

Reasoning about Human-Friendly Strategies in Repeated Keyword Auctions

Francesco Belardinelli
Université d'Evry
Évry, France
francesco.belardinelli@univ-evry.fr

Wojciech Jamroga
SnT, University of Luxembourg
Institute of Computer Science, PAS
wojciech.jamroga@uni.lu

Vadim Malvone
Télécom Paris
Paris, France
vadim.malvone@telecom-paris.fr

Munyeque Mittelmann
Université de Toulouse - IRIT
Toulouse, France
munyeque.mittelmann@irit.fr

Aniello Murano
University of Naples Federico II
Naples, Italy
nello.murano@gmail.com

Laurent Perrussel
Université de Toulouse - IRIT
Toulouse, France
laurent.perrussel@irit.fr

ABSTRACT

In online advertising, search engines sell ad placements for keywords continuously through auctions. This problem can be seen as an infinitely repeated game since the auction is executed whenever a user performs a query with the keyword. As advertisers may frequently change their bids, the game will have a large set of equilibria with potentially complex strategies. In this paper, we propose the use of natural strategies for reasoning in such setting as they are processable by artificial agents with limited memory and/or computational power as well as understandable by human users. To reach this goal, we introduce a quantitative version of Strategy Logic with natural strategies in the setting of imperfect information. In a first step, we show how to model strategies for repeated keyword auctions and take advantage of the model for proving properties evaluating this game. In a second step, we study the logic in relation to the distinguishing power, expressivity, and model-checking complexity for strategies with and without recall.

KEYWORDS

Mechanism Design; Auctions; Strategic Reasoning

ACM Reference Format:

Francesco Belardinelli, Wojciech Jamroga, Vadim Malvone, Munyeque Mittelmann, Aniello Murano, and Laurent Perrussel. 2022. Reasoning about Human-Friendly Strategies in Repeated Keyword Auctions. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 10 pages.

1 INTRODUCTION

In recent years a wealth of logic-based languages have been introduced to reason about the strategic abilities of autonomous agents in multi-agent systems (MAS), including Alternating-time Temporal Logic (ATL) [8], Strategy Logic (SL) [27, 54], and Game Logic [60], just to name a few. In conjunction with model checking techniques [10], these formal languages have allowed for the development of efficient verification tools [25, 36, 51], which have been successfully applied to the certification of MAS as different as voting protocols [13, 45], robot swarms [32, 49], and business processes [30, 38].

Still, verification tools and techniques are comparatively less developed for *data-driven* and *data-intensive* systems¹, that is, contexts where the data content of processes, or *agents*, is key to model and account for the evolution of the system [15, 55]. This is the case also for online advertising, where search engines sell ad placements for keywords continuously through auctions. This problem can be seen as an infinitely repeated game since the auction is executed whenever a user performs a query with the keyword. As advertisers may frequently change their bids, the game will have a large set of equilibria with potentially complex strategies, thus making the specification and verification of keyword auctions a complex problem to solve for current model checking methods².

In this paper, we investigate the use of *natural strategies* [46, 47] for reasoning about equilibria in keyword auctions. The work in [46, 47] followed some classical research on human concept learning [18, 35], social norms [67, 68], commonsense reasoning [29], automated planning [37, 66] and the psychology of planning [56], usability [57], agent-oriented programming [17, 42], as well as robotics [58, 62]; in short, it proposed to model “human-friendly” strategies by lists of condition-action pairs with bounded complexity. This was in contrast to “combinatorial” strategies, defined as functions from (sequences of) states to actions, and typically used in the semantics of MAS logics [8, 27, 54, 60]. It was argued in [46, 47] that natural strategies provide better models of behaviour for agents with limited memory and computing capacity, such as humans or simple bots. The concept have been already used to redefine some security requirements for voting protocols in [45].

In our case, the bidding strategy in an auction should be executable for a *simple* artificial agent, as well as reasonably transparent to the human user, which makes natural strategies a good match. Moreover, natural strategies provide a way to define complexity (and hence also “simplicity”) metrics for various functionality, security, and usability properties in MAS. By focusing on simple strategies, one can make the verification of equilibrium properties

¹“[Model checking] is mainly appropriate to control-intensive applications and less suited for data-intensive applications” [10, p. 15]

²“In principle, the sets of equilibria in such repeated games can be very large, with players potentially punishing each other for deviations. The strategies required to support such equilibria are usually quite complex, however, requiring precise knowledge of the environment and careful implementation. In theory, advertisers could implement such strategies via automated robots, but in practice they may not be able to: bidding software must first be authorized by the search engines, and search engines are unlikely to permit strategies that would allow advertisers to collude and substantially reduce revenues.” [34]

decidable, or even tractable, despite the prohibitive complexity of the general problem. This is especially evident for strategies with memory, which normally make the synthesis and model checking problems undecidable [31, 72].

Contribution By leveraging on natural strategies, we introduce a quantitative semantics for SL with natural strategies and imperfect information. As a first contribution, we show how to represent popular strategies for repeated keyword auctions in the proposed framework, as well as prove properties pertaining to this game.

Second, we analyse our novel variant of SL in relation with its distinguishing power, expressivity, and complexity of the model checking problem, for natural strategies with and without recall.

Outline In Section 2, we recall the basic definitions. Section 3 introduces the new logic $\text{NatSL}[\mathcal{F}]$. In Section 4 we use it to analyse repeated keyword auctions. The expressivity of $\text{NatSL}[\mathcal{F}]$ is investigated in Section 5. Section 6 establishes the complexity of model checking, and Section 7 concludes the paper. The omitted proofs are available at <http://arxiv.org/abs/2201.09616>.

1.1 Related work

Recently, there have been efforts to apply formal methods to the (semi-)automatic verification of some decision-making mechanisms, including auctions and voting protocols. A number of works [12, 23, 48] expresses these mechanisms in high-level specification languages. However, in contrast with standard model checking techniques [10], their verification is not fully automated, but only assisted by a reasoner. Troquard et al. [69] introduce a framework for fully-automated verification of voting protocols. Still, their approach can only model one-shot mechanisms and thus does not capture multi-stage protocols and repeated auctions. In preliminary works, Pauly and Wooldridge [61] and Wooldridge et al. [75] advocate the use of ATL [8] to reason about decision-making mechanisms. As ATL lacks the expressivity to reason about quantitative aspects such as valuations and payments, and solution concepts such as equilibria, Maubert et al. [53] introduce $\text{SLK}[\mathcal{F}]$, a quantitative and epistemic version of SL [27, 54], and show how it can be used for reasoning about notions such as Nash equilibrium and strategyproofness. Still, their approach considers strategies as functions from states to actions and cannot handle strategies with recall.

A key assumption of the present contribution is that agents have only partial observability of the global state of the system, as it is often the case in real-life applications. Contexts of imperfect information have been extensively considered in the literature on formal verification [22, 31, 44, 50, 64]. Generally speaking, imperfect information immediately entails higher complexity of game solving. In multi-player games, the complexity can go up to being non-elementary [63], or even undecidability when considered in the context of memoryful strategies [31]. Hence, it is of interest to analyse imperfect information systems where agents have finite or bounded memory, in order to retrieve a decidable model checking problem. Works that are closest in spirit to our contribution concern modeling, specification, and reasoning about strategies of bounded-memory agents. We directly build on the research by Jamroga, Malvone, and Murano on natural strategies [46, 47]. We generalize the approach by considering quantitative semantics for

both natural strategies and the logic, which is more suitable for reasoning about mechanisms with monetary transfer (e.g., auctions). We also consider SL instead of ATL, due to its expressive power. In a related vein, Ågotnes and Walther [2] investigate strategic abilities of agents with bounded memory, while Belardinelli et al. [14] consider bounded memory as an approximation of perfect recall. On a related direction, temporal and strategic logics have been extended to handle agents with bounded resources [5, 6, 20, 21]. Issues related to bounded rationality are also investigated in [11, 39, 43].

