A Logic of Knowledge and Strategies with Imperfect Information

Francesco Belardinelli

Laboratoire IBISC, Université d'Evry, France, email: belardinelli@ibisc.fr

Abstract. In this paper we put forward Epistemic Strategy Logic (ESL), a logic of knowledge and strategies in contexts of imperfect information. ESL extends Strategy Logic by adding modal operators for individual and collective knowledge. This enhanced framework allows us to represent explicitly and to reason about the knowledge agents have of their own and other agents' strategies. We provide a semantics to ESL in terms of epistemic concurrent game models, then illustrate the expressive power of ESL as a specification language for games, both of perfect and imperfect information. Notably, we show that some fixed-point characterisations of operators for strategic abilities, which normally fail in contexts of imperfect information, can be recovered in a controlled way through the interplay of epistemic and strategy modalities.

1 Introduction

In recent years logic-based formalisms for representing and reasoning about strategic abilities, both individual and coalitional, have been a thriving area of research in Artificial Intelligence and Multi-agent System [4, 5, 12]. A diverse family of multi-modal logics has been introduced to provide a formal account of complex strategic reasoning and behaviours for individual agents and groups, including Alternating-time Temporal Logic (ATL), Strategy Logic, Coalition Logic to name just a few [2, 7, 27]. In parallel with these developments, a well-established tradition in knowledge representation focuses on extending formalisms for reactive systems with epistemic operators, so as to reason about the systems' evolution, as well as the knowledge agents have thereof [9]. Seminal contributions on extensions of linear- and branching-time temporal logics with agent-indexed epistemic modalities date back to the '80s [13, 14]. Since then, these investigations have matured into a solid body of works, which is nowadays rightly regarded as a key contribution of formal methods to computer science [23], particularly when combined with verification techniques [10, 20, 22].

By bringing together these two lines of research, [15] introduced Alternating-time Temporal Epistemic Logic (ATEL), an extension of ATL with epistemic operators; thus paving the way for a wealth of contributions on sophisticated interactions between strategies and knowledge. More directly related to the present work, [3] and [16] put forward two variants of Epistemic Strategy Logic (ESL) with perfect knowledge, by building on the (non-epistemic) Strategy Logics in

[7, 24]. These investigations are relevant from both a purely theoretical perspective and an applied viewpoint. Indeed, logics of time and knowledge have been successfully deployed in the verification by model checking of distributed and Multi-agent Systems as diverse as security protocols, UAVs, web services, and e-commerce [6, 10, 22, 20].

In this paper we advance the state-of-the-art in Epistemic Strategy Logic by exploring the theoretical properties of ESL in contexts of imperfect information. Specifically, we analyse and compare the expressive power of ESL as a specification language for games, both of perfect and imperfect information. It is well-known that in logics of strategies the latter assumption easily leads to higher complexity and even undecidability of the satisfiability and model checking problems [4, 5, 28]. This is due in part to the weaker semantical properties of strategies under imperfect information. In this paper we describe these features both syntactically and semantically, and show how some fixed-point characterisations of operators for strategic abilities, which normally fail in contexts of imperfect information, can be partially recovered through the interaction of epistemic and strategy modalities. These results might point to the development of decision procedures for the cases in hand [11].

Related Work. The present contribution draws inspiration from the extensive literature on logics of strategic abilities, especially Strategy Logic (SL). This formalism has been introduced in [7] for concurrent game structures (CGS) with two-players. In [24] CGS have been extended to a multi-player setting and bind operators for strategies have been introduced in the syntax. More recently, various notions of strategy (e.g., *behavioural* [25, 26]) have been analysed within SL. Hereafter we adopt multi-agent CGS in line with [24]. However, by drawing on the *Interpreted Systems* semantics for temporal epistemic logic [9], we provide an agent-based semantics to ESL, in which agents have possibly different actions and protocols, in contrast with [24–26].

To the best of our knowledge, epistemic extensions of SL have been recently considered in [3, 6, 16]; however always in contexts of perfect information. In particular, [6] describes MCMAS-SLK, a tool to model check CGS against specifications in an epistemic extension of SL, containing also binding operators. We extend the state-of-the-art w.r.t. [3, 6, 16] by analysing ESL with imperfect information, which has not yet been attempted. We will see that this endeavour requires several syntactical and semantical innovations, and it provides interesting new insight on the epistemic dimension of strategies.

Besides Strategy Logic, the interaction between knowledge and strategies has already been thoroughly studied within ATL. The wealth of contributions on the subject is testament to the interest and relevance of the area, of which it is impossible to give an exhaustive overview. Hereafter we briefly consider the works most directly related to the present contribution. The alternating-time temporal logic ATEL was introduced in [15], and immediately imperfect information variants were considered in [18], which put forward alternating-time temporal observational logic (ATOL) and ATEL-R*, as well as the key notion of *uniform strategy*. Interestingly, [18] discusses the distinction between *de re* and *de dicto* knowledge of strategies, but a formal account is not developed explicitly. The same distinction will be considered later on in the context of ESL with imperfect information. Further, in [17] ATL is enriched with a constructive notion of knowledge. As regards (non-epistemic) ATL, more elaborate notions of strategy have also appeared: [1] introduces commitment in strategies; while [19] defines a notion of "feasible" strategy. In future work it might be worth exploring to what extent, if any, the theoretical results available for the various flavours of ATEL transfer to ESL.

Scheme of the paper. In Section 2 we introduce the epistemic concurrent game models (ECGM), which are used in Section 3 to provide a semantics to Epistemic Strategy Logic (ESL). We show that ECGM are suitable to model games in normal form, while properties such as the existence of Nash equilibria, and knowledge thereof, can be specified in ESL. In Section 4 we analyse ESL in contexts of imperfect information. In particular, we show how some characterisations of SL operators can be recast by combining epistemic and strategy modalities. Finally, in Section 5 we discuss the results obtained and point to future research. For reasons of space, only sketches of proofs are provided.

2 Epistemic Concurrent Game Models

This section is devoted to introducing epistemic concurrent game models (ECGM) and related notions [3], starting from the definition of *agent*.

Definition 1 (Agent). *An* agent *is a tuple* $a = \langle L_a, Act_a, Pr_a \rangle$ *such that*

- L_a *is the set of* local states $l_a, l'_a, \ldots$*;*
- Act_a *is a finite set of* actions $\sigma_a, \sigma'_a, \ldots$*;*
- $Pr_a : L_a \mapsto 2^{Act_a}$ *is the* protocol function*.*

An agent a can be thought of as situated in some local state $l_a \in L_a$, that represents the information she has access to, and as performing the actions in Act_a according to the protocol Pr_a [9]. Differently from [24], hereafter agents might have different actions and protocols. Given a set $Ag = \{a_0, \ldots, a_n\}$ of agents, we define the set $\mathcal{G}$ of *global states* $s, s', \ldots$ (resp. the set Act of *joint actions* $\sigma, \sigma', \ldots$) as the cartesian product $\prod_{a \in Ag} L_a$ (resp. $\prod_{a \in Ag} Act_a$). We now introduce ECGM as the systems generated by the agents' interactions. In what follows we fix a set AP of atomic propositions and denote the jth component of a tuple t as t_j or, equivalently, $t(j)$.

