An Epistemic Logic for Reasoning about Strategies in General Auctions*

Munyque Mittelmann and Laurent Perrussel

Université de Toulouse - IRIT, France {munyque.mittelmann, laurent.perrussel}@irit.fr

Abstract. In this paper, we present the Epistemic Auction Description Language (E-ADL), a language for epistemic and strategic reasoning in auctions from the player's perspective. An automated auction player faces the challenge of understanding and processing several different auction-based markets. With E-ADL, an agent can evaluate the mechanism with well-known properties of economic theory, such as strategy-proofness and efficiency. Moreover, with the epistemic component, the agent can reason about other agents' private valuation and their awareness of the protocol properties.

Keywords: Logics for Multi-agents · Game Description Language · General Game Playing · Auction-based Markets.

1 Introduction

Auctions are well-defined environments that provide a valuable testing-ground for economic theory. They are important for understanding methods of price formation and negotiations in which both buyers and sellers are actively involved in determining the price [9]. Typically, an auction-based market is described by a set of rules stating how the participants bid, how the winners are determined, and what should be their payment. Any autonomous auction agent will face the challenge of understanding and processing a number of different auction-based markets. There are variants that differ on the participants' type (e.g. only buyers, both buyers and sellers, ...), the kind and amount of goods been auctioned, the bidding protocol, and the allocation and payment rules [10].

The great variety of auction protocols prevents any autonomous agent to easily switch between different auction-based markets. Building intelligent agents that can switch between different auctions and process their rules is a key issue for designing automated auction-based market places. For this reason, we previously proposed a general language to describe auction-based markets from the auctioneer perspective [11]. *Auction Description Language* (ADL) is based on the Game Description Language (GDL), which is a logic programming language for representing and reasoning about game rules [6].

In the players' perspective, the agents may be able to process the protocol to collect the definition of its main components: the bid legality, the payment, and allocation rules. With respect to these components, the bidder evaluates the auction market: the impact of her participation (individual rationality), the objectives of the auctioneer (to maximize revenue or efficiency), and the possibility of manipulation (strategy-proofness). Finally, the bidder may use her knowledge about other agents' private valuations and awareness of the auction properties.

In this paper, we focus on the epistemic and strategic reasoning of such auction players. We extend ADL with knowledge operators from the Epistemic GDL [8] and the action modality from the GDL variant proposed in [18]. Our goal is to provide the ground for the design of General Auction Players: (i) such player should be able to evaluate the mechanism and its strategy-proofness dimension and (ii) if not, she should then consider her knowledge about other players in order to define her action.

1.1 Related Work

To the best of our knowledge, there is no contribution that focuses on the strategic dimension of auctions through a logical perspective. However, numerous contributions are defining logical systems for Strategic Reasoning. The Alternating time Temporal Logic (ATL) [1] provides a logic-based analysis of strategic decisions. For representing games, the Propositional Logic of Games (GPL) [13] specifies the effects of game playing by using inference mechanisms from propositional dynamic logic (PDL). A more practical approach to specify a game is to use the Game Description Language (GDL) [6]. The Auction Description Language (ADL) [11] extends GDL by handling numerical variables, a key feature for representing the allocation and payment rules.

To represent strategies, the Strategy Logic (SL) uses first-order quantifications over strategies in turn-based (i.e. *asynchronous*) games [2]. This approach cannot model the internal structures of strategies, which prevents to easily design strategies aiming to achieve a goal state. [14] introduces a logic for reasoning about composite strategies in extensive form turn-based games: strategies are treated as programs that are combined by PDL-like connectives to ensure an outcome. Zhang and Thielscher [19] present a variant of GDL to describe game strategies, where formulas can be understood as move recommendations for a player. However, their work can only model turn-based games.

To incorporate imperfect information games, GDL has been extended to GDL-II [16] and GDL-III [17]. As purely descriptive languages, GDL-II and GDL-III aim at describing the rules of an imperfect information game, but do not provide tools for reasoning about how a player infers information based on these rules. All these logics face decidability and tractability issues: their expressive power prevents them to be implemented in realistically in an artificial agent. Jiang et al. [8] propose an epistemic extension of GDL (EGDL) to represent and reason about imperfect information games. The language allows representing the rules of an imperfect information game.

Structure of the paper Due to the space limitation, we omitted the Propositions proofs. All the proofs are available at https://epistemicadl.page.link/EELP2020. The remainder of the paper proceeds as follows. In Section 2, we define the base terminology and describe the State Transition structures that are used to evaluate E-ADL semantics. In Section 3, we present the language syntax and semantics and illustrate our approach by describing and deriving properties about a standard type of auction, a First-Price Blind protocol. In Section 4, we define strategy rules, which are formulas in E-ADL assigning a unique action to be taken in each state. In Section 5, we show how to verify classic properties of mechanism design in E-ADL auctions, such as efficiency and strategy-proofness. Finally, Section 6 concludes the paper.

2 State Transition Auctions

In this section, we introduce a logical framework for reasoning in general auction protocols. The framework is based on ADL [11] and Epistemic GDL [7]. We call the framework *Epistemic Auction Description Language*, denoted E-ADL. We restrict our definition to single-unit and single-good auctions.