Also relevant for the present contribution are papers that study explicit representations of strategies. This category is much richer and includes extensions of ATL^* with explicit reasoning about actions and strategies [1, 41, 70, 73], as well as logics that combine features of temporal and dynamic logic [40, 59]. Duijf and Broersen [33] present a variant of STIT logic, that enables reasoning about strategies and their performance in the object language. Also, plans in agent-oriented programming are in fact rule-based descriptions of strategies. In particular, reasoning about agent programs using strategic logics was investigated in [3, 4, 16, 28, 76].

2 PRELIMINARIES

We first recall basic notions. For the remainder of the paper, we fix a set of atomic propositions AP , a set of agents Ag and a set of strategy variables Var . We let n be the number of agents in Ag . Finally, let $\mathcal{F} \subseteq \{f: [-1, 1]^m \rightarrow [-1, 1] \mid m \in \mathbb{N}\}$ be a set of functions over $[-1, 1]$ of possibly different arities.

2.1 Weighted Concurrent Game Structures

The semantics of natural strategies and $\text{NatSL}[\mathcal{F}]$ are interpreted over weighted concurrent game structures (wCGS). A difference from classical structures is that the labelling of atomic propositions is replaced by a weight function. We consider weighted propositions for easily handling quantitative aspects (such as prices).

Definition 1. A *weighted concurrent game structure with imperfect information* (wCGS) is a tuple $\mathcal{G} = (\text{Ac}, V, L, \delta, \ell, V_i, \{\sim_a\}_{a \in \text{Ag}})$ where: (i) Ac is a finite set of actions; (ii) V is a finite set of states; (iii) $L: \text{Ag} \times V \rightarrow 2^{\text{Ac}}$ is a *legality function*, defining the availability of actions; (iv) δ is a transition function assigning a successor state $v' = \delta(v, (c_a)_{a \in \text{Ag}})$ to each state $v \in V$ and any tuple of actions $(c_a)_{a \in \text{Ag}}$, where $c_a \in L(a, v)$; (v) $\ell: V \times \text{AP} \rightarrow [-1, 1]$ is a *weight function*; (vi) $V_i \subseteq V$ is a set of *initial states*; and (vii) $\sim_a \subseteq V \times V$ is an equivalence relation called the *observation relation* of agent a .

We require that the wCGS is uniform, that is $v \sim_a v'$ implies $L(a, v) = L(a, v')$. We write \mathbf{o} for a tuple of objects $(o_a)_{a \in \text{Ag}}$, one for each agent, and such tuples are called *profiles*. Given a profile \mathbf{o} and $a \in \text{Ag}$, we let o_a be agent a 's component, and o_{-a} is $(o_i)_{i \in \text{Ag} \setminus \{a\}}$. Similarly, we let $\text{Ag}_{-a} = \text{Ag} \setminus \{a\}$.

In a state $v \in V$, each player a chooses an available action $c_a \in L(a, v)$, and the game proceeds to state $\delta(v, \mathbf{c})$ where \mathbf{c} is the *action profile* $(c_a)_{a \in \text{Ag}}$. A *play* $\pi = v_0 v_1 v_2 \dots$ is an infinite sequence of states such that for every $i \geq 0$ there exists an action profile \mathbf{c} such that $\delta(v_i, \mathbf{c}) = v_{i+1}$. We write $\pi_i = v_i$ for the state at index i in play π . A *history* $h = v_0 v_1 v_2 \dots v_n$ is a finite sequence of states. The last element of a history is denoted by $\text{last}(h) = v_n$. $H_{\mathcal{G}}$ denotes the set of all histories in the wCGS \mathcal{G} .

2.2 Natural Strategies

In this section we recall the notion of uniform natural strategies from [47]. Natural strategies are conditional plans, represented through an ordered list of condition-action rules [47]. The intuition is that the first rule whose condition holds in the history of the game is selected, and the corresponding action is executed. As we are considering the setting of imperfect information, the conditions are regular expressions over *weighted epistemic (WE) formulas*. Given an agent a , the WE formulas over AP, denoted $WE(AP)$, are conditions on a 's knowledge and are expressed by the following Backus-Naur Form grammar:

$$\begin{aligned}\psi &::= \top \mid K_a \varphi \mid f(\psi, \dots, \psi) \\ \varphi &::= p \mid f(\varphi, \dots, \varphi) \mid K_i \varphi\end{aligned}$$

where $f \in \mathcal{F}$ is a function, $p \in AP$ is an atomic proposition and $i \in Ag$ is an agent.

Given a wCGS \mathcal{G} , a state $v \in V$ and a $WE(AP)$ formula φ , we inductively define the satisfaction value of φ in v , denoted $\llbracket \varphi \rrbracket(v)$:

$$\begin{aligned}\llbracket p \rrbracket(v) &= \ell(v, p) \\ \llbracket K_a \varphi \rrbracket(v) &= \min_{v' \sim_a v} \llbracket \varphi \rrbracket(v') \\ \llbracket f(\varphi_1, \dots, \varphi_m) \rrbracket(v) &= f(\llbracket \varphi_1 \rrbracket(v), \dots, \llbracket \varphi_m \rrbracket(v))\end{aligned}$$

The semantics for the knowledge modality is the standard in the literature on fuzzy epistemic logic (e.g. [52]). Let $Reg(WE(AP))$ be the set of regular expressions over the weighted epistemic conditions $WE(AP)$, defined with the constructors $\cdot, \cup, *$ representing concatenation, nondeterministic choice, and finite iteration, respectively. Given a regular expression r and the language $\mathcal{L}(r)$ on words generated by r , a history h is *consistent* with r iff there exists $b \in \mathcal{L}(r)$ such that $|h| = |b|$ and $\llbracket b[i] \rrbracket(h[i]) = 1$, for all $0 \leq i \leq |h|$. Intuitively, a history h is consistent with a regular expression r if the i -th weighted epistemic condition in r ‘‘holds’’ in the i -th state of h (for any position i in h).

A *uniform natural strategy with recall* σ_a for agent a is a sequence of pairs (r, c) , where $r \in Reg(WE(AP))$ is a regular expression, and c is an action available in $last(h)$, for all histories $h \in H_{\mathcal{G}}$ consistent with r . The last pair on the sequence is required to be (\top^*, c) , with $c \in L(a, v)$ for every $v \in V$ and some $c \in Ac$.

A *uniform memoryless natural strategy* is a special case of natural strategy in which each condition is a weighted epistemic formula (i.e., no regular operators are allowed).

Natural strategies are uniform in the sense they specify the same actions in indistinguishable states (see [47]). We define $Str_a^\rho = \cup_{a \in Ag} Str_a^\rho$, where $\rho \in \{ir, iR\}^3$. Let $size(\sigma_a)$ denote the number of guarded actions in σ_a , $cond_i(\sigma_a)$ be the i -th guarded condition on σ_a , $cond_i(\sigma_a)[j]$ be the j -th WE formula of the guarded condition σ_a , and $act_i(\sigma_a)$ be the corresponding action. Finally, $match(h, \sigma_a)$ is the smallest index $i \leq size(\sigma_a)$ such that for all $0 \leq j \leq |last(h)|$, $\llbracket cond_i(\sigma_a)[j] \rrbracket(h[j]) = 1^4$ and $act_i(\sigma_a) \in L(a, last(h))$. In other

³As usual in the verification process, we denote imperfect recall with r , perfect recall with R , imperfect information with i , and perfect information with I .

