along runs, the following formula $\phi_{num}(x)$ expresses that the number en-
 1369 coded at the present position in the run is the same as the number encoded in the strategy of agent
 1370 1 in the global state denoted by variable x :

$$\alpha_{T,Q} \wedge \bigwedge_{i=0 \dots N-1} (\circ^{i+1}c) \Leftrightarrow \neg D_\emptyset \neg (e_{\sigma(1)}(x) \wedge t_i \wedge \circ c).$$

1371 Note that, by the definition of the transition system, the value of $\circ c$ at a state where t_i holds
 1372 encodes whether the strategy selects a_0 or a_1 on observation $i = 0 \dots N - 1$. Note also that since
 1373 all states of the form (t_i, j) are initial, for every strategy α , the value of $\alpha(t_i, j)$ is represented in this
 1374 way at some point of some run. We may now check that the transitions of the TM are correctly
 1375 computed along the run by means of the following formula ϕ_{trans} :

$$\Box \left(\bigwedge_{(a,b,c,d) \in \Delta} (a \wedge \neg \phi_{max} \wedge \circ^{N+1}(b \wedge \neg \phi_{max} \wedge \circ^{N+1}c) \Rightarrow \right. \\ \left. \circ^{N+1} \exists x [\phi_{num}(x) \wedge \circ((\neg \phi_{num}(x))U(\phi_{num}(x) \wedge d))] \right).$$

1376 Intuitively, here x captures the number encoded at the cell containing the symbol b , and the U
 1377 operator is used to find the next occurrence in the run of this number. The occurrences of ϕ_{max}
 1378 ensure that the three positions considered in the formula do not span across a boundary between
 1379 two configurations. In Figure 1, the bottom part represents a strategy α of agent 1 encoded in some
 1380 global state x . The behaviour of this strategy at the t -state runs represents a number, using the
 1381 statement $e_{\sigma(1)}(x)$ in the formula ϕ_{num} . The formula ϕ_{num} is used to assert that this representation
 1382 of a binary number in α encodes the counter values at a position in a run. Asserting that two
 1383 positions have the same counter number by this device allows us to check the yields relation at
 1384 corresponding positions in the run representation of a computation of the Turing machine.

1385 To express that the machine accepts, we just need to assert that the accepting state is reached;
 1386 this is done by the formula $\phi_{accept} = \diamond \bigvee_{a \in A_T} (a, q_f)$.

Combining these pieces, we get that the TM accepts input w if and only if 1387

$$E_{T,N} \models (\phi_{clock}^N \wedge \phi_{count}^N \wedge \phi_{init}^w \wedge \phi_{trans}) \Rightarrow \phi_{accept}$$

holds, i.e., when every run that correctly encodes a computation of the machine is accepting. \square 1388

Combining Theorem 4.1 and Theorem 4.2 we obtain the following characterization of the com- 1389
plexity of ESL model checking. 1390

COROLLARY 2. *Let Σ be an EXPSPACE presented class of strategies for environments, containing 1391
 $\Sigma^{unif,det}$. The complexity of deciding, given an environment E , an ESL formula ϕ and a context Γ for 1392
the free variables in an ESL formula ϕ relative to E and $\Sigma(E)$, whether $\Gamma, E, \Sigma(E) \models \phi$, is EXPSPACE- 1393
complete. 1394*

The high complexity for ESL model checking motivates the consideration of fragments that 1395
have lower model-checking complexity. We demonstrate two orthogonal fragments for which the 1396
complexity of model checking is in a lower complexity class. One is the fragment ESL^- , where we 1397
allow the operators $\exists x.\phi$ and $e_i(x)$ but restrict the use of the temporal operators to be those of the 1398
branching-time temporal logic CTL. In this case, we have the following result: 1399

THEOREM 4.3. *Let Σ be a PSPACE-presented class of strategies. The problem of deciding, given an 1400
environment E , a formula ϕ of ESL^- , and a context Γ for the free variables of ϕ relative to E and $\Sigma(E)$, 1401
whether $\Gamma, E, \Sigma(E) \models \phi$, is in PSPACE. 1402*

PROOF. We observe that the following fact follows straightforwardly from the semantics for 1403
formulas ϕ of ESL^- : For a context Γ for the free variables of ϕ relative to E and $\Sigma(E)$, and for two 1404
points (r, n) and (r', n') of $\mathcal{I}(E, \Sigma(E))$ with $r(n) = r'(n')$, we have that $\Gamma, \mathcal{I}(E, \Sigma(E)), (r, n) \models \phi$ iff 1405
 $\Gamma, \mathcal{I}(E, \Sigma(E)), (r', n') \models \phi$. That is, satisfaction of a formula relative to a context at a point depends 1406
only on the global state at the point and not on other details of the run containing the point. For 1407
a global state (s, α) of $\mathcal{I}(E, \Sigma(E))$, define the Boolean $SAT(\Gamma, E, \Sigma, (s, \alpha), \phi)$ to be TRUE just when 1408
 $\Gamma, \mathcal{I}(E, \Sigma(E)), (r, n) \models \phi$ holds for some point (r, n) of $\mathcal{I}(E, \Sigma(E))$ with $r(n) = (s, \alpha)$. By the above 1409
observation, we have that $\Gamma, E, \Sigma(E) \models \phi$ iff $SAT(\Gamma, E, \Sigma, (s, \alpha), \phi)$ holds for all initial states s of E 1410
and all strategies $\alpha \in \Sigma(E)$. Since we may check these conditions one at a time, strategies α can 1411
be represented in space linear in $|E|$, and deciding $\alpha \in \Sigma(E)$ is in PSPACE, it suffices to show that 1412
 $SAT(\Gamma, E, \Sigma, (s, \alpha), \phi)$ is decidable in PSPACE. 1413