Definition 2 (ECGM). *Given a set* $Ag = \{a_0, \ldots, a_n\}$ *of agents* $a = \langle L_a, Act_a, Pr_a \rangle$*, an* epistemic concurrent game model *is a tuple* $\mathcal{P} = \langle Ag, I, \tau, \pi \rangle$ *such that*

- $I \subseteq \mathcal{G}$ *is the set of* initial global states s_0*;*
- $\tau : \mathcal{G} \times Act \mapsto \mathcal{G}$ *is the* global transition function*, where* $\tau(s, \sigma)$ *is defined iff* $\sigma_a \in Pr_a(l_a)$ *for every* $a \in Ag$*;*
- $\pi : AP \to 2^{\mathcal{G}}$ *is the* interpretation function*.*

Intuitively, ECGM can be seen as an agent-based variant of concurrent game structures [2, 15]. An ECGM describes the evolution of a system from an initial state $s_0 \in I$, according to the transition function τ.

We now introduce some notation to be used in the following. The *transition relation* $s \to s'$ on states holds iff $\tau(s, \sigma) = s'$ for some $\sigma \in Act$. A *run* λ from a state s, or s-run, is an infinite sequence $s^0 \to s^1 \to \ldots$, where $s^0 = s$. For $n, m \in \mathbb{N}$, with $n \leq m$, we define $\lambda(n) = s^n$ and $\lambda[n, m] = s^n, s^{n+1}, \ldots, s^m$. A state s' is *reachable from* s iff $\lambda(i) = s'$ for some s-run λ and $i \geq 0$. We define $\mathcal{S}$ as the set of reachable states. Further, let $\sharp$ be a placeholder for arbitrary individual actions. Given a subset $A \subseteq Ag$ of agents, an A-*action* σ_A is an $|Ag|$-tuple s.t. *(i)* $\sigma_A(a) \in Act_a$ for $a \in A$, and *(ii)* $\sigma_A(b) = \sharp$ for $b \notin A$. Then, Act_A is the set of all A-actions and $d_A(s) = \{\sigma_A \in Act_A \mid \text{for every } a \in A, \sigma_a \in Pr_a(l_a)\}$ is the set of all A-actions enabled at $s = \langle l_0, \ldots, l_n \rangle$. A joint action σ *extends* an A-action σ_A, or $\sigma_A \sqsubseteq \sigma$, iff $\sigma(a) = \sigma_A(a)$ for all $a \in A$. The *outcome* $out(s, \sigma_A)$ *of action* σ_A *at state* s is the set of all states s' s.t. $\tau(s, \sigma) = s'$ for some joint action $\sigma \sqsupseteq \sigma_A$. Finally, two global states $s = \langle l_0, \ldots, l_n \rangle$ and $s' = \langle l'_0, \ldots, l'_n \rangle$ are *indistinguishable* for agent a, or $s \sim_a s'$, iff $l_a = l'_a$ [9]. Similarly, for a set A of agents, $s \sim_A s'$ iff $l_a = l'_a$ for every $a \in A$, i.e., $(s, s') \in \bigcap_{a \in A} \sim_a$.

To illustrate the modelling power of ECGM, we show how a relevant class of games can be modelled within this framework. We start with a game of perfect information to compare and contrast with imperfect information later on.

Example 1. Epistemic concurrent game models are expressive enough to model games in normal form. As an example, we consider the classic version of the prisoner's dilemma [8], with players 1 and 2, who can either cooperate (**C**) or defect (**D**), and payoff ordering $\alpha > \beta > \gamma > \delta$. This can be represented in normal form as:

		2	
		Cooperate	**D**efect
1	**C**ooperate	β, β	δ, α
	Defect	α, δ	γ, γ

The prisoner's dilemma can be modelled as an ECGM by considering agents $1 = \langle L_1, Act_1, Pr_1 \rangle$ and $2 = \langle L_2, Act_2, Pr_2 \rangle$ s.t. for $a \in \{1, 2\}$, *(i)* $L_a = \{\epsilon_a, \alpha, \beta, \gamma, \delta\}$; *(ii)* $Act_a = \{\mathbf{C}, \mathbf{D}, *\}$; *(iii)* $Pr_a(\epsilon_a) = \{\mathbf{C}, \mathbf{D}\}$ and $Pr_a(\alpha) = Pr_a(\beta) = Pr_a(\gamma) = Pr_a(\delta) = \{*\}$, where $*$ represents the skip action. Also, for each payoff p we consider an atomic propositions p_a, which intuitively expresses that the local state of agent a is equal to the corresponding payoff. We then introduce the ECGM $\mathcal{P}_{pd} = \langle \{1, 2\}, \{s_0\}, \tau, \pi \rangle$ for the prisoner's dilemma, where *(i)* $s_0 = (\epsilon_1, \epsilon_2)$ is the initial state; *(ii)* the transition function τ is given as

- $\tau(s_0, (\mathbf{C}, \mathbf{C})) = (\beta, \beta)$
- $\tau(s_0, (\mathbf{C}, \mathbf{D})) = (\delta, \alpha)$
- $\tau(s_0, (\mathbf{D}, \mathbf{C})) = (\alpha, \delta)$
- $\tau(s_0, (\mathbf{D}, \mathbf{D})) = (\gamma, \gamma)$
- $\tau(s, (*, *)) = s$, for every s different from s_0

and *(iii)* $\pi(p_a) = \{s = (l_1, l_2) \mid l_a = p\}$. The ECGM $\mathcal{P}_{pd}$ is depicted in Fig. 1. Notice that, for technical purposes, we assume that at the end of the game the agents loop on their final state.

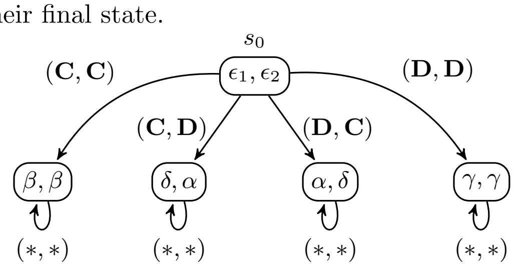

Fig. 1. The ECGM $\mathcal{P}_{pd}$ for the prisoner's dilemma.

3 Epistemic Strategy Logic

We now present an epistemic version of Strategy Logic, as a specification language for ECGM. First, for every set $A \subseteq Ag$ of agents, we introduce a set Var_A of strategy variables $x_A, x'_A, \ldots$.