⁴Note that, we considered the case in which the condition have the same length of the history. There is also the case in which the condition is shorter than the history. This is due to the usage of the finite iteration operator. In the latter case, we need to check a finite number of times the same weighted epistemic formula in different states of the history. For more details on this aspect see [46, 47].

words, $match(h, \sigma_a)$ matches the state $last(h)$ with the first condition in σ_a that holds in h , and action available in $last(h)$.

Measurement of Natural Strategies. The complexity of the strategy σ is the total size of its representation and is denoted as follows: $compl(\sigma) := \sum_{(r,c) \in \sigma} |r|$, where $|r|$ is the number of symbols in r , except by parentheses. If r is a n -ary function in \mathcal{F} , then $|r| = n + 1$.

3 NATURAL STRATEGY LOGIC

SL[\mathcal{F}] [19] proposes a quantitative semantics for Strategy Logic, in which strategies are functions mapping histories to actions. For reasoning about intuitive and simple strategies, we introduce SL[\mathcal{F}] with natural strategies and imperfect information, denoted NatSL[\mathcal{F}]. Throughout this section, let $\rho \in \{ir, iR\}$ denote whether the semantics considers memoryless or recall strategies.

An *assignment* $\chi : Ag \cup Var \rightarrow Str^\rho$ is a function from players and variables to strategies. For an assignment χ , an agent a and a strategy σ for a , $\chi[a \mapsto \sigma]$ is the assignment that maps a to σ and is otherwise equal to χ , and $\chi[s \mapsto \sigma]$ is defined similarly, where s is a variable. For an assignment χ and a state v we let $Out(\chi, v)$ be the unique play that starts in v and follows the strategies assigned by χ . Formally, $Out(\chi, v)$ is the play $v_0 v_1 \dots$ such that $v_0 = v$ and for all $i \geq 0$, $v_{i+1} = \delta(v_i, c)$ where for all $a \in Ag$, $c_a = act_{match(v_i, \chi(a))}(\chi(a))$.

3.1 NatSL[\mathcal{F}] Syntax

Definition 2. The syntax of NatSL[\mathcal{F}] is defined as follows:

$$\varphi ::= p \mid \exists s_a^{\leq k}. \varphi \mid (a, s_a)\varphi \mid f(\varphi, \dots, \varphi) \mid X\varphi \mid \varphi U \varphi$$

where $p \in AP$, $s_a \in Var \cup Str_a^\rho$, $a \in Ag$, and $f \in \mathcal{F}$.

The intuitive reading of the operators is as follows: $\exists s_a^{\leq k}. \varphi$ means that there exists a strategy with complexity less or equal than k for agent a such that φ holds; $(a, s_a)\varphi$ means that when strategy s_a is assigned to agent a , φ holds; X and U are the usual temporal operators ‘‘next’’ and ‘‘until’’. The meaning of $f(\varphi_1, \dots, \varphi_n)$ depends on the function f . We use \top, \vee , and \neg to denote, respectively, function 1, function $x, y \mapsto \max(x, y)$ and function $x \mapsto \neg x$.

A variable is *free* in formula φ if it is bound to an agent without being quantified upon, and an agent a is free in φ if φ contains a temporal operator (X or U) that is not in the scope of any binding for a . The set of free variables and agents in φ is written $free(\varphi)$, and a formula φ is a *sentence* if $free(\varphi) = \emptyset$. The strategy quantifier $\exists s_a^{\leq k}. \varphi$ quantifies on strategies for agent a .

3.2 NatSL[\mathcal{F}] Semantics

Definition 3. Let $\mathcal{G} = (Ac, V, \delta, \ell, V_i, \{\sim_a\}_{a \in Ag})$ be a wCGS, and χ an assignment. The satisfaction value $\llbracket \varphi \rrbracket_\chi^{\mathcal{G}, \rho}(v) \in [-1, 1]$ of a NatSL[\mathcal{F}] formula φ in a state v is defined as follows, where π denotes $Out(v, \chi)$:

$$\begin{aligned}\llbracket p \rrbracket_\chi^{\mathcal{G}, \rho}(v) &= \ell(v, p) \\ \llbracket \exists s_a^{\leq k}. \varphi \rrbracket_\chi^{\mathcal{G}, \rho}(v) &= \max_{\sigma \in \{\alpha \in Str_a^\rho : compl(\alpha) \leq k\}} \llbracket \varphi \rrbracket_{\chi[s_a \mapsto \sigma]}^{\mathcal{G}, \rho}(v) \\ \llbracket (a, s_a)\varphi \rrbracket_\chi^{\mathcal{G}, \rho}(v) &= \llbracket \varphi \rrbracket_{\chi[a \mapsto \chi(s_a)]}^{\mathcal{G}, \rho}(v) \text{ if } s_a \in Var \\ \llbracket (a, \sigma_a)\varphi \rrbracket_\chi^{\mathcal{G}, \rho}(v) &= \llbracket \varphi \rrbracket_{\chi[a \mapsto \sigma_a]}^{\mathcal{G}, \rho}(v) \text{ if } \sigma_a \notin Var\end{aligned}$$

$$\begin{aligned} \llbracket f(\varphi_1, \dots, \varphi_m) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) &= f(\llbracket \varphi_1 \rrbracket_{\chi}^{\mathcal{G}, \rho}(v), \dots, \llbracket \varphi_m \rrbracket_{\chi}^{\mathcal{G}, \rho}(v)) \\ \llbracket \mathbf{X}\varphi \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) &= \llbracket \varphi \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_1) \\ \llbracket \varphi_1 \mathbf{U} \varphi_2 \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) &= \sup_{i \geq 0} \min \left(\llbracket \varphi_2 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_i), \min_{0 \leq j < i} \llbracket \varphi_1 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_j) \right) \end{aligned}$$

If φ is a sentence, its satisfaction value does not depend on the assignment, and we write $\llbracket \varphi \rrbracket^{\mathcal{G}, \rho}(v)$ for $\llbracket \varphi \rrbracket_{\chi}^{\mathcal{G}, \rho}(v)$ where χ is any assignment. We also let $\llbracket \varphi \rrbracket^{\mathcal{G}, \rho} = \min_{v_i \in V_i} \llbracket \varphi \rrbracket^{\mathcal{G}, \rho}(v_i)$.

Remark 1. When propositions only take values in $\{-1, 1\}$ and $\mathcal{F} = \{\top, \vee, \neg\}$, $\text{NatSL}[\mathcal{F}]$ corresponds to a Boolean-valuated extension of SL with Natural Strategies.

We define the classic abbreviations: $\perp := \neg\top$, $\varphi \rightarrow \varphi' := \neg\varphi \vee \varphi'$, $\varphi \wedge \varphi' := \neg(\neg\varphi \vee \neg\varphi')$, $\mathbf{F}\psi := \top \mathbf{U}\psi$, $\mathbf{G}\psi := \neg\mathbf{F}\neg\psi$ and $\forall s^{\leq k}. \varphi := \neg\exists s^{\leq k}. \neg\varphi$, and check that they correspond to the intuition. For instance, \wedge corresponds to \min , $\mathbf{F}\psi$ computes the supremum of the satisfaction value of ψ over all future points in time, $\mathbf{G}\psi$ computes the infimum of these values, and $\forall s^{\leq k}. \varphi$ minimizes the value of φ over all possible strategies s .

4 REPEATED KEYWORD AUCTIONS

Modeling mechanisms with monetary transfer and private valuations require handling quantitative features and imperfect information. Memoryless strategies are enough for mechanisms in which all relevant information is encoded in the current state (e.g. English auction). In repeated auctions, agents may, as well, use information from the previous states for choosing their strategies.