We proceed by describing an APTIME algorithm for $SAT(\Gamma, E, \Sigma, (s, \alpha), \phi)$ and using the fact 1414
that $APTIME = PSPACE$ [12]. The algorithm operates recursively, with the following cases: 1415

- (1) If $\phi = p$, for $p \in Prop$, then return TRUE if $p \in \pi(s)$, else return FALSE. 1416
- (2) If $\phi = e_i(x)$, then return TRUE if $(s, \alpha)_i = \Gamma(x)_i$, else return FALSE. 1417
- (3) If $\phi = \phi_1 \wedge \phi_2$, then universally call $SAT(\Gamma, E, \Sigma, (s, \alpha), \phi_1)$ and 1418
 $SAT(\Gamma, E, \Sigma, (s, \alpha), \phi_2)$. 1419
- (4) If $\phi = \neg\phi_1$, then return the complement of $SAT(\Gamma, E, \Sigma, (s, \alpha), \phi_1)$. 1420
- (5) If $\phi = A \circ \phi_1$ then universally choose a state t such that $s \xrightarrow{a} t$ for some for some joint ac- 1421
tion a , and call $SAT(\Gamma, E, \Sigma, (t, \alpha), \phi_1)$. The other temporal operators from CTL are handled 1422
similarly. (In the case of operators using U , we need to run a search for a path through 1423
the set of states of E generated by the strategy α , but this is easily handled in APTIME.) 1424
- (6) If $\phi = D_G \phi_1$, then universally choose a global state (t, β) such that $(s, \alpha) \sim_G^D (t, \beta)$ and 1425
universally 1426
 - (a) decide whether $\beta \in \Sigma(E)$, and 1427
 - (b) call $REACH(t, \beta)$, and 1428
 - (c) call $SAT(\Gamma, E, \Sigma, (t, \beta), \phi_1)$. 1429

1430 (Here $REACH(t, \beta)$ decides whether state t is reachable in E from some initial state
 1431 when the agents run the joint strategy β ; this is trivially in PSPACE. Deciding $\beta \in \Sigma(E)$
 1432 is in PSPACE by the assumption that Σ is PSPACE-presented.)

- 1433 (7) If $\phi = C_G \phi_1$, then universally guess a global state (t, β) and universally do the following:
 1434 (a) Decide whether $(s, \alpha) \sim_G^C (t, \beta)$ using an existentially branching binary search for a
 1435 path of length at most $|S| \times |\Sigma(E)|$. For all states (u, γ) on this path it should be verified
 1436 that $REACH(u, \gamma)$ and that $\gamma \in \Sigma(E)$. The maximal length of the path is in the worst
 1437 case exponential in $|E|$, but the binary search can handle this in APTIME.
 1438 (b) call $SAT(\Gamma, E, \Sigma, (t, \beta), \phi_1)$.
 1439 (8) If $\phi = \exists x(\phi_1)$, then existentially guess a global state (t, β) , and universally
 1440 (a) decide if $\beta \in \Sigma(E)$, and
 1441 (b) call $REACH(t, \beta)$, and
 1442 (c) call $SAT(\Gamma[(t, \beta)/x], E, \Sigma, (s, \alpha), \phi_1)$.

1443 A straightforward argument based on the semantics of the logic shows that the above correctly
 1444 computes SAT.

1445 We remark that a more efficient procedure for checking that $(s, \alpha) \sim_G^C (t, \beta)$ is possible in the
 1446 typical case where $\Sigma(E)$ is a Cartesian product of sets of strategies for each of the agents. In this
 1447 case, if there exists a witness chain then there is one of length at most $|S|$. Let $G = G_1 \cup \sigma(G_2)$
 1448 such that $G_1, G_2 \subseteq \text{Ags}$. The number of steps through the relation $\cup_{i \in G} \sim_i$ required to witness
 1449 $(s, \alpha) \sim_G^C (t, \beta)$ depends on the sets G_1, G_2 as follows:

- 1450 (1) If $G_1 = G_2 = \emptyset$, then we must have $(s, \alpha) = (t, \beta)$ and a chain of length 0 suffices.
 1451 (2) If G_1 is nonempty and $G_2 = \emptyset$, then we must have $s(\cup_{i \in G_1} \sim_i)^* t$, but β can be arbitrary,
 1452 and this component can be changed in any step. A path of length $|S|$ suffices in this case.
 1453 (3) If $G_1 = \emptyset$ and $G_2 = \{i\}$ is a singleton, then we must have $\alpha_i = \beta_i$, but s and t can be arbitrary.
 1454 A path of length one suffices in this case.
 1455 (4) If $|G_1| \geq 1$, say, $i \in G_1$, and $G_2 = \{j\}$ is a singleton, then $(\cup_{i \in G} \sim_i)^*$ is the universal rela-
 1456 tion and a path of length 2 suffices. In particular, for any $(s, \alpha), (t, \beta)$ we have $(s, \alpha) \sim_i$
 1457 $(s, \beta) \sim_{\sigma(j)} (t, \beta)$.
 1458 (5) If $|G_2| \geq 2$, then $(\cup_{i \in G} \sim_i)^*$ is the universal relation and a path of length 2 suffices. In
 1459 particular, for any $(s, \alpha), (t, \beta)$ and distinct $i, j \in G_2$, there exists α' such that $\alpha'_i = \alpha_i$ and
 1460 $\alpha'_k = \beta_k$ for all $k \in \text{Ags}$ with $k \neq i$, and $(s, \alpha) \sim_{\sigma(i)} (s, \alpha') \sim_{\sigma(j)} (t, \beta)$. \square