Definition 1. *An auction signature $\mathcal{S}$ is a tuple $(N, V, \mathcal{A}, \Phi, Y)$, where: (i) $N = \{1, 2, \cdots, n\}$ is a nonempty finite set of agents; (ii) $V \subset \mathbb{N}$ is a finite subset of natural numbers representing the range of valuations, bids and payments; (iii) $\mathcal{A} = \bigcup_{r\in N} A^r$, where $A^r = \{bid^r(x) : x \in V\}$ consists of a nonempty set of actions (or bids) performed by agent $r \in N$ and $A^r \cap A^i = \emptyset$ if $r \neq i$. For convenience, we occasionally write a^r for denoting an action in A^r; (iv) $\Phi = \{p, q, \cdots\}$ is a finite set of atomic propositions for specifying individual features of a state; and (v) $Y = \{y_1, y_2, \cdots\}$ is a finite set of numerical variables for specifying numerical features of a state.*

Through the rest of the paper, we will fix an auction signature $\mathcal{S}$ and all concepts will be based on this auction signature, except if stated otherwise. In this paper, we adopt a semantics based on state-transition models which is more suitable for describing the dynamics as the one based on stable models initially considered for GDL and General Game Playing [6].

Definition 2. *A state transition ST-model M is a tuple $(W, I, T, \{R_r\}_{r\in N}, L, U, \pi_\Phi, \pi_Y)$, where: (i) W is a finite nonempty set of states; (ii) $I \subseteq W$ is a set of initial states; (iii) $T \subseteq W \setminus I$ is a set of terminal states; (iv) $R_r \subseteq W \times W$ is an equivalence relation for agent r, indicating the states that are indistinguishable for r; (v) $L \subseteq W \times \mathcal{A}$ is a legality relation, describing the legal actions at each state; (vi) $U : W \times \prod_{r\in N} A^r \to W$ is an update function, specifying the transitions for each combination of joint actions; (vii) $\pi_\Phi : W \to 2^\Phi$ is the valuation function for the state propositions; and (viii) $\pi_Y : W \times Y \to V$ is the valuation function for the numerical variables.*

Given $d \in \prod_{r\in N} A^r$, let $d(r)$ be the individual action for agent r in the joint action d. Let $L(w) = \{a \in \mathcal{A} \mid (w, a) \in L\}$ be the set of all legal actions at state

w. Let $R_r(w)$ denote the set of all states that agent r cannot distinguish from w, i.e. $R_r(w) = \{u \in W : wR_ru\}$.

For any $w \in W$ and $d \in \prod_{r \in N} A^r$, we call (w, d) a *move*. It is a *legal move* if $(w, d(r)) \in L$, for all $r \in N$. We use the notation $(w, d(r), d(-r))$ when it is more convenient, where $d(-r) \in \prod_{i \neq r \in N} A^i$ denotes the actions of all agents except by r in the joint action d. Any set $S_r \subseteq \{(w, d) : d(r) \in L(w)$ & $w \in W$ & $d \in \prod_{i \in N} A^i\}$ of *moves* is a *strategy* of a player $r \in N$. Our notion of move and strategy is based on the asynchronous definition from [18] and [19].

Definition 3. *Two moves $(w, d), (w', d')$ are imperfect recall equivalent for agent r, written $(w, d) \approx_r (w', d')$, iff wR_rw' and $d(r) = d(r)'$.*

An agent with imperfect recall is only aware of the present state but forgets everything that happened [7]. We only consider imperfect recall because differently from the standard GDL, our semantics is based on moves instead of paths. This semantics allows the agent to reason about the effects of actions without exploring all ways the game could proceed (i.e. all the reachable states in each complete path where she takes this action). Since a GDL path is a sequence of states and (legal) joint actions, the set of complete paths for a model M can have exponential size. For a formal definition of reachable states and complete paths in GDL, please refer to [8]. In E-ADL, we define the action execution modality in synchronous games[1]. The idea of move-based semantics and action modalities comes from [18]. Their approach is restricted to synchronous games, where only one action can be performed at a given state.

Given an agent $r \in N$, a strategy S_r is *complete* if there is a move $(w, d) \in S_r$ unless $L(w) \cap A^r = \emptyset$, for each state $w \in W$. In other words, the strategy S_r always provides at least one action to be taken in any state, except if there is no legal action. A strategy S_r is *deterministic* if $(w, d(r), d(-r)) \in S_r$ and $(w, d(r)', d(-r)') \in S_r$, then $d(r) = d(r)'$, for any $w \in W$. Finally, a strategy S_r is *functional* if it is complete and deterministic. Intuitively, a functional strategy provides a unique action to be taken by r in any state.

A *run* in an E-ADL model is a finite sequence of legal moves $(w_0, d_0), (w_1, d_1), \cdots, (w_e, d_e)$, such that $w_0 \in I$ and $w_{i+1} = U(w_i, d_i)$, for any $0 \leq i < e$. Although the agents evaluate the semantics based on moves, a run represents an execution of an auction protocol.

3 Epistemic Auction Description Language

The *Epistemic Auction Description Language* (E-ADL) is a framework to allow epistemic reasoning for auction players. Let $z \in \mathcal{L}_z$ be a numerical term defined as follows: $z ::= z' \mid add(z, z) \mid sub(z, z) \mid min(z, z) \mid max(z, z) \mid times(z, z) \mid y \mid maxbid$, where $z' \in \mathbb{N}, y \in Y$, $a \in \mathcal{A}$.

The numerical terms $add(z_1, z_2)$, $sub(z_1, z_2)$ and $times(z_1,\ z_2)$ specify the value obtained by adding, subtracting and multiplying z_2 from z_1, respectively.

[1] Asynchronous games can be simulated as synchronous games in GDL by turn-based legality. An example of how to simulate asynchronous games is given at [8].

The terms $min(z_1, z_2)$ and $max(z_1, z_2)$ specify the minimum and maximum value between z_1 and z_2, resp. The numerical term $maxbid$ represents the highest bid in a move. Finally, y denotes the value of the variable $y \in Y$ in the current state.