Definition 3 (ESL). *For $p \in AP$ and $A \subseteq Ag$, the ESL formulas ϕ are defined in BNF as follows:*

$$\phi ::= p \mid \neg\phi \mid \phi \to \phi \mid X\phi \mid \phi U \phi \mid D_A\phi \mid \exists x_A \phi$$

The language ESL extends the Strategy Logic in [7] by adding the epistemic operator D_A for distributed knowledge in the set A of agents. Equivalently, ESL can be thought of as an epistemic extension of the Strategy Logic in [24], *modulo* the binding operator. The ESL formula $\exists x_A \phi$ is read as "the agents in A have a strategy to achieve ϕ". The interpretation of LTL operators X and U is standard. Observe that epistemic formulas $K_a\phi$ for "agent a knows ϕ", can be defined as $D_{\{a\}}\phi$. Similarly, we write $\exists x_a \phi$ for $\exists x_{\{a\}}\phi$, i.e., "agent a has a strategy to achieve ϕ". The other propositional connectives, LTL operators G and F, and the universal strategy quantifier $\forall$ can be defined as standard. Also, notice that the *nested-goal* fragment ESL[NG], the *boolean-goal* fragment ESL[BG], and the *one-goal* fragment ESL[1G] can be introduced in analogy to Strategy Logic [24]. Finally, the *free* variables $fr(\phi) \subseteq Ag$ of an ESL formula ϕ are inductively defined as follows:

$$\begin{array}{ll}
fr(p) & = \emptyset \\
fr(\neg\phi) = fr(D_A\phi) & = fr(\phi) \\
fr(\phi \to \phi') & = fr(\phi) \cup fr(\phi') \\
fr(X\phi) = fr(\phi U \phi') & = Ag \\
fr(\exists x_A \phi) & = fr(\phi) \setminus A
\end{array}$$

A *sentence* is an ESL formula ϕ with $fr(\phi) = \emptyset$.

We make use of ECGM to provide a semantics to ESL formulas, starting with the notion of strategy. We first consider the case of perfect information, as it provides us with a term of comparison with contexts of imperfect information. We will elaborate on their differences.

Definition 4 (Perfect Information Strategy). *Let γ be an ordinal s.t. $1 \leq \gamma \leq \omega$ and $A \subseteq Ag$ a set of agents. A γ-*recall (perfect information) *A*-strategy *is a function $f_A[\gamma] : \bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n \mapsto Act_A$ s.t. $f_A[\gamma](\kappa) \in d_A(last(\kappa))$ for every $\kappa \in \bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n$, where $last(\kappa)$ is the last element of κ.*

Thus, a γ-recall A-strategy returns an enabled A-action for every sequence of states of length at most γ (possibly infinite). For $A = \{a\}$, $f_A[\gamma]$ can be seen as a function from $\bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n$ to Act_a s.t. $f_A[\gamma](\kappa) \in Pr_a(last(\kappa)(a))$ for $\kappa \in \bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n$. One key feature of perfect information contexts – which fails for the imperfect information semantics provided below – is that, for every set $A = \{a_0, \ldots, a_m\}$ of agents, $f_A[\gamma]$ is tantamount to $f_{a_0}[\gamma] \times \ldots \times f_{a_m}[\gamma]$, where for every $\kappa \in \bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n$, $(f_{a_0}[\gamma] \times \ldots \times f_{a_m}[\gamma])(\kappa)$ is defined as the set of actions $\sigma \in Act_A$ s.t. $\sigma_a = f_a[\gamma](\kappa)$ if $a \in A$, and $\sigma_a = \sharp$ otherwise. Therefore, whenever we assume perfect information, a group's strategy is really the composition of its members' strategies; we will observe that this is not always the case for imperfect information. Also, the *outcome* of strategy $f_A[\gamma]$ at state s, or $out(s, f_A[\gamma])$, is the set of all s-runs λ s.t. $\lambda(i+1) \in out(\lambda(i), f_A[\gamma](\lambda[j,i]))$ for all $i \geq 0$ and $j = max(i - \gamma + 1, 0)$. With an abuse of notation we write $s' \in out(s, f_A[\gamma])$ to indicate that $s' \in \lambda(i)$ for some $\lambda \in out(s, f_A[\gamma])$ and $i \geq 0$. By varying γ we can define positional strategies ($\gamma = 1$), strategies with perfect recall ($\gamma = \omega$), etc. [12]. However, these different choices do not affect the following definitions and results, so we assume that γ is fixed and omit it, unless specified otherwise. Finally, given an Ag-strategy f and an A-strategy g, let f^A_g denote the Ag-strategy s.t. for every $\kappa \in \bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n$, $f^A_g(\kappa) = \sigma \in Act$, where $\sigma_A = g(\kappa)$ and $\sigma_{\bar{A}} = f(\kappa)$. Since $|out(s, f)| = 1$ for any Ag-strategy f, in the following we simply write $\lambda = out(s, f)$.

Definition 5 (Semantics of ESL). *We define whether an ECGM $\mathcal{P}$* satisfies *a formula φ at state s according to context $V_0, \ldots, V_n \subseteq \mathcal{S}$ and Ag-strategy f, or $(\mathcal{P}, s, \vec{V}, f) \models \varphi$, as follows (clauses for propositional connectives are straightforward and thus omitted):*

$(\mathcal{P}, s, \vec{V}, f) \models p$	*iff*	$s \in \pi(p)$
$(\mathcal{P}, s, \vec{V}, f) \models X\psi$	*iff*	*for* $\lambda = out(s, f)$, $(\mathcal{P}, \lambda(1), \vec{V}, f) \models \psi$
$(\mathcal{P}, s, \vec{V}, f) \models \psi U \psi'$	*iff*	*for* $\lambda = out(s, f)$ *and some* $k \geq 0$, $(\mathcal{P}, \lambda(k), \vec{V}, f) \models \psi'$, *and for every* j, $0 \leq j < k$ *implies* $(\mathcal{P}, \lambda(j), \vec{V}, f) \models \psi$
$(\mathcal{P}, s, \vec{V}, f) \models D_A\psi$	*iff*	*for all* $s' \in V_A$, $s' \sim_A s$ *implies* $(\mathcal{P}, s', \vec{V}, f) \models \psi$
$(\mathcal{P}, s, \vec{V}, f) \models \exists x_A \psi$	*iff*	*for some A-strategy* g, $(\mathcal{P}, s, \vec{V}, f^A_g) \models \psi$

where $V_A = \bigcap_{a \in A} V_a$.

An ESL formula φ is *true* at state s according to f, or $(\mathcal{P}, s, f) \models \varphi$, if $(\mathcal{P}, s, out(s, f_0), \ldots, out(s, f_n), f) \models \varphi$; it is *true* at state s, or $(\mathcal{P}, s) \models \varphi$, if $(\mathcal{P}, s, f) \models \varphi$ for all Ag-strategies f; it is *true* in $\mathcal{P}$, or $\mathcal{P} \models \varphi$, if $(\mathcal{P}, s_0) \models \varphi$ for all $s_0 \in I$; and it is *valid*, or $\models \varphi$, if φ is true in every ECGM. We remark that the satisfaction of formulas is independent from bound variables, that is, if for every $\kappa \in \bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n$, strategies f and f' return the same $fr(\phi)$-actions, i.e., $(f(\kappa))_{fr(\phi)} = (f'(\kappa))_{fr(\phi)}$, then $(\mathcal{P}, s, f) \models \phi$ iff $(\mathcal{P}, s, f') \models \phi$. In particular, the satisfaction of sentences is independent from strategies.