1461 The following result shows that the PSPACE upper bound from this result is tight, already for
 1462 formulas that use strategy indices in the CTLK operators, but make no direct uses of the constructs
 1463 $\exists x$ and $e_i(x)$.

1464 **THEOREM 4.4.** *The problem of deciding, given an environment E for two agents and a formula ϕ of*
 1465 *CTLK(Ags \cup σ (Ags), Prop), whether $E, \Sigma^{unif, det}(E) \models \phi$ is PSPACE hard.*

1466 **PROOF.** We proceed by a reduction from the satisfiability of Quantified Boolean Formulae (QBF).
 1467 An instance of QBF is a formula ϕ of form

$$Q_1 x_1 \dots Q_n x_n (\gamma),$$

1468 where $Q_1, \dots, Q_n \in \{\exists, \forall\}$ and γ is a formula of propositional logic over propositions x_1, \dots, x_n .
 1469 The QBF problem is to decide, given a QBF instance ϕ , whether it is true. We construct an envi-
 1470 ronment E_ϕ and a formula ϕ^* of CTLK using strategic indices $\sigma(i)$ such that the QBF formula ϕ is
 1471 true iff we have $E_\phi, \Sigma^{unif, det}(E_\phi) \models \phi^*$.

1472 Given the QBF formula ϕ , we construct the environment $E_\phi = \langle S, I, \{Acts_i\}_{i \in \text{Ags}}, \rightarrow, \{O_i\}_{i \in \text{Ags}}, \pi \rangle$
 1473 for 2 agents $\text{Ags} = \{1, 2\}$ and propositions $\text{Prop} = \{p_0, \dots, p_n, q_1, q_2\}$ as follows:

1506 choices for variable values while preserving previous choices. It can then be shown that the QBF
 1507 formula ϕ is true iff $E_\phi, \Sigma^{unif, det} \models \phi^*$. \square

1508 Combining Theorem 4.3 and Theorem 4.4, we obtain the following:

1509 **COROLLARY 3.** *Let Σ be a PSPACE-presented class of strategies. The problem of deciding if*
 1510 *$\Gamma, E, \Sigma(E) \models \phi$, given an environment E , a formula ϕ of ESL^- and a context Γ for the free variables of*
 1511 *ϕ relative to E and $\Sigma(E)$, is PSPACE complete.*

1512 Since PSPACE is strictly contained in EXPSPACE, this result shows a strict improvement in
 1513 complexity as a result of the restriction to the CTL-based fragment. We remark that, by a triv-
 1514 ial generalization of the standard state labelling algorithm for model checking CTL to handle the
 1515 knowledge operators, the problem of model checking the logic $CTLK(Ags, Prop)$ in the system
 1516 $\mathcal{I}(\mathcal{E})$ generated by an epistemic transition system \mathcal{E} is in PTIME. Thus, there is a jump in com-
 1517 plexity from CTLK as a result of the move to the strategic setting, even without the addition of the
 1518 operators $\exists x.\phi$ and $e_i(x)$. However, this jump is not so large as the jump to the the full logic ESL.

1519 An orthogonal restriction of ESL is to retain the CTL* temporal basis, i.e., to allow full use of
 1520 LTL operators, but to allow epistemic operators and strategy indices but omit use of the operators
 1521 $\exists x.\phi$ and $e_i(x)$. This gives the logic $CTL^*K(Ags \cup \sigma(Ags), Prop)$. For this logic, we also see an im-
 1522 provement in the complexity of model checking compared to full ESL, as is shown in the following
 1523 result.

1524 **THEOREM 4.5.** *Let $\Sigma(E)$ be a PSPACE presented class of strategies for environments E . The com-*
 1525 *plexity of deciding, given an environment E and a CTL*K formula ϕ for agents $Ags(E)^+ \cup \sigma(Ags(E))$,*
 1526 *whether $E, \Sigma(E) \models \phi$, is PSPACE-complete.*

1527 **PROOF.** The lower bound is straightforward from the fact that linear time temporal logic LTL
 1528 is a sublanguage of CTL*K, and model checking LTL is already PSPACE-hard [52]. For the upper
 1529 bound, we describe an alternating PTIME algorithm and invoke the fact that $APTIME = PSPACE$
 1530 [12]. We abbreviate $\mathcal{I}(E, \Sigma(E))$ to \mathcal{I} .