The Epistemic Auction Description language is denoted by $\mathcal{L}_{E-ADL}$ and a formula φ in $\mathcal{L}_{E-ADL}$ is defined by the following BNF grammar:

$$\varphi ::= p \mid initial \mid terminal \mid legal(a^r) \mid does(a^r) \mid \mathsf{K}_r\varphi \mid \mathsf{C}\varphi \mid [\,a^r\,]\varphi \mid \\ \neg\varphi \mid \varphi \wedge \varphi \mid z < z \mid z > z \mid z = z \mid r \prec r$$

where $p \in \Phi$, $r \in N$, $a^r \in \mathcal{A}$ and $z \in \mathcal{L}_z$.

Other connectives $\vee, \rightarrow, \leftrightarrow, \top$ and $\bot$ are defined by $\neg$ and $\wedge$ in the standard way. The comparison operators $\leq$, $\geq$ and $\neq$ are defined by $\vee, >, <$ and $=$. The extension of the comparison operators $>, <, =$, $\leq$, $\geq$, $\neq$ and numerical terms $max(z_1, z_2), min(z_1, z_2)$, $add(z_1, z_2)$ to multiple arguments is straightforward.

We use set notation to compactly represent numerical terms with two or more arguments. For instance, $max(\{\vartheta_i : i \in N\})$ is a representation of the numerical term $max(\vartheta_1, \vartheta_2, \cdots, \vartheta_n)$, i.e. the maximum private value among the agents. Note that this notation also includes conditions over elements of a set. As an example, assume z_1, z_2, $bound$ and $maxset$ are natural numbers, then the formula $maxset = max(\{v_1, v_2 : v_1 < bound \;\&\; v_2 < bound\})$ is a compact representation of the E-ADL formula $(v_1 < bound \wedge v_2 < bound \wedge maxset = max(v_1, v_2)) \vee (v_1 < bound \wedge maxset = v_1) \vee (v_2 < bound \wedge maxset = v_2)$, i.e. $maxset$ is the maximum between two values complying with a logical condition.

Intuitively, $initial$ and $terminal$ specify the initial and the terminal states, respectively; $does(a^r)$ asserts that agent r takes action a^r at the current move; $legal(a^r)$ asserts that agent r is allowed to take action a^r the current state. Given an agent $r \in N$, we denote $does(bid^r(\vartheta_r))$ to represent that r did the action of bidding its own value. Similarly, $legal(bid^r(\vartheta_r))$ denotes that this action is legal. The epistemic operators K_r and C are taken from the Modal Epistemic Logic [5]. The formulas $\mathsf{K}_r\varphi$ and $\mathsf{C}\varphi$ are read as "agent r knows φ" and "φ is common knowledge among all the agents in N" (i.e. every agent knows φ, knows that every other agent knows φ, and so on), respectively. The action execution operator comes from the GDL variant with action modalities [18] and a formula $[\,a^r\,]\varphi$ means that if action a^r is executed at the current state, φ will be true in the next state. The formulas $z_1 > z_2$, $z_1 < z_2$, $z_1 = z_2$ means that a numerical term z_1 is greater, less and equal to a numerical term z_2, respectively. The tie-breaking priority is represented by the formula $r_1 \prec r_2$, i.e. agent r_1 precedes r_2 in the lexicographical order.

Instead of using the temporal operator $\bigcirc$ from GDL, we use the action modality to define an abbreviation with similar meaning [18]:

$$\bigcirc\varphi =_{def} \bigwedge_{r \in N} \bigvee_{a^r \in A^r} (does(a^r) \wedge [\,a^r\,]\varphi)$$

The formula $\bigcirc\varphi$ means "φ holds at the next state". We also use the following abbreviations from the Modal Epistemic Logic: $\widehat{\mathsf{K}}_r\varphi =_{def} \neg\mathsf{K}_r\neg\varphi$ and

$\mathsf{E}\varphi =_{def} \bigwedge_{r\in N} \mathsf{K}_r\varphi$, where $\widehat{\mathsf{K}}_r\varphi$ represents that "φ is compatible with agent r's knowledge". The formula $\mathsf{E}\varphi$ represents that "every agent in N knows φ".

3.1 Semantics

The semantics for the ADL language is given in two steps. First, we define function f to define the meaning of numerical terms $z \in \mathcal{L}_z$. Next, a formula $\varphi \in \mathcal{L}_{E-ADL}$ is interpreted with respect to a move.

Through the rest of this paper, the function $maximum(a, b, \cdots)$ returns the maximum value between a finite sequence $a, b, \cdots \in \mathbb{Z}$. Let $Y^+ = Y \cup \{maxbid\}$. Numerical terms $z \in \mathcal{L}_z \setminus Y^+$ have a constant evaluation, independently from a move. Their valuation can be simply assigned by function $f_{\mathbb{Z}}$ (Definition 4). In Definition 5, we specify the more general function f to evaluate any $z \in \mathcal{L}_z$.

Definition 4. *. Define Function $f_{\mathbb{Z}} : \mathcal{L}_z \setminus Y^+ \to \mathbb{Z}$, assigning any formula $z \in \mathcal{L}_z \setminus Y^+$ to a number in $\mathbb{Z}$:*

If z is in the form $add(z', z'')$, $sub(z', z'')$, $min(z', z'')$, $max(z', z'')$, $times(z', z'')$ or $mod(z')$, then $f_{\mathbb{Z}}(z)$ is defined through the application of the corresponding mathematical operators and functions. Otherwise, $f_{\mathbb{Z}}(z) = z$ if $z \in \mathbb{Z}$.