Hereafter we analyse the expressiveness of ESL. First of all, notice that the satisfaction of epistemic formulas depends on each set $V_a \subseteq \mathcal{S}$ of alternative states. Typically, in the literature each V_a is taken as the set $\mathcal{S}$ of accessible states [9, 22], as only these are considered as actual epistemic alternatives. An agent won't think possible to be situated in an unaccessible state $s \in \mathcal{G} \setminus \mathcal{S}$. However, by coherently developing this line of reasoning, in each state s only the accessible states in $out(s, f_a)$ may represent actual epistemic alternatives for agent a. Hence, by Def. 5 $(\mathcal{P}, s, f) \models D_A \psi$ iff for all accessible $s' \in \bigcap_{a \in A} out(s, f_a)$, $s' \sim_A s$ implies $(\mathcal{P}, s', out(s, f_0), \ldots, out(s, f_n), f) \models \psi$. That is, the notion of knowledge defined on ECGM can be sum up as indistinguishability given the current state of the system and strategies. We also remark that according to Def. 5, the operator D_A is indeed an $S5$ modality, and it satisfies all standard properties of distributed knowledge [9, 23]. Moreover, it shows some nice features when interacting with strategy modalities.

By the discussion above and definition of $f_{a_0} \times \ldots \times f_{a_m}$, we can derive that variables for group strategies are actually redundant in ESL with perfect information. Intuitively, there is a strategy for a group $A = \{a_0, \ldots, a_m\}$ of agents iff there are strategies for each agent $a \in A$. As a result, each formula $\exists x_A \phi$ (resp. $\forall x_A \phi$) quantifying on strategy variable x_A, can be rewritten as $\exists x_{a_0} \ldots \exists x_{a_m} \phi$ (resp. $\forall x_{a_0} \ldots \forall x_{a_m} \phi$). We state this fact formally.

Remark 1. For every ESL formula ϕ and $A = \{a_0, \ldots, a_m\}$,

$$\models \exists x_A \phi \leftrightarrow \exists x_{a_0} \ldots \exists x_{a_m} \phi \tag{1}$$

We will see that this represents a major difference w.r.t. contexts of imperfect information, where (1) does not normally hold.

Example 2. It is well-known that ATL modalities are expressible in (non-epistemic) Strategy Logic. Intuitively, an ATL formula $\langle\langle A \rangle\rangle \phi$ can be translated into ESL as $\exists x_A \forall x_{\bar{A}} \phi$, where, as discussed above, $\exists x_A$ is actually a shorthand for $\exists x_{a_0} \ldots \exists x_{a_m}$ whenever $A = \{a_0, \ldots, a_m\}$, and $\bar{A} = Ag \setminus A$.

As an example of the expressive power of ESL, we consider the following equivalences, which provide fixed-point characterisations of formulas $\theta_1 ::= \exists x_A \forall x_{\bar{A}} G \phi$, $\theta_2 ::= \exists x_A \forall x_{\bar{A}} F \phi$, and $\theta_3 ::= \exists x_A \forall x_{\bar{A}} (\phi U \phi')$, where ϕ and ϕ' are sentences:

$$\models \theta_1 \leftrightarrow \phi \wedge \exists x_A \forall x_{\bar{A}} X \theta_1 \tag{2}$$

$$\models \theta_2 \leftrightarrow \phi \vee \exists x_A \forall x_{\bar{A}} X \theta_2 \tag{3}$$

$$\models \theta_3 \leftrightarrow \phi' \vee (\phi \wedge \exists x_A \forall x_{\bar{A}} X \theta_3) \tag{4}$$

Formulas (2)-(4) are validities in the perfect information semantics. However, we will see that this is no longer the case in contexts of imperfect information.

Example 3. Strategy Logic is known to be expressive enough to specify Nash equilibria [7, 24]. Specifically, given the payoff ordering $\alpha > \beta > \gamma > \delta$ of the prisoner's dilemma in Example 1, we define the formula ψ_{NE} as follows:

$$\begin{aligned}\psi_{NE} ::= \textstyle\bigwedge_{a=1}^{2} & ((\exists y_a X\alpha_a \to X\alpha_a)\wedge \\ & (\neg\exists y_a X\alpha_a \wedge \exists y_a X\beta_a \to X\beta_a)\wedge \\ & (\neg\exists y_a X\alpha_a \wedge \neg\exists y_a X\beta_a \wedge \exists y_a X\gamma_a \to X\gamma_a)\wedge \\ & (\neg\exists y_a X\alpha_a \wedge \neg\exists y_a X\beta_a \wedge \neg\exists y_a X\gamma_a \wedge \exists y_a X\delta_a \to X\delta_a))\end{aligned}$$

where by the definition of ECGM $\mathcal{P}_{pd}$, each atom $p_a \in \{\alpha_a, \beta_a, \gamma_a, \delta_a\}$ is true at state s iff the local state of agent a is equal to payoff p. Also, $fr(\psi_{NE}) = \{1, 2\}$.

Then, we can check that for the prisoner's dilemma ECGM $\mathcal{P}_{pd}$, $(\mathcal{P}_{pd}, s_0, f) \models \psi_{NE}$ iff the strategy profile $f(s_0)$ is a Nash equilibrium. In general, given a game in normal form with n players and payoff ordering $\alpha^1 > \ldots > \alpha^k$, we define the SL formula ψ_{NE} characterizing Nash equilibria as follows:

$$\psi_{NE} ::= \bigwedge_{a=1}^{n} \bigwedge_{i=1}^{k} \left(\left(\bigwedge_{j=1}^{i-1} \neg\exists y_a X\alpha_a^j \right) \wedge \exists y_a X\alpha_a^i \to X\alpha_a^i \right)$$

Furthermore, the SL formula $\exists x_{Ag}\psi_{NE}$ expresses the existence of Nash equilibria. In particular, we can check that for $a \in \{1, 2\}$, $\mathcal{P}_{pd} \models K_a\exists x_1\exists x_2\psi_{NE}$, that is, each player knows that the prisoner's dilemma admits Nash equilibria. More specifically, for $a \in \{1, 2\}$, $\mathcal{P}_{pd} \models \exists x_1\exists x_2 K_a\psi_{NE}$, i.e., every agent knows what the Nash equilibrium actually is for the game. By a closer comparison of ESL formulas $K_a\exists x_{Ag}\psi_{NE}$ and $\exists x_{Ag}K_a\psi_{NE}$ we can see that the former expresses *de dicto* knowledge of strategies, while the latter asserts knowledge *de re*: the players not only know that there is some strategy that is a Nash equilibrium, but they are actually capable of pointing it out. This is indeed the case as we assume that our players are perfect rational reasoners, and the prisoner's dilemma is a game of perfect information. In the following section we will see that in contexts of imperfect information it is not always possible to infer knowledge *de re* from *de dicto* knowledge.

4 ESL with Imperfect Information

In Section 3 we observed that the notion of satisfaction provided in Def. 5 assumes that agents have perfect information about the state they are situated in. Specifically, their strategies are determined by the system's global state. However, in many use cases of interest agents might have access to only a partial view of the system, as represented in their local state. Along this line, we introduce an imperfect information semantics for ESL, starting with some formal definitions. First, two *histories* $\kappa, \kappa' \in \mathcal{S}^n$ are *indistinguishable* for a set A of agents, or $\kappa \sim_A \kappa'$, iff for every $m \leq n$, $\kappa(m) \sim_A \kappa'(m)$.