We now focus on using $\text{NatSL}[\mathcal{F}]$ to model and verify repeated keyword auctions and related strategies. Repeated keyword auctions are used by online search engines for selling advertising slots when users perform a search with a keyword [24]. For a keyword of interest, the advertisers (bidders) submit a bid stating the maximum amount she is willing to pay for a click on her sponsored link. When a user submits a query, an auction is run to determinate the slot allocation among the advertisers bidding on the keyword of interest. The most common mechanism for keyword auctions is the Generalized Second Price (GSP) [24], in which the agents are allocated slots in decreasing order of bids and the payment for the slot s is the bid of the agent allocated to the slot $s + 1$.

We assume that \mathcal{F} contains the function $\leq : (x, y) \mapsto 1$ if $x \leq y$ and $\leq : (x, y) \mapsto -1$ otherwise; and for readability we use the infix notation $x \leq y$ in the formula. We also assume that \mathcal{F} contains the equality = and comparison functions $<$, $>$, \geq (defined similarly). Finally, we assume \mathcal{F} contains functions \neg , \sum , \times , \setminus , \min , \max and argmax with the standard meaning (for details, see [53]).

Let us fix a price increment $\text{inc} \in (0, 1]$, a set of slots $S = \{1, \dots, m\}$, where $m \in \mathbb{N} \setminus \{0\}$. Each slot has a click-through rate $\theta_1 > \dots > \theta_m$, where $\theta_s \in [0, 1]$ is the probability that the user will click on the advertisement in slot s . The agents in Ag are the advertisers, each one having a private valuation $v_a \in V_a$ for a click, where $V_a \subset [0, 1]$ is a finite set of possible valuations. We assume the valuations are distinct, that is, if $a \neq a'$, then $v_a \neq v_{a'}$. We denote by $<$ an arbitrary order among the agents in Ag , used in case of ties. The atomic propositional set is $\text{AP} = \{\text{all}_{a,s}, p_s, \vartheta_a : a \in \text{Ag}, s \in S\}$, where $\text{all}_{a,s}$ represents whether agent a is allocated to slot s , p_s

denotes the price of slot s and ϑ_a denotes a 's valuation. Define $\mathcal{G}_{\text{GSP}} = (\text{Ac}, V, L, \delta, \ell, V_i, \{\sim_a\}_{a \in \text{Ag}})$, where:

- $\text{Ac} = \{0 + x \times \text{inc} : 0 \leq x \leq \frac{1}{\text{inc}}\}$, where $b \in \text{Ac}$ denotes a bid with price b for a click; given $c = (c_a)_{a \in \text{Ag}}$, let $\text{rank}_c = (a_1, \dots, a_n)$ be the sequence of distinct agents in Ag ordered by their bid, that is, $i < j$ if $c_{a_i} > c_{a_j}$ or $c_{a_i} = c_{a_j}$ and $a_i < a_j$ for $i, j \in \{1, \dots, n\}$ with $i \neq j$. In case of draws, the sequence is determined with respect to $<$. We let $\text{rank}_c(i)$ denote the agent in the i -th position of the sequence rank_c .
- $V = \{\langle \text{al}_1, \dots, \text{al}_m, \text{pr}_1, \dots, \text{pr}_m, (\text{vl}_a)_{a \in \text{Ag}} \rangle : \text{al}_s \in \text{Ag} \cup \{\text{none}\} \ \& \ \text{pr}_s \in \text{Ac} \ \& \ \text{vl}_a \in V_a \ \& \ a \in \text{Ag} \ \& \ 1 \leq s \leq m\}$, where each state represents the current slot allocation and prices, with al_s, pr_s , and vl_a denoting the winner of slot s , the price per click of s and a 's valuation, resp.;
- For each $a \in \text{Ag}$ and $v \in V$, $L(a, v) = \text{Ac}$;
- For each $v \in V$ and $c = (c_a)_{a \in \text{Ag}}$ such that $c_a \in L(a, v)$, the transition function uses the agent's bids to chose the next allocations and prices and is defined as follows: $\delta(v, (c_a)_{a \in \text{Ag}}) = \langle \text{al}'_1, \dots, \text{al}'_m, \text{pr}'_1, \dots, \text{pr}'_m, (\text{vl}'_a)_{a \in \text{Ag}} \rangle$, where for each agent a and slot s . (i) $\text{al}'_s = \text{rank}_c(s)$ if $s \leq n$, and $\text{al}'_s = \text{none}$ otherwise; (ii) $\text{pr}'_s = c_{\text{rank}_c(s+1)}$ if $s+1 \leq n$, and $\text{pr}'_s = 0$ otherwise.
- For each agent a , slot $s \in S$ and state $v = \langle \text{al}_1, \dots, \text{al}_m, \text{pr}_1, \dots, \text{pr}_m, (\text{vl}_a)_{a \in \text{Ag}} \rangle$, the weight function is defined as follows: (i) $\ell(v, \text{all}_{a,s}) = 1$ if $\text{al}_s = a$, and $\ell(v, \text{all}_{a,s}) = 0$ otherwise; (ii) $\ell(v, p_s) = \text{pr}_s$; and (iii) $\ell(v, \vartheta_a) = \text{vl}_a$.
- In an initial state, the prices are 0 and the slots are allocated to *none*, that is, $V_i = \{\langle \text{none}, \dots, \text{none}, 0, \dots, 0, \text{vl}_1, \dots, \text{vl}_n \rangle \in V\}$;
- For each agent a and two states $v = \langle \text{al}_1, \dots, \text{al}_m, \text{pr}_1, \dots, \text{pr}_m, (\text{vl}_a)_{a \in \text{Ag}} \rangle$ and $v' = \langle \text{al}'_1, \dots, \text{al}'_m, \text{pr}'_1, \dots, \text{pr}'_m, (\text{vl}'_a)_{a \in \text{Ag}} \rangle$ in V , the observation relation \sim_a is such that if $v \sim_a v'$ then (i) $\text{al}_s = \text{al}'_s$, for each $1 \leq s \leq m$; (ii) $p_s = p'_s$, for each $1 \leq s \leq m$; (iii) $\text{vl}_a = \text{vl}'_a$.

Notice there is exactly one initial state for each possible valuation profile in $(\prod_{a \in \text{Ag}} V_a)$. Additionally, valuations remain unchanged after the initial state. We use the formula $s^{-1} := 1 \setminus s$ when it is convenient to obtain a value in $[-1, 1]$ for representing a slot s . The utility of agent a when she is assigned to slot s is denoted by the formula $\text{ut}_{a,s} := \theta_s \times (\vartheta_a - p_s)$. The expected utility for agent a depends on her actual allocation, that is, $\text{ut}_a := \sum_{s \in S} \text{all}_{a,s} \times \text{ut}_{a,s}$.

4.1 Solution concepts for \mathcal{G}_{GSP}

In this section, we show how $\text{NatSL}[\mathcal{F}]$ can be used for the verification of mechanisms with natural strategies. In sight of our motivating example, we aim at rephrasing conditions and properties usually considered in the analysis of keyword auctions [24, 34, 71]. **Nash equilibrium** Since auctions are noncooperative, the solution concept in the pure strategy setting usually considered is the Nash equilibrium (NE). The NE captures the notion of stable solution: a strategy profile is NE if no player can improve her utility through an unilateral change of strategy [65]. With $\text{NatSL}[\mathcal{F}]$, we restrict the range of strategies to simple ones, as it enables us to reason about artificial agents with limited capabilities and human-friendly strategies. Let $\sigma = (\sigma_a)_{a \in \text{Ag}}$ be a profile of strategies and $k > 0$

and define the formula

$$\text{NE}(\sigma, k) := \bigwedge_{a \in \text{Ag}} \forall t \leq k. [(\text{Ag}_{-a}, \sigma_{-a})(a, t) \text{Xut}_a \leq (\text{Ag}, \sigma) \text{Xut}_a]$$

The formula $\text{NE}(\sigma, k)$ means that, for every agent and alternative strategy t of complexity at most k , binding to t when everyone else binds to their strategies in σ leads to at most the same utility as when she also binds to her strategy in σ . In relation to strategies with complexity at most k , the strategy profile σ leads to a NE in the next state of v if $\llbracket \text{NE}(\sigma, k) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) = 1$.