1531 For a formula ϕ , write $\maxk(\phi)$ for the maximal epistemic subformulas of ϕ , defined to be the
 1532 set of subformulas of the form $A\psi$ or $C_G\psi$ or $D_G\psi$ for some set G of basic and strategic indices,
 1533 which are themselves not a subformula of a larger subformula of ϕ of one of these forms. Note
 1534 that $A\psi$ can be taken to be epistemic, because it is equivalent to $D_{\{e\} \cup \sigma(Ags)}\psi$; in the following,
 1535 we assume that $A\psi$ is written in this form. Also note that for epistemic formulas ψ , satisfaction at
 1536 a point depends only on the global state, i.e., for all points (r, m) and (r', m') of \mathcal{I} , we have that
 1537 if $r(m) = r'(m')$ then $\mathcal{I}, (r, m) \models \psi$ iff $\mathcal{I}, (r', m') \models \psi$. Thus, for global states (s, α) of \mathcal{I} , we may
 1538 write $\mathcal{I}, (s, \alpha) \models \psi$ to mean that $\mathcal{I}, (r, m) \models \psi$ for some point (r, m) with $r(m) = (s, \alpha)$.

1539 Define a ϕ -labelling of E to be a mapping $L : S \times \maxk(\phi) \rightarrow \{0, 1\}$, giving a truth value for each
 1540 maximal epistemic subformula of ϕ . A ϕ -labelling can be represented in space $|S| \times |\phi|$. Note that
 1541 if we treat the maximal epistemic subformulas of ϕ as if they were atomic propositions, evaluated
 1542 at the states of E using the ϕ -labelling L , then ϕ becomes an LTL formula, evaluable on any path in
 1543 E with respect to the labelling L . Verifying that all α -paths from a state s satisfy ϕ with respect to L
 1544 is then exactly the problem of LTL model checking, for which there exists an APTIME procedure
 1545 $ASAT(E, (s, \alpha), L, \phi)$, since model checking LTL is in PSPACE [52] and $APTIME = PSPACE$ [12].
 1546 For this to correspond to model checking in \mathcal{I} , we require that the ϕ -labelling L gives the correct
 1547 answers for the truth value of the formula at each state (s, α) , i.e., that $L(\psi) = 1$ iff $\mathcal{I}, (s, \alpha) \models \psi$.
 1548 We handle this by means of a guess and verify technique.

1549 To handle the verification, an alternating PTIME algorithm $KSAT(E, \Sigma, (s, \alpha), \phi)$ is defined, for ϕ
 1550 an epistemic formula, such that $KSAT(E, \Sigma, (s, \alpha), \phi)$ returns TRUE iff $\mathcal{I}, (s, \alpha) \models \phi$. The definition

is recursive and uses a call to the procedure ASAT. Specifically, $\text{KSAT}(E, \Sigma, (s, \alpha), D_G\phi)$ operates as follows:

- (1) universally guess a state t of E and a joint strategy β in E , then
- (2) verify that t is reachable in E using joint strategy β , that $(s, \alpha) \sim_G (t, \beta)$, and that β is in $\Sigma(E)$, then
- (3) existentially guess a ϕ -labelling L of E , then
- (4) universally,
 - (a) call $\text{ASAT}(E, (t, \beta), L, \phi)$, and
 - (b) for each state w and formula $\psi \in \text{maxk}(\phi)$, call $\text{KSAT}(E, \Sigma, (w, \beta), \psi)$.

Note that step 4(b) verifies that the ϕ -labelling L is correct.

For $\text{KSAT}(E, \Sigma, (s, \alpha), C_G\phi)$, the procedure is similar, except that instead of verifying that $(s, \alpha) \sim_G (t, \beta)$ in the second step, we need to verify that $(s, \alpha) (\cup_{i \in G} \sim_i)^* (t, \beta)$. This is easily handled in APTIME by a standard recursive procedure that guesses a midpoint of the path and branches universally to verify the existence of the left and right halves of the chain. (See the proof of Theorem 4.3 for some further discussion on this point.)

To solve the model-checking problem in \mathcal{I} , we can now apply the following alternating procedure:

- (1) universally guess a global state (s, α) of \mathcal{I} , then branch existentially to the following cases:
 - (a) if s is an initial state of E return FALSE, else return TRUE,
 - (b) if $\alpha \in \Sigma(E)$, return FALSE, else return TRUE,
 - (c) call $\text{KSAT}(E, \Sigma, (s, \alpha), A\phi)$.

Evidently, each of the alternating procedures runs in polynomial time internally, and the number of recursive calls is $O(|\phi|)$. It follows that the entire computation is in $\text{APTIME} = \text{PSPACE}$. \square

It is interesting to note that, although $\text{CTL}^*K(\text{Ags} \cup \sigma(\text{Ags}))$ is significantly richer than the temporal logic LTL, the added expressiveness comes without an increase in complexity: model checking LTL is already PSPACE-complete [52].

5 CONCLUSION

We now discuss some related work and remark on some questions for future research. The sections above have already made some references and comparisons to related work on each of the topics that we cover. Beside these references, the following are also worth mentioning.

Semantics that explicitly encode strategies in runs have been used previously in the literature on knowledge in information flow security [25]; what is novel in our approach is to develop a logic that enables explicit reference to these strategies.

A variant of propositional dynamic logic (PDL) for describing strategy profiles in normal form games subject to preference relations is introduced in Reference [54]. This work does not cover temporal aspects as we have done in this article. Another approach based on PDL is given in Reference [47], which describes strategies by means of formulas.