We adapt ideas in [18] to define a notion of strategy suitable for contexts of imperfect information.

Definition 6 (Uniform Strategies). *Let $A \subseteq Ag$ be a set of agents. A* γ-recall uniform A-strategy *is a γ-recall A-strategy $f^u_A[\gamma] : \bigcup_{1 \leq n \leq \gamma} \mathcal{S}^n \mapsto Act_A$ s.t. for all histories κ and κ' in $\mathcal{S}^n$, if $\kappa \sim_A \kappa'$ then $f^u_A[\gamma](\kappa) = f^u_A[\gamma](\kappa')$.*

In the case of a single agent a, a γ-recall uniform a-strategy can be seen as a mapping $f^u_a[\gamma] : \bigcup_{1 \leq n \leq \gamma} L^n_a \mapsto Act_a$ s.t. $f^u_a[\gamma](\kappa) \in Pr_a(last(\kappa))$ for every $\kappa \in \bigcup_{1 \leq n \leq \gamma} L^n_a$. That is, uniform strategies are functions from sequences of local states, rather than global states. This accounts for the "locality" of strategies.

We can now introduce a corresponding notion of satisfaction. In the literature on logics of strategic abilities we can identify two main variants of the satisfaction relation, depending on whether a strategy has to be successful for all undistinguishable states. As regards ATL, [28] assumes this to be the case, while [2] relaxes this constraint. Both alternatives are analysed in [18]. Hereafter we follow [2] and define a satisfaction relation $\models^i$ for imperfect information. Then we show how the relation $\models^i_S$, formalising the account put forward in [28], can be represented in terms of $\models^i$. Formally, for an ECGM $\mathcal{P}$, an ESL formula φ, a state s, a context $\vec{V} \subseteq \mathcal{S}$, and a uniform Ag-strategy f, all clauses for the relation $\models^i$ (resp. $\models^i_S$) are the same as in Def. 5, but the interpretation of $\exists$-formulas is restricted to uniform strategies as follows:

$$
\begin{array}{lll}
(\mathcal{P}, s, \vec{V}, f) \models^i \exists x_A \phi & \text{iff} & \text{for some uniform } A\text{-strategy } g,\ (\mathcal{P}, s, \vec{V}, f^A_g) \models^i \phi \\
(\mathcal{P}, s, \vec{V}, f) \models^i_S \exists x_A \phi & \text{iff} & \text{for some uniform } A\text{-strategy } g, \text{ for all } s' \in \vec{V}, \\
 & & s' \sim_A s \text{ implies } (\mathcal{P}, s', \vec{V}, f^A_g) \models^i_S \phi
\end{array}
$$

It is apparent that the definition of $\models^i_S$ has an epistemic flavour. Indeed, it has been shown that knowledge operators are definable in terms of ATL modalities when the satisfaction relation $\models^i_S$ is considered [5].

Next, we state a result on the relationship between these two notions of satisfaction. First, consider the translation function $\rho : \text{ESL} \mapsto \text{ESL}$ defined as:

$$
\begin{array}{lll}
\rho(p) & = & p \\
\rho(\phi \star \phi') & = & \rho(\phi) \star \rho(\phi') \\
\rho(\flat \phi) & = & \flat \rho(\phi) \\
\rho(\exists x_A \phi) & = & \exists x_A D_A \rho(\phi)
\end{array}
$$

where $\star$ is any binary operator, while $\flat$ is any unary operator different from $\exists$. Then, we can check that the two notions of satisfaction for imperfect information are related as described by the following lemma.

Lemma 1. *For every ESL formula ϕ,*

$$(\mathcal{P}, s, f) \models^i_S \phi \textit{ iff } (\mathcal{P}, s, f) \models^i \rho(\phi)$$

Sketch of Proof. The proof is by induction on the structure of ϕ. The case of interest is for existential formulas. For $\phi \equiv \exists x_A \psi$, $(\mathcal{P}, s, \vec{V}, f) \models^i_S \phi$ iff for some uniform A-strategy g, for all $s' \in \bigcap_{a \in A} V_a$, $s' \sim_A s$ implies $(\mathcal{P}, s', \vec{V}, f^A_g) \models^i_S \psi$. By induction hypothesis, $(\mathcal{P}, s', \vec{V}, f^A_g) \models^i \rho(\psi)$. Further, for all $s' \in \bigcap_{a \in A} V_a$, $s' \sim_A s$ implies $(\mathcal{P}, s', \vec{V}, f^A_g) \models^i \rho(\psi)$ iff $(\mathcal{P}, s, \vec{V}, f^A_g) \models^i$

$D_A\rho(\psi)$, iff $(\mathcal{P}, s, \vec{V}, f) \models^i \exists x_A D_A \rho(\psi)$, i.e., $(\mathcal{P}, s, \vec{V}, f) \models^i \rho(\phi)$. In particular, the above holds for $\vec{V} = out(s, f_0), \ldots, out(s, f_n)$. □

As a result, the strategy operator $\exists$ interpreted according to $\models^i_S$ can be simulated in $\models^i$ by using epistemic modalities in a *de re* manner. Lemma 1 provides one more example of the expressiveness of Epistemic Strategy Logic.

Another noteworthy result is that formulas (2)-(4) are no longer valid in the semantics of imperfect information.

Lemma 2. *For every ESL formula* ϕ,

$$\not\models^i \phi \wedge \exists x_A \forall x_{\bar{A}} X\theta_1 \to \theta_1 \tag{5}$$

$$\not\models^i \phi \vee \exists x_A \forall x_{\bar{A}} X\theta_2 \to \theta_2 \tag{6}$$

$$\not\models^i \phi' \vee (\phi \wedge \exists x_A \forall x_{\bar{A}} X\theta_3) \to \theta_3 \tag{7}$$

Sketch of Proof. We illustrate the case for (6) by considering the ECGM $\mathcal{Q}$ depicted in Fig. 2.

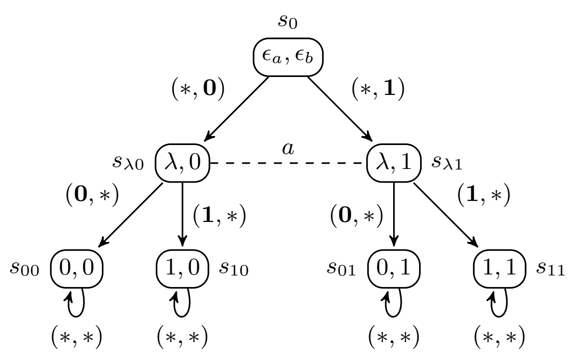

Fig. 2. The ECGM $\mathcal{Q}$.

In this game player b chooses secretly between values 0 and 1. Then, at the successive stage, a also chooses between 0 and 1. The game is won by a if the values provided by the two players coincide, otherwise b wins. To model this game as an ECGM we introduce agents $a = \langle L_a, Act_a, Pr_a \rangle$ and $b = \langle L_b, Act_b, Pr_b \rangle$ defined as

- $L_a = \{\epsilon_a, \lambda, 0, 1\}$ and $L_b = \{\epsilon_b, 0, 1\}$;
- $Act_a = Act_b = \{\mathbf{0}, \mathbf{1}, *\}$;
- $Pr_a(\epsilon_a) = Pr_a(0) = Pr_a(1) = \{*\}$ and $Pr_a(\lambda) = \{\mathbf{0}, \mathbf{1}\}$; while $Pr_b(\epsilon_b) = \{\mathbf{0}, \mathbf{1}\}$, $Pr_b(0) = Pr_b(1) = \{*\}$.