Predicting outcomes of a keyword auction is a difficult task given the infinite nature of NE continuum [77]. For this reason, refined solution concepts have been proposed to reduce the NE continuum to subsets. Edelman *et al.* [34] studied the subset called locally envy-free equilibrium (LEFE), in which no advertiser can improve her utility by exchanging her current slot to the one ranked one position above, given the current prices.

Locally envy free equilibrium Let $\sigma = (\sigma_a)_{a \in \text{Ag}}$ be a profile of strategies, we define the formula

$$\text{LEFE}(\sigma) := \bigwedge_{a \in \text{Ag}} (\text{Ag}, \sigma) \text{X} [\text{LEF}_{\text{wins}}^a \wedge \text{LEF}_{\text{loses}}^a]$$

where $\text{LEF}_{\text{wins}}^a := \bigwedge_{1 < s \leq m} (\text{all}_{a,s} = 1 \rightarrow \text{ut}_{a,s} \geq \text{ut}_{a,s-1})$ indicates that when an agent is allocated to a slot, she does not prefer to switch to the slot right above and $\text{LEF}_{\text{loses}}^a := (\bigwedge_{s \in S} \text{all}_{a,s} = 0) \rightarrow 0 \geq \text{ut}_{a,m}$ denotes that agents who were not assigned to any slot do not prefer to get the last slot.

$\text{LEFE}(\sigma)$ means that, for any agent, when everyone follows the strategies in σ , it holds that (i) if she wins s , her utility for s is greater than for slot $s - 1$ (at current prices) and (ii) if she does not get any slot, then her utility for the last slot is at most zero. Strategy profile σ leads to LEFE in the next state of v if $\llbracket \text{LEFE}(\sigma) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) = 1$. Based on [34, 71], we have that any LEFE is also a NE:

Proposition 1. *For any complexity $k \geq 0$, state $v \in V$, $\rho \in \{iR, ir\}$ and strategy profile $\sigma = (\sigma_a)_{a \in \text{Ag}}$ with $\sigma_a \in \text{Str}_a^{\rho}$ for each agent a , $\llbracket \text{LEFE}(\sigma) \rightarrow \text{NE}(\sigma, k) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) = 1$.*

As LEFE is still an equilibrium continuum, Edelman *et al.* [34] characterize an equilibrium in which the slot allocation and payments coincide with the ones in the dominant-strategy equilibrium (DSE) of the Vickrey–Clarke–Groves (VCG) mechanism.

Let $\mathbf{v} = (v_a)_{a \in \text{Ag}}$ be a valuation profile. Truthfully reporting \mathbf{v} is the DSE of VCG [65]. For each slot s and agent a , the allocation rule for VCG in the keyword auction is the same as under GSP [34]: $\text{all}_{a,s}^*(\mathbf{v}) = 1$ if $\text{rank}_{\mathbf{v}}(s) = a$ and $s \leq n$. Otherwise, $\text{all}_{a,s}^*(\mathbf{v}) = 0$. The payment for the last slot m is $p_m^*(\mathbf{v}) = \theta_m \cdot v_{\text{rank}_{\mathbf{v}}(m+1)}$ if $m+1 \leq n$ and $p_m^* = 0$ otherwise. For the remaining slots $1 \leq s < m$, $p_s^*(\mathbf{v}) = (\theta_s - \theta_{s+1}) \cdot v_{\text{rank}_{\mathbf{v}}(s+1)} + p_{s+1}^*(\mathbf{v})$. We assume $p_s^*(\mathbf{v})$ and $\text{all}_{s,a}^*(\mathbf{v})$ are functions in \mathcal{F} .

VCG outcome The following formula denotes whether the allocation and payments in the next state are the same as the ones for the VCG when agents bid truthfully:

$$\varphi_{\text{VCG}}(\sigma) := (\text{Ag}, \sigma) \text{X} \left[\bigwedge_{s \in S} (p_s = p_s^*(\mathbf{\theta}) \wedge \bigwedge_{a \in \text{Ag}} \text{all}_{s,a} = \text{all}_{s,a}^*(\mathbf{\theta})) \right]$$

If a strategy profile leads to the VCG outcome, then it is a LEFE:

Proposition 2. *For any state $v \in V$, $\rho \in \{iR, ir\}$ and strategy profile $\sigma = (\sigma_a)_{a \in \text{Ag}}$ with $\sigma_a \in \text{Str}_a^{\rho}$ for each agent a , $\llbracket \varphi_{\text{VCG}}(\sigma) \rightarrow \text{LEFE}(\sigma) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) = 1$.*

In fact, from [34, 71] the VCG payments are the lower bound of locally envy-free equilibrium. Thus, in any other locally envy-free equilibrium the total revenue obtained by GSP is at least as high as the one obtained by VCG in equilibrium.

Corollary 1. *For any state $v \in V$, $\rho \in \{iR, ir\}$ and strategy profile $\sigma = (\sigma_a)_{a \in \text{Ag}}$ with $\sigma_a \in \text{Str}_a^{\rho}$ for each agent a , $\llbracket \text{LEFE}(\sigma) \rightarrow \sum_{s \in S} (p_s) \geq \sum_{s \in S} (p_s^*) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) = 1$.*

The solution concepts characterized in the previous section are considered in a single stage of the game. Since the auction is repeated, advertisers can change their bids very frequently and one may investigate whether the prices stabilize and at what values [34]. Stable bids must be best responses to each other, that is, the bids form an (one-shot) equilibrium. Cary *et al.* [24] raises the problem on whether there exists a “natural bidding strategy” for the advertisers that would lead to equilibrium.

Convergence The concept of convergence or stabilization can be easily encoded in NatSL[\mathcal{F}]: we say a wCGS \mathcal{G} converge to a property φ if the initial states lead to φ being eventually always the case. Formally, a wCGS converge to a condition φ if $\llbracket \text{FG}(\varphi) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v_i) = 1$ for each initial state $v_i \in V_i$.

4.2 Natural Strategies for \mathcal{G}_{GSP}

Given agent a and the wCGS \mathcal{G}_{GSP} , we exemplify strategies for a in a repeated keyword auction. For readability, we omit the epistemic operator K_a from an epistemic condition $K_a \varphi$ when the satisfaction value of φ is known by a in all states. A common approach for an advertiser is to assume that all the other bids will remain fixed in the next round and target the slot that maximizes her utility at current prices. This mechanism allows a range of bids that will result in the same outcome from a 's perspective, so a number of strategies are distinguished by the bid choice within this range.