Definition 5. *Define Function $f : W \times \prod_{r\in N} A^r \times \mathcal{L}_z \to \mathbb{Z}$, assigning any state $w \in W$, joint action $d \in \prod_{r\in N} A^r$, and formula $z \in \mathcal{L}_z$ to a number in $\mathbb{Z}$:*

$$f(w,d,z) = \begin{cases} f_{\mathbb{Z}}(z) & \text{if } z \in \mathbb{Z} \setminus Y^+ \\ \pi_Y(w,z) & \text{if } z \in Y \\ maximum(\{x : d(r) = bid^r(x) \ \&\ r \in N\}) & \text{if } z = maxbid \end{cases}$$

Definition 6. *Let M be an ST-Model. Given a move (w, d), where $w \in W$ and $d \in \prod_{r\in N} A^r$, and a formula $\varphi \in \mathcal{L}_{ADL}$, we say φ is true (or satisfied) in the move (w, d) under M, denoted by $M \models_{(w,d)} \varphi$, according with the following:*

$$\begin{array}{lll}
M \models_{(w,d)} p & \textit{iff} & p \in \pi_{\Phi}(w) \\
M \models_{(w,d)} \neg\varphi & \textit{iff} & M \not\models_{(w,d)} \varphi \\
M \models_{(w,d)} \varphi_1 \wedge \varphi_2 & \textit{iff} & M \models_{(w,d)} \varphi_1 \text{ and } M \models_{(w,d)} \varphi_2 \\
M \models_{(w,d)} initial & \textit{iff} & w \in I \\
M \models_{(w,d)} terminal & \textit{iff} & w \in T \\
M \models_{(w,d)} r_1 \prec r_2 & \textit{iff} & r_1 \prec_{Lex} r_2 \\
M \models_{(w,d)} legal(a^r) & \textit{iff} & a^r \in L(w) \\
M \models_{(w,d)} does(a^r) & \textit{iff} & d(r) = a^r \\
M \models_{(w,d)} z_1 \ \mathsf{cp}\ z_2 & \textit{iff} & f(w,d,z_1) \ \mathsf{cp}\ f(w,d,z_2), \textit{where } \mathsf{cp} \in \{>,<,=\} \\
M \models_{(w,d)} \mathsf{K}_r\varphi & \textit{iff} & \textit{for any } w' \in W \ \&\ d' \in \prod_{i\in N} A^i, \textit{ if } (w,d) \approx_r (w',d'), \\
 & & \textit{then } M \models_{(w',d')} \varphi \\
M \models_{(w,d)} \mathsf{C}\varphi & \textit{iff} & \textit{for any } w' \in W \ \&\ d' \in \prod_{i\in N} A^i, \textit{ if } (w,d) \approx_N (w',d'), \\
 & & \textit{then } M \models_{(w',d')} \varphi
\end{array}$$

$$M \models_{(w,d)} [a^r]\varphi \text{ iff } M \models_{(U(w,d'),c)} \varphi, \text{ where } d' = (a^r, d(-r)), \text{ for all } c \in \prod_{i\in N} A^i$$

where $\approx_N$ *is the transitive closure of* $\bigcup_{r\in N} \approx_r$ *and* $\prec_{Lex}$ *is a relation denoting the lexicographical order among agents in* N.

In the semantics of $\mathsf{C}\varphi$, note $(w,d) \approx_N (w',d')$ represents the transitive closure of the equivalence relation between states as the joint actions d and d' are the same.

A formula φ is *globally true* in an ST-Model M, written $M \models \varphi$, if $M \models_{(w,d)} \varphi$ for all $w \in W$ and $d \in \prod_{r\in N} A^r$. Finally, let Σ be a set of formulas in $\mathcal{L}_{ADL}$, then M is a *model* of Σ if $M \models \varphi$ for all $\varphi \in \Sigma$.

Given an ST-model M, the Epistemic Properties (Prop. 1) express when a formula is globally known by one agent and when it is globally common knowledge.

Proposition 1. *Let* M *be an ST-Model,* $r \in N$ *be an agent and* $\varphi \in \mathcal{L}_{E-GDL}$ *be a formula, then*

1. $M \models \varphi \to \mathsf{K}_r\varphi$ *if and only if for all* $w, w' \in W$ *and all* $d, d' \in \prod_{i\in N} A^i$ *such that* $(w,d) \approx_r (w',d')$, $M \models_{(w,d)} \varphi$ *iff* $M \models_{(w',d')} \varphi$
2. $M \models \varphi \to \mathsf{C}\varphi$ *if and only if for all* $w, w' \in W$ *and all* $d, d' \in \prod_{i\in N} A^i$ *such that* $(w,d) \approx_N (w',d')$, $M \models_{(w,d)} \varphi$ *iff* $M \models_{(w',d')} \varphi$

3.2 Running Example: a First-Price Blind Auction

In this section, we illustrate how to describe an auction in E-ADL. First, we present the protocol from the auctioneer perspective. This protocol describes strictly the rules of the auction. Next, we present additional epistemic rules for allowing the agents' reasoning.

Auctioneer Perspective To describe a First-Price Blind Auction, we first define the auction signature, written $\mathcal{S}_{bli} = \{N, V, \mathcal{A}, \Phi_{bli}, Y_{bli}\}$, where $\Phi_{bli} = \{\}$, $Y_{bli} = \{payment_r, alloc_r, \vartheta_i : r \in N\}$. The numerical variables $payment_r$, $alloc_r$ and ϑ_r specify the payment, the allocation and the private value for an agent r. The rules of a First-Price Blind Auction are formulated by E-ADL-formulas as shown in Figure 1.