Also, we consider the atomic proposition eq, which intuitively expresses that the two bits in the global state are defined and equal. We then introduce the ECGM $\mathcal{Q} = \langle \{a, b\}, \{s_0\}, \tau, \pi \rangle$ as

- $s_0 = (\epsilon_a, \epsilon_b)$ is the only initial state;
- the transition function τ is given as follows for $i, j \in \{0, 1\}$:

- $\tau((\epsilon_a, \epsilon_b), (*, \mathbf{i})) = (\lambda, i)$
- $\tau((\lambda, i), (\mathbf{j}, *)) = (j, i)$
- $\tau((i, j), (*, *)) = (i, j)$

– $\pi(eq) = \{s_{00}, s_{11}\}$.

In Fig. 2 a dashed line stands for epistemic indistinguishability.

Now we check that $\mathcal{Q} \not\models (3)$. Indeed, it is the case that $(\mathcal{Q}, s_{\lambda 0}, out(s_0, \vec{f}), f) \models^i \exists x_a \forall x_b F eq$ whenever agent a does $\mathbf{0}$ at local state λ, or $g_a(\lambda) = \mathbf{0}$. Also, $(\mathcal{Q}, s_{\lambda 1}, out(s_0, \vec{f}), f) \models^i \exists x_a \forall x_b F eq$ whenever a does $\mathbf{1}$ at λ. Hence, $(\mathcal{Q}, s_0, f) \models^i \exists x_a \forall x_b X(\exists x_a \forall x_b F eq)$. However, it is not the case that $(\mathcal{Q}, s_0, f) \models^i \exists x_a \forall x_b F eq$ if we only consider uniform strategies, as agent a should perform different actions in states $s_{\lambda 0}$ and $s_{\lambda 1}$ that she is not able to distinguish. $\square$

Furthermore, we remark that in each state $s \in \{s_{\lambda 0}, s_{\lambda 1}\}$ agent a has *de dicto* knowledge of a uniform strategy to achieve $F eq$, that is, $(\mathcal{Q}, s, out(s_0, \vec{f}), f) \models^i K_a \exists x_a \forall x_b F eq$. However, she has no *de re* knowledge of such strategy, as it is the case that $(\mathcal{Q}, s, out(s_0, \vec{f}), f) \not\models^i \exists x_a K_a \forall x_b F eq$. Hence, in marked contrast with perfect information, in contexts of imperfect information *de dicto* knowledge of strategies does not imply knowledge *de re* in general.

Now observe that for the satisfaction relation $\models^i_S$ the converse of (2)-(4) fail.

Lemma 3. *For every ESL formula ϕ,*

$$\not\models^i_S \theta_1 \to \phi \wedge \exists x_A \forall x_{\bar{A}} X \theta_1 \tag{8}$$
$$\not\models^i_S \theta_2 \to \phi \vee \exists x_A \forall x_{\bar{A}} X \theta_2 \tag{9}$$
$$\not\models^i_S \theta_3 \to \phi' \vee (\phi \wedge \exists x_A \forall x_{\bar{A}} X \theta_3) \tag{10}$$

Sketch of Proof. We provide a witness for (9) by means of the ECGM $\mathcal{Q}'$ in Fig. 3. Formally, we introduce agents $a' = \langle L_{a'}, Act_{a'}, Pr_{a'} \rangle$ and $b' =$

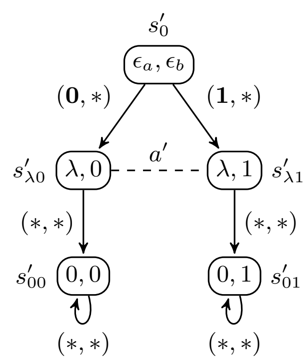

Fig. 3. The ECGM $\mathcal{Q}'$.

$\langle L_{b'}, Act_{b'}, Pr_{b'} \rangle$ s.t. *(i)* $L_{a'} = L_a$ and $L_{b'} = L_b$; *(ii)* $Act_{a'} = Act_a$ and $Act_{b'} = \{*\}$; and *(iii)* $Pr_{a'}(\epsilon_a) = \{\mathbf{0}, \mathbf{1}\}$, and $Pr_{a'}(\lambda) = Pr_{a'}(0) = Pr_{a'}(1) = \{*\}$; while $Pr_{b'}(\epsilon_b) = Pr_{b'}(0) = Pr_{b'}(1) = \{*\}$. Then, the ECGM $\mathcal{Q}' = \langle \{a', b'\}, \{s'_0\}, \tau', \pi' \rangle$ is specified as *(i)* $s'_0 = s_0 = (\epsilon_a, \epsilon_b)$; *(ii)* the transition function τ' is defined as in Fig. 3; and *(iii)* $\pi'(eq) = \{s'_{00}\}$. Also in Fig. 3 a dashed line denotes epistemic indistinguishability.

We show that $\mathcal{Q}' \not\models (3)$. First, notice that $(\mathcal{Q}', s'_0, out(s_0, \vec{f}), f) \models^i_S \exists x_{a'} \forall x_{b'} Feq$ by using the uniform strategy $g_{a'}$ s.t. $g_{a'}(s'_0) = \mathbf{0}$ and $g_{a'}(s'_{\lambda 0}) = *$. However, $(\mathcal{Q}', s'_{\lambda 0}, out(s_0, \vec{f}), f) \not\models^i_S \exists x_{a'} \forall x_{b'} Feq$, as there is no uniform a'-strategy g' s.t. $(\mathcal{Q}', s'_{\lambda 1}, out(s_0, \vec{f}), f^{a'}_{g'}) \models^i_S \forall x_{b'} Feq$. Hence, $(\mathcal{Q}', s'_0, f) \not\models^i_S \exists x_{a'} \forall x_{b'} X\theta_2$. $\square$

We now show that, differently from perfect information and Remark 1, in contexts of imperfect information an ESL formula $\exists x_A \phi$, for $A = \{a_0, \ldots, a_m\}$, is not equivalent to $\exists a_0 \ldots \exists a_m \phi$. In particular, while $\exists a_0 \ldots \exists a_m \phi \rightarrow \exists x_A \phi$ is valid, as the existence of a uniform strategy for each agent $a \in A$ implies the existence of a uniform strategy for group A, the converse does not hold. To see this, consider the modification $\mathcal{Q}''$ of the ECGM $\mathcal{Q}$ above, depicted in Fig. 4, which includes also an agent c, who has no strategic power but complete information on the current state[1].