Balanced bidding In the balanced bidding strategy (BB) [24], the agent bids so as to be indifferent between successfully winning the targeted slot at its current price, or winning a slightly more desirable slot at her bid price. The natural strategy representing balanced bidding for agent a is denoted BB_a and is constructed in three parts. First, include the guarded actions $(BB_{a,1}(b), b)$ for each action $b \in \text{Ac}$. Second, include $(BB_{a,2}(b, s), b)$ for each $b \in \text{Ac}$ and $1 < s \leq m$. Third, the last guarded action is $(\top, 0)$. The condition $BB_{a,1}(b)$ refers to the case in which the slot maximizing a 's utility is the top slot and b is $(\vartheta_a + p_1)/2$:

$$BB_{a,1}(b) := b = \frac{\vartheta_a + p_1}{2} \wedge (\text{argmax}_{s \in S} (\text{ut}_{a,s}))^{-1} = 1$$

Condition $BB_{a,2}(b, s)$ denotes the case in which the slot $s \neq 1$ maximizes a 's utility and b is the bid value that is high enough to force the prices paid by her competitors to rise, but not so high that she would mind getting a higher slot at a price just below b .

$$BB_{a,2}(b, s) := \text{ut}_{a,s} = \theta_{s-1} \times (\vartheta_a - b) \wedge (\text{argmax}_{s' \in S} (\text{ut}_{a,s'}))^{-1} = s^{-1}$$

Notice the guarded action $BB_{a,2}(b, s)$ is defined for $s > 1$ since it compares the utility with the one for $s - 1$. The case $s = 1$ is treated by the guarded action $BB_{a,1}(b)$.

Given a valuation profile $\mathbf{v} = (v_a)_{a \in \text{Ag}}$, let η_x be the agent in the x -th position of rank_v (that is, η_x is the agent with x -th highest valuation). We let $b_{\eta_x}(\mathbf{v})$ be a function in \mathcal{F} defined as follows:

$$b_{\eta_x}(\mathbf{v}) = \begin{cases} \frac{\theta_x}{\theta_{x-1}} \cdot b_{\text{rank}_v(x+1)}(\mathbf{v}) + (1 - \frac{\theta_x}{\theta_{x-1}})v_{\eta_x} & \text{if } x \geq m+1 \\ v_{\eta_x} & \text{if } 2 \leq x \leq m \end{cases}$$

If $\mathbf{BB} = (BB_a)_{a \in \text{Ag}}$ converges to the equilibrium with VCG outcomes, the agent with the highest valuation bids any value above $b_{\eta_2}(\mathbf{v})$. The equilibrium bid for $a \neq \eta_1$ is $b_a(\mathbf{v})$ [24]. When there are two slots and all players update their bids according to BB, the game converges to the equilibrium with VCG outcome. However, this is not the case for more than two slots [24].

Proposition 3. *For any initial state $v_i \in V_i$, state $v \in V$, and $1 < x \leq n$, the following holds, where $\mathbf{v} = (\ell(v, \vartheta_a))_{a \in \text{Ag}}$:*

- (1) If $\llbracket \varphi_{\text{VCG}}(\mathbf{BB}) \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v) = 1$, then $\text{act}_{\text{match}(v, \mathbf{BB}_{\eta_x})} = b_{\eta_x}(\mathbf{v})$ and $\text{act}_{\text{match}(v, \mathbf{BB}_{\eta_1})} > b_{\eta_2}(\mathbf{v})$;
- (2) If $m = 2$, then $\llbracket \text{FG}(\varphi_{\text{VCG}}(\mathbf{BB})) \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v_i) = 1$;
- (3) If $m \geq 3$, then $\llbracket \text{FG}(\varphi_{\text{VCG}}(\mathbf{BB})) \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v_i) \neq 1$.

Restricted BB The restricted balanced bidding strategy (RBB) [24] is a variation of BB in which the agent only targets slots that are no better than her current slot. The natural strategy representing RBB for agent a is denoted RBB_a and is constructed as follows. First, include the guarded actions $(RBB_{a,1}(b), b)$ for each action $b \in \text{Ac}$. Second, include $(RBB_{a,2}(b, s), b)$ for each $b \in \text{Ac}$ and $1 < s \leq m$. Finally, the last guarded action is $(\top, 0)$. Let $s_a = \min(m, \sum_{s' \in S} s' \times \text{all}_{a, s'})$ be the slot assigned to agent a or the last slot if there is no such slot. Define $RBB_{a,1}(b)$ and $RBB_{a,2}(b, s)$:

$$\begin{aligned} RBB_{a,1}(b) &:= b = \frac{\vartheta_a + p_1}{2} \wedge \text{argmax}_{s \in S \& s \geq s_a} (\text{ut}_{a, s}) = 1 \\ RBB_{a,2}(b, s) &:= \text{ut}_{a, s} = \theta_{s-1} \times (\vartheta_a - b) \\ &\wedge (\text{argmax}_{s' \in S \& s' \geq s_a} (\text{ut}_{a, s'}))^{-1} = s^{-1} \end{aligned}$$

Similar to the results in [24], we have that if all agents follow the restricted balanced-bidding strategy, the auction converge to the VCG equilibrium outcome. RBB always converge:

Proposition 4. *For any initial state $v_i \in V_i$, state $v \in V$, and $1 < x \leq n$, the following holds, where $\mathbf{v} = (\ell(v, \vartheta_a))_{a \in \text{Ag}}$:*

- (1) If $\llbracket \varphi_{\text{VCG}}(\mathbf{RBB}) \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v) = 1$, then $\text{act}_{\text{match}(v, \mathbf{RBB}_{\eta_x})} = b_{\eta_x}(\mathbf{v})$ and $\text{act}_{\text{match}(v, \mathbf{RBB}_{\eta_1})} > b_{\eta_2}(\mathbf{v})$;
- (2) $\llbracket \text{FG}(\varphi_{\text{VCG}}(\mathbf{RBB})) \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v_i) = 1$.

Knowledge grounded RBB The knowledge grounded RBB strategy (KBB) is a variation of RBB in which the agent uses her knowledge about the valuation of the player currently at her target slot to ground her bid value. The idea is to avoid bidding more than what she knows her opponent values the slot. The natural strategy representing KBB for agent a is denoted KBB_a and is constructed in three steps. First, include the guarded actions $(KBB_{a,1}(b, c, i), c)$ for each $b, c \in \text{Ac}$ and agent $i \neq a$. Second, include $(KBB_{a,2}(b, s, c, i), c)$ for each $b, c \in \text{Ac}$, slot $1 < s \leq m$ and agent $i \neq a$. Finally, include the guarded actions from RBB_a . The conditions $KBB_{a,1}(b, c, i)$ and $KBB_{a,2}(b, s, c, i)$ are defined as follows:

$$KBB_{a,1}(b, c, i) := K_a(RBB_{a,1}(b) \wedge \text{all}_{i,1} = 1 \wedge c = \min(\vartheta_i, b))$$

$$KBB_{a,2}(b, s, c, i) := K_a(RBB_{a,2}(b, s) \wedge \text{all}_{i,s} = 1 \wedge c = \min(\vartheta_i, b))$$

The prices under KBB are at most the same as under RBB:

Proposition 5. *For any state $v \in V$, slot $s \in S$ and agent a , $\llbracket (\text{Ag}, \mathbf{KBB})p_s \leq (\text{Ag}, \mathbf{RBB})p_s \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v) = 1$.*

Remark 2. With natural strategies, we can easily construct an strategy in which agent η_x plays according to $b_{\eta_x}(\mathbf{v})$ (for $1 < x \leq n$) and agent η_1 bids $b_{\eta_2} + \text{inc}$ when she knows others' valuations.