In an initial state, all agents have the payment and allocation equal to 0 (Rule 1). We are in a terminal state *iff* we are not in an initial state (Rule 2). Rule 3 specifies that it is legal, for all agents, to bid any value between 0 and V if we are not in the terminal state. In the terminal state, the agents can only bid 0. Rules 4 and 5 specify how the payment and allocation are updated. The formula $wins_{r,x}$ represents the condition of whether agent r bids x and for all other agent i, either i does not bid x or r wins the tie-breaking with i. If it is an initial state, $wins_{r,x}$ for an agent r and x is the highest bid, then in the next state she will get the good and pay her bid price. Otherwise, she will not get the good and the payment will be 0. Rule 6 states that after the terminal state, the numerical variables cannot change (self-loop). Let Σ_{bli} be the set of Rules 1-6.

Let $wins_{r,x} =_{def} does(bid^r(x)) \wedge \bigwedge_{i \neq r \in N} (\neg does(bid^i(x)) \vee r \prec i)$ and $r \in N$,
1. $initial \leftrightarrow \bigwedge_{i \in N} payment_i = 0 \wedge alloc_i = 0$
2. $terminal \leftrightarrow \neg initial$
3. $\bigwedge_{x \in V} (legal(bid^r(x)) \leftrightarrow \neg terminal \vee (terminal \wedge x = 0))$
4. $\bigvee_{x \in V} (initial \wedge maxbid = x \wedge wins_{r,x} \leftrightarrow \bigcirc(alloc_r = 1 \wedge payment_r = x))$
5. $\bigwedge_{x \in V} (initial \wedge \neg wins_{r,x} \rightarrow \bigcirc(alloc_r = 0 \wedge payment_r = 0))$
6. $\bigvee_{x \in V, y \in \{0,1\}} terminal \wedge alloc_r = y \wedge payment_r = x \rightarrow \bigcirc(alloc_r = y \wedge payment_r = x)$

Fig. 1. First-Price Blind Auction represented by Σ_{bli} (Auctioneer Protocol)

Agent's perspective Now let us focus on the agents' perspective. The following E-ADL rules says that each agent is aware of its own valuation, each agent has a private in V, and each agent knows her own action, respectively:

Let $r \in N$,
1. $\bigwedge_{x \in V} (\vartheta_r = x \rightarrow \mathsf{K}_r(\vartheta_r = x))$
2. $\bigvee_{x \in V} \vartheta_r = x$
3. $\bigwedge_{x \in V} (does(bid^r(x)) \rightarrow \mathsf{K}_r does(bid^r(x)))$

Fig. 2. First-Price Blind Auction represented by $\Sigma_{bli,N}$ (Agent's Protocol)

Let Σ_{bli} be the set of Rules 1-3. We assume that each agent knows the agent's protocol, i.e. $\mathsf{E}\Sigma_{bli,N}$ and the Auctioneer Protocol is common knowledge, i.e. $\mathsf{C}\Sigma_{bli}$.

Model Representation Given $\mathcal{S}_{bli} = (N, V, \mathcal{A}, \Phi_{bli}, Y_{bli})$, the state transition ST-model $M_{bli} = (W_{bli}, I_{bli}, T_{bli}, \{R_{r,bli}\}_{r \in N}, L_{bli}, U_{bli}, \pi_{\Phi,bli}, \pi_{Y,bli})$ is the model representation of the First-Price Blind Auction. By space limitation, we omit its construction, which is available at https://epistemicadl.page.link/EELP2020. Figure 3 illustrates a run in M_{bli}, where $N = \{a, b, c\}$. In state w_0, the joint action d_0 states the agents' bids. In state w_1, the good is allocated the agent with highest bid and she pays her bid.

In the next section, we derive properties from the blind auction model M_{bli} and the protocols represented by the set of rules Σ_{bli} and $\Sigma_{bli,N}$.

Protocol Valuation The next proposition shows that soundness does hold, i.e. the framework provides a sound description for Σ_{bli} and $\Sigma_{bli,N}$.

Proposition 2. *M_{bli} is an ST-Model and it is a model of Σ_{bli} and $\Sigma_{bli,N}$.*

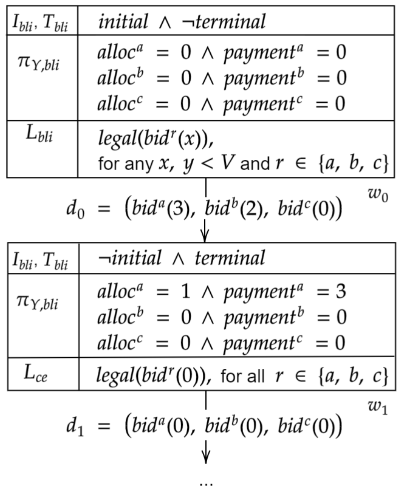

Fig. 3. A run $(w_0, d_0), (w_1, d_1)$ in M_{bli}

Proposition 3 shows that Σ_{bli} (and resp. $\Sigma_{bli,N}$) represents a one-shot protocol.

Proposition 3. $M_{bli} \models initial \rightarrow \bigcirc terminal$

If the auctioneer protocol Σ_{bli} entails a formula, then this formula is commom knowledge. Similarly, if the agents' perspective protocol $\Sigma_{bli,N}$ entails a formula, then every agent knows the formula (Proposition 4).

Proposition 4. *Given* $\varphi \in \mathcal{L}_{E-GDL}$*,*

1. *If* $M_{bli} \models \Sigma_{bli} \rightarrow \varphi$*, then* $\mathsf{C}\varphi$
2. *If* $M_{bli} \models \Sigma_{bli,N} \rightarrow \varphi$*, then* $\mathsf{E}\varphi$

Up to now, we focus on the protocol definition and semantics. Next, we address how an agent can use $\mathcal{L}_{E-ADL}$ to *choose* her actions during an auction. In other words, we describe strategy rules: E-ADL formulas which assign a unique action to be taken in each state.