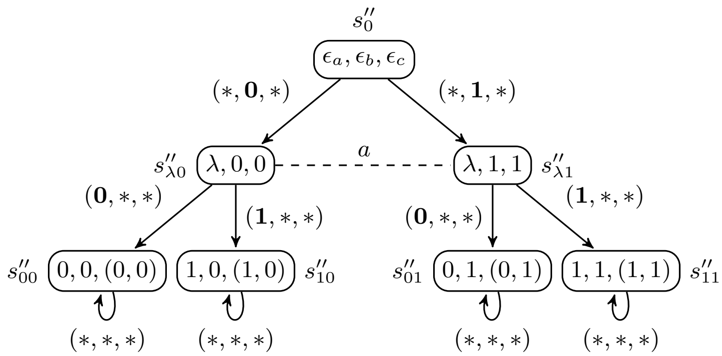

Fig. 4. The ECGM $\mathcal{Q}''$.

Clearly, in s''_0 the group $A = \{a, c\}$ of agents has a uniform A-strategy g to achieve Feq, that is, $(\mathcal{Q}, s_0, f) \models^i \exists x_A \forall x_b Feq$ for $g(s_0) = (*, \sharp, *)$, $g(s_{\lambda 0}) = (\mathbf{0}, \sharp, *)$, and $g(s_{\lambda 1}) = (\mathbf{1}, \sharp, *)$. However, agent a alone has no uniform strategy to achieve $\exists x_c \forall x_b Feq$ in s''_0, as states $s''_{\lambda 0}$ and $s''_{\lambda 1}$ are still undistinguishable to a. Hence, $(\mathcal{Q}, s_0, f) \not\models^i \exists x_a \exists x_c \forall x_b Feq)$. As a consequence, variables for group strategies cannot be dispensed with in contexts of imperfect information.

Finally, we show the main result of this paper, namely the fixed-point characterisations (2)-(4) of strategy operators can be partially recovered in ESL. Specifically, this result is achieved by combining epistemic and strategy modalities as stated in the following theorem.

Theorem 1. *Recall that* $\theta_1 ::= \exists x_A \forall x_{\bar{A}} G\phi$*,* $\theta_2 ::= \exists x_A \forall x_{\bar{A}} F\phi$*, and* $\theta_3 ::= \exists x_A \forall x_{\bar{A}} (\phi U \phi')$*. If we consider positional strategies (*$\gamma = 1$*), then*

$$\models^i \theta_1 \leftrightarrow \phi \wedge \exists x_A \forall x_{\bar{A}} X(\exists x_A \forall x_{\bar{A}} (G\phi \wedge D_A(\theta_1 \rightarrow G\phi))) \quad (11)$$

$$\models^i \theta_2 \leftrightarrow \phi \vee \exists x_A \forall x_{\bar{A}} X(\exists x_A \forall x_{\bar{A}} (F\phi \wedge D_A(\theta_2 \rightarrow F\phi))) \quad (12)$$

$$\models^i \theta_3 \leftrightarrow \phi' \vee (\phi \wedge \exists x_A \forall x_{\bar{A}} X(\exists x_A \forall x_{\bar{A}} (\phi U \phi' \wedge D_A(\theta_3 \rightarrow \phi U \phi'))) \quad (13)$$

[1] We thank an anonymous reviewer at ECAI14 for pointing out this model.

Sketch of Proof. We briefly show (12), the proof for the other cases is similar.

$\Rightarrow$ We can prove that $(\mathcal{P}, s, f) \models \theta_2 \to (\neg\phi \to \exists x_A \forall x_{\bar{A}} X \theta_2)$ as in Example 2. Further, $(\mathcal{P}, s, f) \models \exists x_A \forall x_{\bar{A}} D_A(\theta_2 \to F\phi)$ as we are using uniform strategies. In particular, we can choose the same strategies to witness the existential quantifiers in the two formulas.

$\Leftarrow$ Suppose that $(\mathcal{P}, s, f) \models \exists x_A \forall x_{\bar{A}} X(\exists x_A \forall x_{\bar{A}}(F\phi \wedge D_A(\theta_2 \to F\phi)))$, that is, for some uniform A-strategy g, for all uniform $\bar{A}$-strategy h and $\lambda = out(s, f^{A,\bar{A}}_{g,h}))$, $(\mathcal{P}, \lambda(1), out(s, \vec{f}), f^{A,\bar{A}}_{g,h}) \models \exists x_A \forall x_{\bar{A}}(F\phi \wedge D_A(\theta_2 \to F\phi))$, i.e., for some uniform A-strategy g', for all uniform $\bar{A}$-strategy h' and $\lambda' = out(\lambda(1), f^{A,\bar{A}}_{g',h'}))$, $(\mathcal{P}, \lambda'(1), out(s, \vec{f}), f^{A,\bar{A}}_{g',h'}) \models F\phi \wedge D_A(\theta_2 \to F\phi)$. Now we have to show that A-strategies g and g' can be composed so as to obtain a uniform A-strategy g'' witnessing $(\mathcal{P}, s, f) \models \exists x_A \forall x_{\bar{A}} F\phi$. To this end, we notice that $(\mathcal{P}, \lambda'(1), out(s, \vec{f}), f^{A,\bar{A}}_{g',h'}) \models D_A(\theta_2 \to F\phi)$, that is, for all $s' \in \bigcap_{a\in A} out(s, f_a)$, $s' \sim_A \lambda'(1)$ implies $(\mathcal{P}, s', out(s, \vec{f}), f^{A,\bar{A}}_{g',h'}) \models \exists x_A \forall x_{\bar{A}} F\phi \to F\phi$, i.e., the strategy $f^{A,\bar{A}}_{g',h'}$ to achieve $F\phi$ does not depend on the particular $s' \in \bigcap_{a\in A} out(s, f_a)$ s.t. $s' \sim_A \lambda'(1)$ (since $f^{A,\bar{A}}_{g',h'}$ is positional). As a consequence, strategies g and g' can actually be composed into g'' similarly as in Example 2, thus obtaining the desired result. $\square$

We remark that the conjunct $D_A(\theta_a \to \Box(\phi, \phi'))$ in the RHS of (11)-(13), for $\Box \in \{G, F, U\}$, guarantees that we can derive *de re* knowledge of strategies from *de dicto* knowledge, which in turn ensures the existence of uniform strategies; hence ruling out counterexemples such as the ECGM $\mathcal{Q}$ in Fig. 2. Observe that (11)-(13) are not really fixed-points as θ-formulas appear negatively on the RHS. Nonetheless, the interest of validities such as (11)-(13) lies in the fact that similar characterisations feature prominently in decision procedures for the satisfiability and model checking problems of temporal logics [11, 21]. We envisage that (11)-(13) may play a similar role w.r.t ESL with imperfect information. From a more philosophical perspective, such equivalences shed light on the meaning of strategies operators. For instance, according to (12), a group A of agents has a (uniform) strategy to eventually achieve ϕ (irrespectively of the other agents' behaviour) iff either ϕ holds or at the next step they have a strategy g to eventually achieve ϕ, and know that if they can eventually achieve ϕ, then they can do so by playing g.