BB with recall Since BB may not converge to the VCG equilibrium outcome due to loops on the slot allocation and prices, we construct a strategy that behaves according to BB while there is no repetition in the outcome and follows RBB otherwise. Hereafter, we show that this strategy with recall prevents the loops that hinder the convergence of BB. Define the set of weighted conditions $\Psi = \{\bigwedge_{s \in S} (p_s = pr_s \wedge \bigwedge_{a \in \text{Ag}} \text{all}_{a,s} = al_{a,s}) : pr_s \in \text{Ac} \& al_{a,s} \in \{0, 1\}\}$. The natural strategy representing balanced bidding with recall for agent a is denoted BBR_a and is constructed as follows. First, include the guarded actions $(BBR_{a,1}(\psi, b), b)$ for each action $b \in \text{Ac}$ and condition $\psi \in \Psi$. Second, include $(BBR_{a,2}(\psi, b, s), b)$ for each $\psi \in \Psi$, $b \in \text{Ac}$ and $1 < s \leq m$. Third, include $(BBR_{a,3}(\psi, b), b)$ for each action $b \in \text{Ac}$. Fourth, include $(BBR_{a,4}(b, s), b)$ for each $b \in \text{Ac}$ and $1 < s \leq m$. Finally, the last guarded action is $(\top^*, 0)$.

Now we define each guarded condition in BBR_a . If the current allocation and payments have already happen in the past, a plays according to the restricted bidding strategy:

$$BBR_{a,1}(\psi, b) := \top^* \cdot \psi \cdot \top^* \cdot (\psi \wedge RBB_{a,1}(b))$$

$$BBR_{a,2}(\psi, b, s) := \top^* \cdot \psi \cdot \top^* \cdot (\psi \wedge RBB_{a,2}(b, s))$$

If there was no repetition on the payments and slot allocation, she plays according to the balanced bidding strategy:

$$BBR_{a,3}(\psi, b) := \top^* \cdot BB_{a,1}(b)$$

$$BBR_{a,4}(\psi, b, s) := \top^* \cdot BB_{a,2}(b, s)$$

When all agents follow the strategy profile $\mathbf{BBR} = (BBR_a)_{a \in \text{Ag}}$, the game converges to the VCG equilibrium outcome.

Proposition 6. *For any initial state $v_i \in V_i$, state $v \in V$, and $1 < x \leq n$, the following holds, where $\mathbf{v} = (\ell(v, \vartheta_a))_{a \in \text{Ag}}$:*

- (1) If $\llbracket \varphi_{\text{VCG}}(\mathbf{BBR}) \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v) = 1$, then $\text{act}_{\text{match}(v, \mathbf{BBR}_{\eta_x})} = b_{\eta_x}(\mathbf{v})$ and $\text{act}_{\text{match}(v, \mathbf{BBR}_{\eta_1})} > b_{\eta_2}(\mathbf{v})$;
- (2) $\llbracket \text{FG}(\varphi_{\text{VCG}}(\mathbf{BBR})) \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v_i) = 1$.

When other agents are inactive (i.e. they repeat their last action), if BBR_a selects a different bid from the one assigned by RBB_a , the utility of a in the next state is greater under BBR_a .

Proposition 7. *Let $\rho \in \{ir, iR\}$ and $v = \delta(v', c)$, for some state $v' \in V$ and action profile $c = (c)_{a \in \text{Ag}}$. Given an agent $a \in \text{Ag}$, let $\sigma_{-a}^\rho = (\sigma_i^\rho)_{i \in \text{Ag}_{-a}}$ be a ρ -strategy profile, where the strategy σ_i^ρ of agent i is such that $\text{act}_{\text{match}(v, \sigma_i^\rho)} = c_i$. If $\text{act}_{\text{match}(v, \mathbf{BBR}_a)} \neq \text{act}_{\text{match}(v, \mathbf{RBB}_a)}$, then $\llbracket (\text{Ag}_{-a}, \sigma_{-a}^{iR})(a, \mathbf{BBR}_a) \text{Xut}_a \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v) > \llbracket (\text{Ag}_{-a}, \sigma_{-a}^{ir})(a, \mathbf{RBB}_a) \text{Xut}_a \rrbracket^{\mathcal{G}_{\text{GSP}, ir}}(v)$.*

Remark 3. In vindictive bidding [78], the agent bids as high as possible to raise the payment of the advisor in the slot right below hers. Since there is the risk that a change in other agents' bids could result in paying a higher price than expected, the player could use memory to balance the use of aggressive bids.

5 EXPRESSIVITY

In relation to $\text{SL}[\mathcal{F}]$ with combinatorial strategies, $\text{NatSL}[\mathcal{F}]$ introduces a new, broader class of human-friendly strategies *and* a language for expressing properties of agents that use such strategies. Clearly, strategies with quantitative conditions can be used to obtain goals that would not be achievable otherwise. On the other hand, bounded natural strategies of $\text{NatSL}[\mathcal{F}]$ may not achieve some goals that can be enforced with combinatorial strategies of $\text{SL}[\mathcal{F}]$. In this section, we show that the expressive power of $\text{NatSL}[\mathcal{F}]$ is incomparable to that of $\text{SL}[\mathcal{F}]$. In other words, there are properties of quantitative games with natural strategies that cannot be equivalently translated to properties based on combinatorial strategies, and vice versa. From this, we conclude that reasoning about human-friendly strategies offers an inherently different view of a multi-agent system from the “standard” one.

5.1 Expressive and Distinguishing Power

We first adapt the notions of distinguishing power and expressive power to the quantitative case as follows⁵.

Definition 4 (Distinguishing power of real-valued logics). Let $\mathcal{L}_1 = (L_1, \llbracket \cdot \rrbracket_1)$ and $\mathcal{L}_2 = (L_2, \llbracket \cdot \rrbracket_2)$ be two logical systems with syntax L_1, L_2 and real-valued semantics $\llbracket \cdot \rrbracket_1, \llbracket \cdot \rrbracket_2$ over the same class of models \mathcal{M} . We say that \mathcal{L}_2 is *at least as distinguishing* as \mathcal{L}_1 (written: $\mathcal{L}_2 \leq_d \mathcal{L}_1$) iff for every pair of models $M, M' \in \mathcal{M}$, if there exists a formula $\varphi_1 \in L_1$ such that $\llbracket \varphi_1 \rrbracket_1^M \neq \llbracket \varphi_1 \rrbracket_1^{M'}$, then there is also $\varphi_2 \in L_2$ with $\llbracket \varphi_2 \rrbracket_2^M \neq \llbracket \varphi_2 \rrbracket_2^{M'}$. In other words, if there is a formula of \mathcal{L}_1 discerning M from M' , then there must be also a formula of \mathcal{L}_2 doing the same.

Definition 5 (Expressive power of real-valued logics). \mathcal{L}_2 is *at least as expressive* as \mathcal{L}_1 (written: $\mathcal{L}_2 \leq_e \mathcal{L}_1$) iff for every $\varphi_1 \in L_1$ there exists $\varphi_2 \in L_2$ such that, for every model $M \in \mathcal{M}$, we have $\llbracket \varphi_1 \rrbracket_1^M = \llbracket \varphi_2 \rrbracket_2^M$. In other words, every formula of \mathcal{L}_1 has a translation in \mathcal{L}_2 that produces exactly the same truth values on models in \mathcal{M} .

It is easy to see that $\mathcal{L}_1 \leq_e \mathcal{L}_2$ implies $\mathcal{L}_1 \leq_d \mathcal{L}_2$. Thus, by transposition, we also get that $\mathcal{L}_1 \not\leq_d \mathcal{L}_2$ implies $\mathcal{L}_1 \not\leq_e \mathcal{L}_2$.

In the remainder, \mathcal{M} is the class of pointed weighted games, i.e., pairs (\mathcal{G}, v) where \mathcal{G} is a wCGS and v is a state in \mathcal{G} .