A very rich generalization of ATEL for probabilistic environments is described in Reference [49]. This proposal includes variables that refer to *strategy choices*, and strategic operators that may refer to these variables, so that statements of the form “when coalition A runs the strategy represented by variable S1, and coalition B runs the strategy represented by variable S2, and the remaining agents behave arbitrarily, then the probability that ϕ holds is at least δ ” can be expressed. Here a strategy choice maps each state, coalition and formula to a uniform imperfect recall strategy for the coalition. There are a number of syntactic restrictions compared to our logic. The epistemic

1595 operators in this approach apply only to state formulas rather than path formulas (in the sense of
 1596 this distinction from CTL*.) Moreover, the strategic variables may be quantified, but only in the
 1597 prefix of the formula. These constraints imply that notions such as “agent i knows that there exists
 1598 a strategy by which it can achieve ϕ ” and “agent i knows that it has a winning response to every
 1599 strategy chosen by agent j ” cannot be naturally expressed.

1600 The extended temporal epistemic logic ETLK we have introduced, of which our epistemic strat-
 1601 egy logic ESL is an instantiation with respect to a particular semantics, uses constructs that resem-
 1602 ble constructs from *hybrid logic* [4]. Hybrid logic is an approach to the extension of modal logics
 1603 that uses “nominals,” i.e., propositions p that hold at a single world. These can be used in combi-
 1604 nation with operators such as $\exists p$, which marks an arbitrary world as the unique world at which
 1605 nominal p holds. Our construct $\exists x$ is closely related to the hybrid construct $\exists p$, but we work in
 1606 a setting that is richer in both syntax and semantics than previous works. There have been a few
 1607 works using hybrid logic ideas in the context of epistemic logic [27, 48] but none are concerned
 1608 with temporal logic. Hybrid temporal logic has seen a larger amount of study [5, 21, 22, 51], with
 1609 variances in the semantics used for the model-checking problem.

1610 We note that if we were to view the variable x in our logic as a propositional constant, it would
 1611 be true at a set of points in the system $\mathcal{I}(E, \Sigma)$, hence not a nominal in that system. Results in
 1612 Reference [5], where a hybrid linear time temporal logic formula is checked in all paths in a given
 1613 model, suggest that a variant of ESL in which x is treated as a nominal in $\mathcal{I}(E, \Sigma)$ would have a
 1614 complexity of model checking at least non-elementary, compared to our EXPSPACE and PSPACE
 1615 complexity results.

1616 Our PSPACE model-checking result for $\text{CTLK}(\text{Ags} \cup \sigma(\text{Ags}))$ seems to be more closely related to
 1617 the result in Reference [21] that model checking a logic $\text{HL}(\exists, @, F, A)$ is PSPACE-complete. Here
 1618 F is essentially a branching time future operator and A is a universal operator (similar to our D_\emptyset),
 1619 the construct $@_p\phi$ says that ϕ holds at the world marked by the nominal p , and $\exists p(\phi)$ says that ϕ
 1620 holds after marking some world by p . The semantics in this case does not unfold the model into
 1621 either a tree or a set of linear structures before checking the formula, so the semantics of the hybrid
 1622 existential \exists is close to our idea of quantifying over global states. Our language, however, has a
 1623 richer set of operators, even in the temporal dimension, and introduces the strategic dimension in
 1624 the semantics. It would be an interesting question for future work to consider fragments of our
 1625 language to obtain a more precise statement of the relationship with hybrid temporal logics.

1626 Strategy Logic [13] is a (non-epistemic) generalization of ATL for perfect information strategies
 1627 in which strategies may be explicitly named and quantified. Strategy logic has a non-elementary
 1628 model-checking problem. Work on identification of more efficient variants of quantified strategy
 1629 logic includes [42], who formulate a variant with a 2-EXPTIME-complete model-checking problem.
 1630 In both cases, strategies are perfect recall strategies, rather than the imperfect recall strategies that
 1631 form the basis for our PSPACE-completeness result for model checking.

1632 Most closely related to this article are a number of independently developed works that consider
 1633 epistemic extensions of variants of strategy logic. Belardinelli [3] develops a logic, based on linear
 1634 time temporal logic with epistemic operators, that adds an operator $\exists x_i$, the semantics of which
 1635 existentially modifies the strategy associated to agent i in the current strategy profile. It omits
 1636 the binding operator from Reference [42], so provides no other way to refer to the variable x .
 1637 The logic is shown to have nonelementary model-checking complexity. This complexity is higher
 1638 than the results we have presented, because the semantics for strategies allows agents to have
 1639 perfect information and perfect recall (though the semantics for the knowledge operators is based
 1640 on imperfect information and no recall), whereas we have assumed imperfect information and no
 1641 recall for strategies.