4 Strategy rules

For any state-transition model M and a formula $\varphi \in \mathcal{L}_{E-ADL}$, let $S(\varphi) = \{(w, d) : M \models_{(w,d)} \varphi\}$. $S(\varphi)$ denotes all moves under which φ is valid.

Definition 7. *Given a model* M *and a strategy* S^r *for agent* $r \in N$*, a formula* $\varphi \in \mathcal{L}_{E-ADL}$ *is a representation of* S^r *iff* $S^r = S(\varphi)$*. If* $S(\varphi)$ *is a representation of a functional strategy, then* φ *is a strategy rule.*

In the following example, we illustrate a strategy rule in the Blind Auction represented by M_{bli}.

Example 1. Given the blind auction represented in the running example by M_{bli} and an agent $r \in N$, the formula

$$outbid_r =_{def} \bigvee_{v \in V} \Big(does(bid^r(v)) \wedge \\ v = min\big(\{\vartheta_r, x : \mathsf{K}_r\ x = add(max(\{\vartheta_i : i \neq r \in N\}), 1)\ \&\ x \in V\}\big)\Big)$$

is a strategy rule where the agent outbids the higher private value of her opponents or bids its own value if she knows the other agents' private values are greater then her value.

Let us assume the agent $a \in N$ has the equivalence relation R_a illustrated by Figure 4. Agent a knows the other agents, b and c, evaluate the good at most 2 while a evaluates the good at 5. Since it is a first-price auction, we can see that b and c would not maximize their utility by bidding above their private value. Thereby, the strategy $outbid_a$ consists of outbidding the adversaries private value, that is to bid 3.

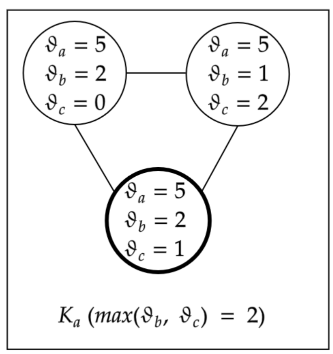

Fig. 4. Indistinguishable states for agent a

The action of $bid^a(3)$ was taken at the run illustrated by Figure 3. In this case, agent a gets the good and pays 3, which leads to a positive utility.

Proposition 5. *Given an agent r, the formula $outbid_r$ is a strategy rule.*

In the next section, we describe how to verify classic properties of mechanism design, such as efficiency, strategy-proofness and individual rationality. For the game-theoretic definition of mechanism design and its properties, please refer to [12], [4] and [15].

5 Representing Classical Properties of Auctions

Assuming the agents have private valuations $\vartheta_1, \cdots, \vartheta_n$ in V. Given an agent $r \in N$, let us denote the numerical term $utility(\vartheta_r, r) =_{def} sub(times(\vartheta_r, alloc_r), payment_r)$ as the utility of r given the private value ϑ_r. Next, we show how to represent strategy-proofness, efficiency and individual rationality properties in E-ADL.

We assume an agent is rational by the standard utility maximization definition: a rational agent has a private valuation and tries to maximize her payoff. Following [3], we consider a weak notion of rationality. An agent r is said to be *weakly rational* if she does a legal action a^r such that, for every other legal action b^r, there exists some state of the world that r considers possible, where a^r performs as well as b^r. That is if r knows that by bidding a^r it is possible to get a utility at least as good as by doing b^r.

Formally, agent $r \in N$ is *weakly rational* in the move (w, d) if

$$M \models_{(w,d)} \bigwedge_{a^r \in A^r} (does(a^r) \wedge legal(a^r) \to \bigwedge_{b^r \neq a^r \in A^r} (\neg legal(b^r) \vee \bigvee_{x,x' \in V} (\hat{\mathsf{K}}_r([a^r]utility(\vartheta_r, r) = x \wedge [b^r]utility(\vartheta_r, r) = x' \wedge x \geq x'))))$$

We denote $rat(w)$ as the set of joint actions $d \in \prod_{r \in N} A^r$ such that (w, d) is a move where all the agents are weakly rational. Notice that this notion of rationality requires an epistemic reasoning for the agents about the possible consequences of their actions.

Strategy-proofness A mechanism is strategy-proof if the agents would prefer to truthfully report their valuation rather than bidding any other possible value [15]. We say that a state $w \in W$ is *strategy-proof* for an agent $r \in N$, if bidding her private values leads to a better (or equal) utility than bidding any other value, for any joint action $d \in \prod_{i \in N} A^i$. In E-ADL, we express this condition by using the action execution operator as follows:

$$SP_r =_{def} \bigvee_{x \in V} ([bid^r(\vartheta_r)]utility(\vartheta_r, r) = x \wedge \bigwedge_{v \neq \vartheta_r \in V} [bid^r(v)]utility(\vartheta_r, r) \leq x)$$

The formula SP_r means that agent r gets the utility x when bidding her own private value and for any other value v, her utility is bellow or equal to x. In this formula, the value of the numerical term $utility(\vartheta_r, r)$ depends on the valuation of the numerical variables in the state resulting from applying the action execution operator. An auction is strategy-proof in a state when it is strategy-proof in that state for all agents and joint actions, i.e. $SP_N =_{def} \bigwedge_{r \in N} SP_r$.

For verifying whether an auction is strategy-proof, we do not make any assumption about the agents' rationality. In the blind auction described by M_{bli}, the winner pays her bid. Thereby, the agents do not have an incentive to be truthful (i.e. bid their private value).