As a special case, for a single agent a we have the following equivalences, where $\theta_4 ::= \exists x_a \forall x_{\bar{a}} G\phi$, $\theta_5 ::= \exists x_a \forall x_{\bar{a}} F\phi$, and $\theta_6 ::= \exists x_a \forall x_{\bar{a}}(\phi U \phi')$ for $\bar{a} = Ag \setminus \{a\}$, and distributed knowledge is replaced with individual knowledge.

$$\models^i \theta_4 \leftrightarrow \phi \wedge \exists x_a \forall x_{\bar{a}} X(\exists x_a \forall x_{\bar{a}}(G\phi \wedge K_a(\theta_4 \to G\phi)))$$
$$\models^i \theta_5 \leftrightarrow \phi \vee \exists x_a \forall x_{\bar{a}} X(\exists x_a \forall x_{\bar{a}}(F\phi \wedge K_a(\theta_5 \to F\phi)))$$
$$\models^i \theta_6 \leftrightarrow \phi' \vee (\phi \wedge \exists x_a \forall x_{\bar{a}} X(\exists x_a \forall x_{\bar{a}}(\phi U \phi' \wedge K_a(\theta_6 \to \phi U \phi')))$$

We conclude by briefly discussing the modifications necessary to extend Theorem 1 to strategies defined on arbitrary γ. Without going into details, we just remark that by defining the relation of indistinguishability on histories, it is possible to introduce a perfect recall semantics for ESL, where the relation of satisfaction is defined on histories, rather than single states. Given this perfect recall version of the imperfect information semantics for ESL, we can prove Theorem 1 for arbitrary γ.

5 Conclusions

In this paper we analysed Epistemic Strategy Logic, an extension of Strategy Logic [7, 24] with epistemic operators for individuals and coalitions, in contexts of imperfect information. We considered uniform strategies within the semantics of epistemic concurrent game models (ECGM), and investigated the modelling and expressive power of both ECGM and ESL. Specifically, we showed that a rich class of games can be modelled as ECGM, while the existence of Nash equilibria and knowledge thereof can be specified in ESL. We discussed some semantical features of Epistemic Strategy Logic. We showed that imperfect information allows to distinguish between *de dicto* and *de re* knowledge of strategies, while these are equivalent in contexts of perfect information. Notably, we proved that characterisations of ATL operators – even though not fixed-points – can be recovered in contexts of imperfect information through the interplay of epistemic and strategy operators. These results might point to future developments w.r.t. the satisfiability and model checking problems.

A number of extensions of the proposed framework are envisaged. Firstly, the nested-goal, boolean-goal, and one-goal fragment of SL are known to have better computational properties than full Strategy Logic [24]. It is likely that the corresponding fragments of ESL enjoy similar qualities. Secondly, ESL can be extended with further modalities for group knowledge, e.g. common knowledge, for enhanced expressiveness. Thirdly, various assumptions can be imposed on ECGM, for instance perfect recall, no learning, and synchronicity. These extensions, while enhancing the modelling power of ECGM, are also likely to increase their computational complexity.

Acknowledgements: we would like to thank the LAMAS reviewers for useful comments on a previous version of the paper.

References

1. Agotnes, T., Goranko, V., Jamroga, W.: Alternating-time temporal logics with irrevocable strategies. In: Proceedings of TARK '07, ACM, (2007)
2. Alur, R., Henzinger, T., Kupferman, O.: Alternating-time temporal logic. J. ACM 49(5), 672–713 (2002)
3. Belardinelli, F.: Reasoning about knowledge and strategies: Epistemic strategy logic. In Proceedings of SR '14 (2014)
4. Bulling, N., Jamroga, W.: Comparing variants of strategic ability. Autonomous Agents and Multi-Agent Systems 28(3), 474–518 (2014)
5. Bulling, N., Dix, J., Jamroga, W.: Model checking logics of strategic ability. In Specification and Verification of Multi-agent Systems, pp. 125–159. Springer (2010)

6. Cermák, P., Lomuscio, A., Mogavero, F., Murano, A.: MCMAS-SLK: A model checker for the verification of strategy logic specifications. In Proceedings of CAV '14, pp. 525–532. Springer (2014).
7. Chatterjee, K., Henzinger, T., Piterman, N.: Strategy logic. Inf. Comput. 208(6), 677–693 (2010)
8. Easley, D., Kleinberg, J.: Networks, Crowds, and Markets. CUP (2010)
9. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning About Knowledge. The MIT Press (1995)
10. Gammie, P., van der Meyden, R.: Mck: Model checking the logic of knowledge. In Proceedings of CAV '04, pp. 479–483. Springer (2004)
11. Goranko, V., Shkatov, D.: Tableau-based decision procedures for logics of strategic ability in multiagent systems. ACM Trans. Comput. Log. 11(1) (2009)
12. Goranko, V., Jamroga, W.: Comparing semantics of logics for multi-agent systems. Synthese 139(2), 241–280 (2004)
13. Halpern, J.Y., Vardi, M.Y.: The complexity of reasoning about knowledge and time: Extended abstract. In: Hartmanis, J. (ed.) STOC. pp. 304–315. ACM (1986)
14. Halpern, J.Y., Vardi, M.Y.: The complexity of reasoning about knowledge and time. i. lower bounds. J. Comput. Syst. Sci. 38(1), 195–237 (1989)
15. van der Hoek, W., Wooldridge, M.: Cooperation, knowledge, and time. Studia Logica 75(1), 125–157 (2003)
16. Huang, X., Meyden, R.: An epistemic strategy logic. In Proc. of SR '14 (2014)
17. Jamroga, W., Ågotnes, T.: Constructive knowledge. Journal of Applied Non-Classical Logics 17(4), 423–475 (2007)
18. Jamroga, W., van der Hoek, W.: Agents that know how to play. Fundam. Inform. 63(2-3), 185–219 (2004)
19. Jonker, G.: Feasible strategies in Alternating-time Temporal Epistemic Logic. Master's thesis, University of Utrecht (2003)
20. Kacprzak, M., Nabialek, W., Niewiadomski, A., Penczek, W., Pólrola, A., Szreter, M., Wozna, B., Zbrzezny, A.: Verics 2007 - a model checker for knowledge and real-time. Fundamenta Informaticae 85(1-4), 313–328 (2008)
21. Kupferman, O., Vardi, M.Y., Wolper, P.: An Automata-Theoretic Approach to Branching-Time Model Checking. Journal of the ACM 47(2), 312–360 (2000)
22. Lomuscio, A., Qu, H., Raimondi, F.: MCMAS: A model checker for the verification of multi-agent systems. In Proceedings of CAV '09, pp. 682–688. Springer (2009)
23. Meyer, J., Hoek, W.: Epistemic Logic for AI and Computer Science. CUP (1995)
24. Mogavero, F., Murano, A., Perelli, G., Vardi, M.: Reasoning about strategies. ACM Trans. Comput. Log. 15(4), 34 (2014).
25. Mogavero, F., Murano, A., Sauro, L.: On the boundary of behavioral strategies. In Proceedings of LICS '13, pp. 263–272 (2013),
26. Mogavero, F., Murano, A., Sauro, L.: Strategy games: a renewed framework. In Proceedings of AAMAS '14, pp. 869–876. IFAAMAS/ACM (2014).
27. Pauly, M.: A modal logic for coalitional power in games. J. Log. Comput. 12(1), 149–166 (2002)
28. Schobbens, P.Y.: Alternating-time logic with imperfect recall. Electr. Notes Theor. Comput. Sci. 85(2), 82–93 (2004)