Another extension of strategy logic with epistemic operators has been independently developed 1642 by Čermák et al. [10, 11]. Their syntax and semantics differs from ours in a number of respects. 1643 Although the syntax appears superficially in the form of an extension of LTL, it is more like CTL 1644 in some regards. The transition relation is deterministic in the sense that for each joint action, 1645 each state has a unique successor when that action is performed. Strategies are also assumed to 1646 be deterministic (whereas we allow nondeterministic strategies.) This means that, like CTL, the 1647 semantics of a formula depends only on the current global state and the current strategy profile, 1648 whereas for LTL it is generally the case that the future structure of the run from a given global state 1649 can vary, and the truth value of the formula depends on how it does so. Although it seems that 1650 non-determinism could be modelled, as is commonly done, through the choice of actions of the 1651 environment, treated as an agent, the fact that strategies are deterministic, uniform and memory- 1652 less means that the environment must choose the same alternative each time a global state occurs 1653 in a run. This means that this standard approach to modelling of non-determinism does not work 1654 for this logic. The syntax of the logic moreover prevents epistemic operators from being applied 1655 to formulas with free strategy variables, whereas we allow fully recursive mixing of the constructs 1656 of our logic. Consequently, epistemic notions from our logic like $D_{\{i, \sigma(i)\}}$, expressing an agent's 1657 knowledge about the effects of its own strategy, which are used in several of our applications, do 1658 not appear to be expressible in this logic. Finally, the notion of "interpreted system" in this work, 1659 which corresponds most closely to our notion of "environment," also seems less general than our 1660 notion of environment, because it defines the accessibility relations for the knowledge operators 1661 in a way that makes the environment state known to all agents. 1662

In another article [32], we have implemented a symbolic algorithm that handles model checking 1663 for the fragment $\text{CTLK}(Ags \cup \sigma(Ags))$, which, as shown above, encompasses the expressiveness 1664 of ATEL. Existing algorithms described in the literature for ATEL model checking [8, 9, 40] are 1665 based either on explicit-state model checking or are only partially symbolic in that they iterate 1666 over all strategies, explicitly represented. Our experimental results in Reference [32] show that by 1667 comparison with the partially symbolic approach, a fully-symbolic algorithm can greatly improve 1668 the performance and therefore scalability of model checking. The approach to model-checking 1669 epistemic strategy logic implemented in Reference [10, 11] is fully symbolic, but as already men- 1670 tioned, this logic has a more limited expressive power than ours and its semantics does not permit 1671 representation of a nondeterministic environment. (It does not seem that the semantics could be 1672 extended to allow nondeterminism while retaining correctness of their algorithm.) 1673

Our focus on this article has been on an observational, or imperfect recall, semantics for knowl- 1674 edge. Other semantics for knowledge are also worth considering, but are left for future work. We 1675 note one issue in relation to the connection to ATEL that we have established, should we consider 1676 a perfect recall version of our logic. ATEL operators effectively allow reference to situations in 1677 which agents switch their strategy after some actions have already been taken, whereas in our 1678 model an agent's strategy is fixed for the entire run. When switching to a new strategy, there is 1679 the possibility that the given state is not reachable under this new strategy. We have handled this 1680 issue in our translation by assuming that all states are initial, so that the run can be reinitialized if 1681 necessary to make the desired state reachable. This is consistent with an imperfect recall interpre- 1682 tation of ATEL, but it is not clear that this approach is available on a perfect recall interpretation. 1683 We leave a resolution of this issue to future work. 1684

REFERENCES

- [1] Thomas Ágotnes, Valentin Goranko, and Wojciech Jamroga. 2007. Alternating-time temporal logics with irrevocable 1685 strategies. In *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'07)*. 15–24. 1686 DOI: <http://dx.doi.org/10.1145/1324249.1324256> 1687