Proposition 6. *For any $w \in I_{bli}$, $d \in \prod_{r \in N} A^r$ and $r \in N$, $M_{bli} \not\models SP_r$, and consequently, $M_{bli} \not\models SP_N$.*

Example 2. We can construct a strategy-proof blind auction by changing the payment rule from Σ_{bli}. Let us define a Vickrey Auction Σ_{vic}, such that it is exactly the same as Σ_{bli}, except by Rule 4, which is replaced by the following Rule 4′:

$$\bigvee_{second, first, x_1, \cdots, x_n \in V} \big(initial \wedge maxbid = first \wedge wins_{r,first} \wedge \\ second = max(\{0, x_r : x_r \neq maxbid \;\&\; does(bid^r(x_r) \;\&\; r \in N)\}) \rightarrow \\ \bigcirc (alloc_r = 1 \wedge payment_r = second)\big)$$

The model M_{vic} is constructed in a similar way than M_{bli}, except the update function, which assigns the second highest bid as the winner's payment.

Proposition 7. *For any $w \in I_{vic}$, $d \in \prod_{r \in N} A^r$ and $r \in N$, $M_{vic} \models_{(w,d)} SP_r$, and consequently, $M_{vic} \models_{(w,d)} SP_N$.*

Efficiency We say that a mechanism is *efficient* if it maximizes the social welfare [4]. In a single-good and single-unit auction, it means that the good should be allocated to the agent who valuates it the most, i.e. the agent with the highest private value. Here, we make an assumption about the agents' behavior: we assume they are weakly rational and only consider moves according to this assumption. Without this restriction, agents could perform random actions, and thus it would not be possible to ensure that the winning agent has the highest valuation. Efficiency (EF) in a state $w \in W$, is defined by the validity of the following E-ADL formula for any joint action $d \in rat(w)$:

$$EF =_{def} \bigwedge_{r \in N} \big(alloc_r = 1 \rightarrow \vartheta_r = max(\{\vartheta_i : i \in N\})\big)$$

Efficiency is an epistemic property: to check if an auction is efficient, an agent should reason about the knowledge of all the agents about the possible consequences of their own actions.

Proposition 8. *Given the ST-models M_{bli} and M_{vic}, then for any $w \in W_{bli}$, $w' \in W_{vic}$ and for all $d \in rat(w)$, $d' \in rat(w')$, (i) $M_{bli} \not\models_{(w,d)} EF$ and (ii) $M_{vic} \models_{(w',d')} EF$.*

If we did not assume weakly rationality of the agents, the auction represented by M_{vic} would not be efficient. Even if it is strategy-proof, we still need to link this property to the assumption they will behave rationally.

Individual Rationality A mechanism is individual-rational if an agent can always achieve as much utility as from participating as without participating [12]. We consider that if an agent does not participate in the auction than she would neither have the good allocated or a payment assigned and her utility is zero. The auction is individual-rational in a state $w \in W$ for an agent r if by participating she can achieve a utility greater or equal to zero, for any joint action $d \in \prod_{i \in N} A^i$. The individual-rationality constraint is defined by the following E-ADL formula:

$$IR_r =_{def} \bigvee_{a \in A^r} [a^r] utility(\vartheta_r, r) \geq 0$$

Similarly, we can verify whether the action is individual-rational for every agent in N by the following formula: $IR_N =_{def} \bigwedge_{r \in N} IR_r$.

The individual-rationality property consists of checking whether an agent has a reason to participate in the auction. It does not make any assumptions about her committing to a rational behavior once the auction starts. The following proposition shows M_{bli} and M_{vic} are individual-rational.

Proposition 9. *For any $w' \in W_{bli}$, $w \in W_{vic}$, $d \in \prod_{r \in N} A^r$ and $r \in N$, (i) $M_{bli} \models_{(w,d)} IR_r$, and consequently, $M_{bli} \models_{(w,d)} IR_N$; (ii) $M_{vic} \models_{(w,d)} IR_r$, and consequently, $M_{vic} \models_{(w,d)} IR_N$.*

Note that even if an auction has a property (such as strategy-proofness, efficiency, or individual rationality), it does not mean that the agents are individually or collectively aware of these properties. This knowledge may come from reasoning with previous background knowledge. Let us now discuss how knowing some classical auction properties can be meaningful for defining strategies.

5.1 Knowledge about auction properties

Agents can have different levels of knowledge over auction properties, by combining the C and K_r operators and properties formulas. For instance, $\mathsf{K}_r SP_r$ represents the agent knows the auction is strategy-proof. When that is the case, the agent can avoid any additional reasoning about her strategies and other agents' behavior: she knows that she cannot increase her utility by bidding any value different from her private value. If a weakly rational agent knows an auction is strategy-proof, then she will bid her own private value (Proposition 10).

Proposition 10. *Given an ST-model M, a state $w \in W$, a joint action $d \in \prod_{r \in N} A^i$ and an agent $r \in N$, if r weakly rational then*

$$M \models_{(w,d)} \mathsf{K}_r SP_r \wedge legal(bid^r(\vartheta_r)) \rightarrow does(bid^r(\vartheta_r))$$

Corollary 1. *Given $r \in N$, the formula $does(bid^r(\vartheta_r))$ is a strategy rule.*

To check whether the auction is strategy-proof, i.e. $M \models_{(w,d)} \mathsf{K}_r SP_r$, there is no epistemic requirement about the other agents. The agent can derive Proposition 10 by simply reasoning about the possible outcomes from her actions.