- 1688 [2] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (2002),
1689 672–713.
- 1690 [3] Francesco Belardinelli. 2014. Reasoning about knowledge and strategies: Epistemic strategy logic. In *Proceedings of*
1691 *the 2nd International Workshop on Strategic Reasoning (SR'14)*. 27–33.
- 1692 [4] Patrick Blackburn and Jerry Seligman. 1998. What are hybrid languages? In *Advances in Modal Logic*, M. de Rijke, H.
1693 Wansing, and M. Zakharyashev (Eds.), Vol. 1. CSLI Publications, 41–62.
- 1694 [5] Laura Bozzelli and Ruggero Lanotte. 2010. Complexity and succinctness issues for linear-time hybrid logics. *Theor.*
1695 *Comput. Sci.* 411, 2 (2010), 454–469. DOI: <http://dx.doi.org/10.1016/j.tcs.2009.08.009>
- 1696 [6] Ronen I. Brafman, Jean-Claude Latombe, Yoram Moses, and Yoav Shoham. 1997. Applications of a logic of knowledge
1697 to motion planning under uncertainty. *J. ACM* 44, 5 (1997), 633–668.
- 1698 [7] Thomas Brihaye, Arnaud Da Costa, François Laroussinie, and Nicolas Markey. 2009. ATL with strategy contexts and
1699 bounded memory. In *Proceedings of the International Symposium on Logical Foundations of Computer Science*. 92–106.
- 1700 [8] Simon Busard, Charles Pecheur, Hongyang Qu, and Franco Raimondi. 2013. Reasoning about strategies under partial
1701 observability and fairness constraints. In *Proceedings of the 1st International Workshop on Strategic Reasoning (SR'13)*.
1702 71–79.
- 1703 [9] Jan Calta, Dmitry Shkatov, and Bernd-Holger Schlingloff. 2010. Finding uniform strategies for multi-agent systems.
1704 In *Computational Logic in Multi-Agent Systems (CLIMA XI)*. 135–152.
- 1705 [10] Petr Čermák. 2014. *A Model Checker for Strategy Logic*. MEng Individual Project thesis, Department of Computing,
1706 Imperial College London.
- 1707 [11] Petr Čermák, Alessio Lomuscio, Fabio Mogavero, and Aniello Murano. 2014. MCMAS-SLK: A model checker for the
1708 verification of strategy logic specifications. In *Proceedings of the 26th International Conference on Computer-Aided*
1709 *Verification (CAV'14)*. 525–532.
- 1710 [12] Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. 1981. Alternation. *J. ACM* 28, 1 (1981), 114–133.
- 1711 [13] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. 2010. Strategy logic. *Inf. Comput.* 208, 6 (2010),
1712 677–693. DOI: <http://dx.doi.org/10.1016/j.ic.2009.07.004>
- 1713 [14] Stephen Chong and Andrew C. Myers. 2005. Language-based information erasure. In *Proceedings of the IEEE Computer*
1714 *Security Foundations Workshop*. 241–254.
- 1715 [15] Edmund M. Clarke, E. Allen Emerson, and A. Prasad Sistla. 1986. Automatic verification of finite-state concurrent
1716 systems using temporal logic specifications. *ACM Trans. Program. Lang. Syst.* 8, 2 (1986), 244–263.
- 1717 [16] Catalin Dima and Ferucio Laurentiu Tiplea. 2011. Model-checking ATL under imperfect information and perfect recall
Q6 1718 semantics is undecidable. *CoRR* abs/1102.4225 (2011).
- 1719 [17] E. Allen Emerson and Joseph Y. Halpern. 1986. “Sometimes” and “not never” revisited: On branching versus linear
1720 time temporal logic. *J. ACM* 33, 1 (1986), 151–178.
- 1721 [18] Kai Engelhardt, Peter Gammie, and Ron van der Meyden. 2007. Model checking knowledge and linear time: PSPACE
1722 cases. In *Proceedings of the Symposium on Logical Foundations of Computer Science*. LNCS. Springer, 195–211.
- 1723 [19] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. 1995. *Reasoning about Knowledge*. The MIT Press.
- 1724 [20] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. 1997. Knowledge-based programs. *Distrib. Comput.*
1725 10, 4 (1997), 199–225.
- 1726 [21] Massimo Franceschet and Maarten de Rijke. 2006. Model checking hybrid logics (with an application to semistruc-
1727 tured data). *J. Appl. Logic* 4, 3 (2006), 279–304. DOI: <http://dx.doi.org/10.1016/j.jal.2005.06.010>
- 1728 [22] Massimo Franceschet, Maarten de Rijke, and Bernd-Holger Schlingloff. 2003. Hybrid logics on linear structures: Ex-
1729 pressivity and complexity. In *Proceedings of the 10th International Symposium on Temporal Representation and Rea-*
1730 *soning and the 4th International Conference on Temporal Logic (TIME-ICTL'03)*. 166–173.
- 1731 [23] Joseph Y. Halpern and Yoram Moses. 1990. Knowledge and common knowledge in a distributed environment. *J. ACM*
1732 37, 3 (1990), 549–587.
- 1733 [24] Joseph Y. Halpern and Yoram Moses. 2007. Characterizing solution concepts in games using knowledge-based pro-
1734 grams. In *Proceedings of the 20nd International Joint Conference on Artificial Intelligence (IJCAI'07)*. 1300–1307.
- 1735 [25] Joseph Y. Halpern and Kevin R. O’Neill. 2008. Secrecy in multiagent systems. *ACM Trans. Inf. Syst. Secur.* 12, 1 (Article
1736 No. 5 2008).
- 1737 [26] Joseph Y. Halpern and Nan Rong. 2010. Cooperative equilibrium (extended abstract). In *Proceedings of the 9th Inter-*
1738 *national Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*. 1465–1466.
- 1739 [27] Jens Ulrik Hansen. 2011. A hybrid public announcement logic with distributed knowledge. *Electr. Not. Theor. Comput.*
Q7 1740 *Sci.* 273 (2011), 33–50. DOI: <http://dx.doi.org/10.1016/j.entcs.2011.06.011>
- 1741 [28] Wiebe van der Hoek, Wojciech Jamroga, and Michael Wooldridge. 2005. A logic for strategic reasoning. In *Proceedings*
1742 *of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'05)*. 157–164.
- 1743 [29] Wiebe van der Hoek and Michael Wooldridge. 2002. Tractable multiagent planning for epistemic goals. In *Proceedings*
1744 *of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'02)*. 1167–1174.

- [30] John F. Horty. 2001. *Agency and Deontic Logic*. Oxford University Press. 1745
- [31] Xiaowei Huang and Ron van der Meyden. 2014. An epistemic strategy logic (extended abstract). In *Proceedings of the 2nd International Workshop on Strategic Reasoning (SR14)*. 35–41. 1746
- [32] Xiaowei Huang and Ron van der Meyden. 2014. Symbolic model checking epistemic strategy logic. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI'14)*. 1426–1432. 1747
- [33] Xiaowei Huang and Ron van der Meyden. 2014. A temporal logic of strategic knowledge. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*. 1750
- [34] Wojciech Jamroga. 2003. Some remarks on alternating temporal epistemic logic. In *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS'03)*. 1752
- [35] Wojciech Jamroga and Thomas Ågotnes. 2007. Constructive knowledge: What agents can achieve under imperfect information. *J. Appl. Non-Class. Logics* 17, 4 (2007), 423–475. DOI : <http://dx.doi.org/10.3166/jancl.17.423-475> 1753
- [36] Wojciech Jamroga, Thomas Ågotnes, and Wiebe van der Hoek. 2008. *A Simpler Semantics for Abilities under Uncertainty*. Technical Report IfI-08-01. Department of Informatics, Clausthal University of Technology. 1754
- [37] Wojciech Jamroga and Jürgen Dix. 2006. Model checking abilities under incomplete information is indeed delta-2-complete. In *Proceedings of the 4th European Workshop on Multi-Agent Systems (EUMAS'06)*. 1755
- [38] Wojciech Jamroga and Wiebe van der Hoek. 2004. Agents that know how to play. *Fund. Inf.* 62 (2004), 1–35. 1756
- [39] Geert Jonker. 2003. *Feasible Strategies in Alternating-time Temporal*. Master's thesis. University of Utrecht, The Netherlands. 1757
- [40] Alessio Lomuscio and Franco Raimondi. 2006. Model checking knowledge, strategies, and games in multi-agent systems. In *Proceedings of the Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06)*. 161–168. 1758
- [41] Ron van der Meyden. 1996. Knowledge based programs: On the complexity of perfect recall in finite environments. In *Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'96)*. 31–49. 1759
- [42] Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi. 2010. Reasoning about strategies. In *Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'10)*. 133–144. DOI : <http://dx.doi.org/10.4230/LIPICs.FSTTCS.2010.133> 1760
- [43] Rohit Parikh. 1983. Propositional game logic. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*. 195–200. 1761
- [44] Rohit Parikh and Ramaswamy Ramanujam. 1985. Distributed processes and the logic of knowledge. In *Logics of Programs 1985*. 256–268. 1762
- [45] Marc Pauly. 2002. A modal logic for coalitional power in games. *J. Logic Comput.* 12, 1 (2002), 149–166. 1763
- [46] Amir Pnueli. 1977. The temporal logic of programs. In *Proceedings of the Symposium on Foundations of Computer Science*. 46–57. 1764
- [47] Ramaswamy Ramanujam and Sunil Easaw Simon. 2008. Dynamic logic on games with structured strategies. In *Proceedings of the 1th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*. 49–58. 1765
- [48] Olivier Roy. 2009. A dynamic-epistemic hybrid logic for intentions and information changes in strategic games. *Synthese* 171, 2 (2009), 291–320. DOI : <http://dx.doi.org/10.1007/s11229-009-9644-3> 1766
- [49] Henning Schnoor. 2010. *Explicit Strategies and Quantification for ATL with Incomplete Information and Probabilistic Games*. Technical Report 1008. Institut für Informatik, Christian-Albrechts Universität zu Kiel. 1767
- [50] Pierre-Yves Schobbens. 2004. Alternating-time logic with imperfect recall. *Electr. Not. Theor. Comput. Sci.* 85, 2 (2004), 82–93. DOI : [http://dx.doi.org/10.1016/S1571-0661\(05\)82604-0](http://dx.doi.org/10.1016/S1571-0661(05)82604-0) 1768
- [51] Thomas Schwentick and Volker Weber. 2007. Bounded-variable fragments of hybrid logics. In *Proceedings of the 24th Annual Symposium on Theoretical Aspects of Computer Science (STACS'07)*. Springer LNCS, Vol. 4393. 561–572. DOI : http://dx.doi.org/10.1007/978-3-540-70918-3_48 1769
- [52] A. Prasad Sistla and Edmund M. Clarke. 1985. The complexity of propositional linear temporal logics. *J. ACM* 32, 3 (1985), 733–749. 1770
- [53] D. Sutherland. 1986. A model of information. In *Proceedings of the 9th National Computer Security Conference*. 175–183. 1771
- [54] Jan van Eijck. 2013. PDL as a multi-agent strategy logic. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowledge*. 1772
- [55] Sieuwert van Otterloo and Geert Jonker. 2005. On epistemic temporal strategic logic. *Electronic Notes in Theoretical Computer Science* 126 (2005), 77–92. DOI : <http://dx.doi.org/10.1016/j.entcs.2004.11.014> 1773
- [56] Moshe Y. Vardi. 1996. Implementing knowledge-based programs. In *Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge*. 15–30. 1774
- [57] J. Todd Wittbold and Dale M. Johnson. 1990. Information flow in nondeterministic systems. In *Proceedings of the IEEE Symposium on Security and Privacy*. 144–161. 1775

Received November 2017; revised March 2018; accepted June 2018

1800

Author Queries

- Q1:** AU: Please provide CCS 2012 Concepts per author guidelines and provide XML codes as well.
- Q2:** AU: Please provide Additional Key Words and Phrases section.
- Q3:** AU: Please provide complete mailing and email addresses for all authors.
- Q4:** AU: Initial cap “low” and “high” with “low-level attacker” and “high-level secret”?
- Q5:** AU: Sentence “Another work by van der Hoek, Jamroga...”: use a comma after “work” and then after “[28]”?
- Q6:** AU: Ref. 16: If possible, please update all CoRR and arXiv references in this list.
- Q7:** AU: Ref. 27: Please provide missing volume or issue number.
- Q8:** AU: Ref. 38: please provide missing issue or volume number.
- Q9:** AU: Ref. 54: Please update and complete.