If an agent r knows that an auction is efficient and that she is not the agent with the highest valuation, then she knows she will not win the auction (Proposition 11). In this situation, assuming the payment for losing agents is zero, bidding any value below her private value will lead to the same payoff. This information about efficiency may be useful if the agent needs to choose to participate in different auctions.

Proposition 11. *Assuming the agents are weakly rational, given an ST-model M, a state $w \in W$ and any joint action in $d \in rat(w)$,*

$$\mathsf{C}EF \wedge \mathsf{K}_r \bigvee_{i \in N} \vartheta_i > \vartheta_r \to \mathsf{K}_r \bigcirc alloc_r = 0$$

If it is not IR_r, any agent with utility maximization rationality would not participate. Furthermore, if for some utility maximization rational agent i, agent r knows agent i knows it is not IR_i, then r knows i should not participate and r does not have to reason about i's bid.

E-ADL can provide an interplay between properties and agents' strategies, as illustrated by Proposition 10. When the reasoning about the classical properties is not enough to decide their strategy, agents can use the epistemic component to choose between different weakly rational actions. For instance, the First-Price Blind Auction denoted by M_{bli} is neither strategy-proof nor efficient and there may be several actions complying with the weakly rationality condition. This is a trade-off situation: the less the agent bids, the higher her utility, but the lower her chance of winning (in terms of the outcomes in the possible worlds). The strategy rule $outbid_r$ presented in Example 1 illustrates how an agent can decide among different actions in the blind auction.

6 Conclusion

In this paper, we present E-ADL, a language to allow epistemic and strategic reasoning in single-unit and single-good auctions. The language enables an agent to evaluate the auction through well-known properties from the economic theory: strategy-proofness, efficiency, and individual rationality. With E-ADL an agent can choose her bid with respect to the auction properties and her knowledge about the other players' private valuation and awareness of these properties.

For future work, we intend to generalize the definitions for describing other types of auctions, from multi-units auctions to combinatorial exchange. We also intend to investigate the complexity of the model-checking problem for E-ADL formulas, that is the problem of determining whether an E-ADL formula holds at a move under an ST-model. Promising starting points are the results of ADL and EGDL. The model checking complexity is in PTIME-complete [11] for ADL formulas and in Δ_2^p-complete for EGDL [8]. These complexity results are reasonable when compared to other languages for strategic reasoning, such as ATL and its variants. The main difference between the model-checking for E-ADL and EGDL is the action modality, whose truth condition refers to every possible joint action. Finally, an interesting line of work is to explore orders of rationality [3]: how a rational agent would strategically bid when she knows the other agents are rational, and how should she behave when they are aware of her rationality.

Acknowledgments

This research is supported by the ANR project AGAPE ANR-18-CE23-0013.

References

1. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. Journal of the ACM (JACM) **49**(5), 672–713 (2002)
2. Chatterjee, K., Henzinger, T.A., Piterman, N.: Strategy logic. Information and Computation **208**(6), 677–693 (2010)
3. Chen, J., Micali, S., Pass, R.: Tight revenue bounds with possibilistic beliefs and level-k rationality. Econometrica **83**(4), 1619–1639 (2015)
4. De Vries, S., Vohra, R.V.: Combinatorial auctions: A survey. INFORMS Journal on Computing **15**(3), 284–309 (2003). https://doi.org/10.1287/ijoc.15.3.284.16077
5. Fagin, R., Moses, Y., Halpern, J.Y., Vardi, M.Y.: Reasoning about knowledge. MIT press (2003)
6. Genesereth, M., Thielscher, M.: General game playing. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers (2014)
7. Jiang, G., Perrussel, L., Zhang, D.: On axiomatization of epistemic gdl. In: Baltag, A., Seligman, J., Yamada, T. (eds.) International Workshop on Logic, Rationality and Interaction. pp. 598–613. Springer, Berlin, Heidelberg (2017)
8. Jiang, G., Zhang, D., Perrussel, L., Zhang, H.: Epistemic GDL: A logic for representing and reasoning about imperfect information games. In: Proc. of IJCAI-2016 (2016)
9. Klemperer, P.: Auction Theory: A Guide to the Literature. Journal of Economic Surveys (1999). https://doi.org/10.1111/1467-6419.00083
10. Krishna, V.: Auction Theory. Academic Press (2009)
11. Mittelmann, M., Perrussel, L.: Auction description language (ADL): a general framework for representing auction-based markets. In: de Giacomo, G. (ed.) ECAI 2020. IOS Press, Santiago de Compostela (2020)
12. Parkes, D.C., Ungar, L.H.: Iterative combinatorial auctions: Achieving economic and computational efficiency. University of Pennsylvania Philadelphia, PA (2001)
13. Pauly, M., Parikh, R.: Game logic-an overview. Studia Logica **75**(2), 165–182 (2003)
14. Ramanujam, R., Simon, S.: Dynamic logic on games with structured strategies. In: Proc. of KR-2008. p. 49–58. AAAI Press (2008)
15. Roughgarden, T.: Algorithmic game theory. Communications of the ACM **53**(7), 78–86 (2010)
16. Thielscher, M.: A general game description language for incomplete information games. Proc. of AAAI-10 pp. 994–999 (2010)
17. Thielscher, M.: GDL-III: A description language for epistemic general game playing. Proceeding of IJCAI-2017 pp. 1276–1282 (2017)
18. Zhang, D., Thielscher, M.: A logic for reasoning about game strategies. In: Proc. of AAAI-2015. p. 1671–1677. AAAI Press (2015)
19. Zhang, D., Thielscher, M.: Representing and Reasoning about Game Strategies. Journal of Philosophical Logic **44**(2), 203–236 (2015)