5.2 Expressivity of $\text{NatSL}[\mathcal{F}]$ vs. $\text{SL}[\mathcal{F}]$

$\text{NatSL}[\mathcal{F}]$ and $\text{SL}[\mathcal{F}]$ are based on different notions of strategic ability. The former refers to “natural” strategies, represented as mappings from regular expressions over atomic propositions to actions. The latter uses “combinatorial” strategies, represented by mappings from sequences of states to actions. Each natural strategy can be translated to a combinatorial one, but not vice versa. Consequently, $\text{SL}[\mathcal{F}]$ can express that a given coalition has a combinatorial strategy to achieve their goal (which is not expressible in $\text{NatSL}[\mathcal{F}]$). On the other hand, $\text{NatSL}[\mathcal{F}]$ allows expressing that a winning natural strategy does not exist (which cannot be captured in $\text{SL}[\mathcal{F}]$). Now we show that $\text{NatSL}[\mathcal{F}]$ allows to express properties that cannot be captured in $\text{SL}[\mathcal{F}]$, and vice versa.

Proposition 8. $\text{NatSL}[\mathcal{F}] \not\leq_d \text{SL}[\mathcal{F}]$ in both *ir* and *iR* semantics.

⁵Cf., e.g., [74] for a detailed discussion of standard notions of expressivity.




Figure 1: Model \mathcal{G}_1 . Its counterpart \mathcal{G}'_1 is obtained by fixing p to hold only in q_1, q'_1 . Underscore fits any action label

PROOF SKETCH. Consider model \mathcal{G}_1 in Figure 1, with agents $\text{Ag} = \{1, 2\}$, actions $\text{Ac}_1 = \{a_1, b_1, c_1, d_1, e_1\}$ and $\text{Ac}_2 = \{a_2, b_2\}$ available at all positions, and propositions $\text{AP} = \{p, \text{win}\}$. Both propositions are qualitative (that is, the propositions have only values in $\{-1, 1\}$). For each proposition, the states where it evaluates to 1 are indicated; otherwise its truth value is assumed to be -1 . The outgoing transitions in q'_1, q''_1 (resp. q'_2) are exact copies of those at q_1 (resp. q_2). Moreover, model \mathcal{G}'_1 is obtained by fixing proposition p to hold only in q_1, q'_1 , but not in q''_1 . As all the propositions are qualitative, formulas of $\text{NatSL}[\mathcal{F}]$ and $\text{SL}[\mathcal{F}]$ evaluate to -1 or 1 . Note also that the sets of *ir* and *iR* strategies in each model coincide, so we can concentrate on the *ir* case w.l.o.g.

Let $\mathcal{G} \dagger \sigma$ denote the model obtained by fixing the (memoryless) strategy σ in \mathcal{G} . In order to prove that (\mathcal{G}_1, q_0) and (\mathcal{G}'_1, q_0) satisfy the same formulas of $\text{SL}[\mathcal{F}]$, it suffices to observe that:

- (1) For every strategy σ_1 of agent 1 in \mathcal{G}_1 , there is σ'_1 in \mathcal{G}'_1 such that agent 2 has the same strategic abilities in $(\mathcal{G}'_1 \dagger \sigma_1, q_0)$ and $(\mathcal{G}'_1 \dagger \sigma'_1, q_0)$ (and vice versa). For instance, playing c_1 in \mathcal{G} obtains the same abilities of 1 as playing a_1 in \mathcal{G}' .
- (2) Analogously for strategies of agent 2, e.g., strategy $a_2 a_2 a_2 b_2 b_2$ in \mathcal{G}_1 can be simulated by strategy $a_2 a_2 b_2 b_2 b_2$ in \mathcal{G}'_1 .

On the other hand, the formula $\exists s_2 \leq 2 \forall s_1 \leq 1 (1, s_1)(2, s_2) \text{Fwin}$ of $\text{NatSL}[\mathcal{F}]$ holds in (\mathcal{G}_1, q_0) , but not in (\mathcal{G}'_1, q_0) . The winning natural strategy for agent 2 in \mathcal{G}_1 is $((\tau^* p, a_2), (\tau^*, b_2))$; clearly, it does not succeed in \mathcal{G}'_1 . \square

Proposition 9. $\text{SL}[\mathcal{F}] \not\leq_d \text{NatSL}[\mathcal{F}]$ in both *ir* and *iR* semantics.

PROOF SKETCH. Consider models \mathcal{G}_2 and \mathcal{G}'_2 in Figure 2. They have isomorphic action/transition structures, the only difference being the indistinguishability of states q_1, q_2 in \mathcal{G}'_2 (but not in \mathcal{G}_2). Since the two states have the same valuations of propositions, each natural strategy must specify the same decision in q_1, q_2 . Thus, both players have exactly the same available natural strategies in \mathcal{G}_2 and \mathcal{G}'_2 , and hence (\mathcal{G}_2, q_0) and (\mathcal{G}'_2, q_0) produce the same valuations of $\text{NatSL}[\mathcal{F}]$ formulas.

On the other hand, we have that $\exists s_2 \forall s_1 (1, s_1)(2, s_2) \text{Fwin}$ of $\text{SL}[\mathcal{F}]$ holds in (\mathcal{G}_2, q_0) , but not in (\mathcal{G}'_2, q_0) . \square




Figure 2: Models \mathcal{G}_2 (left) and \mathcal{G}'_2 (right)

The following is an immediate consequence.

Theorem 1. *NatSL[\mathcal{F}] and SL[\mathcal{F}] have incomparable distinguishing power over the class of pointed wCGS (in both ir and iR semantics).*

Corollary 2. *NatSL[\mathcal{F}] and SL[\mathcal{F}] have incomparable expressive power over the class of pointed wCGS (in both ir and iR semantics).*

6 MODEL CHECKING

In this section we show that the model checking problem for NatSL[\mathcal{F}] with imperfect information is no harder than model checking LTL or classic SL with memoryless agents. First of all, we define the quantitative model-checking problem for NatSL[\mathcal{F}].

Definition 6. Given $\rho \in \{ir, iR\}$, the model-checking problem for NatSL[\mathcal{F}] consists in deciding, for a given sentence φ , wCGS \mathcal{G} , state $v \in V$ and predicate $P \subseteq [-1; 1]$, whether $\llbracket \varphi \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) \in P$.

Now, we have all the ingredients to prove the following result.

Theorem 2. *Assuming that functions in \mathcal{F} can be computed in polynomial space, model checking NatSL[\mathcal{F}] with imperfect information, natural strategies with recall, and k as parameter of the problem is PSPACE-complete.*

PROOF. For the lower-bound we recall that LTL[\mathcal{F}] model checking is PSPACE-complete [7]. For the upper-bound, to verify that a given NatSL[\mathcal{F}] formula φ is satisfied over a wCGS \mathcal{G} at a state $v \in V$ under assignments χ over uniform natural strategies with recall, we make use of a recursive function as is done in [26]. We start by showing that each recursive call only needs at most polynomial space. First, observe that each assignment χ has a strategy s_a for each agent $a \in Ag$ ⁶. We know that each strategy s_a that can be assigned to agent a is bounded, and we have that $compl(s_a) \leq k$. Thus, each strategy can be stored in $O(k \cdot |Act|)$ and, by consequence, any assignment can be stored in space $O((|Ag| \cdot |free(\varphi)|) \cdot (k \cdot |Act|))$. Now, we can analyse the recursive function. For the base case, $\llbracket p \rrbracket_{\chi}^{\mathcal{G}, \rho}(v)$ can be computed in constant space via the weight function. For strategy quantification $\llbracket \exists s_a^{\leq k} . \varphi \rrbracket_{\chi}^{\mathcal{G}, \rho}(v)$, besides the recursive call to $\llbracket \varphi \rrbracket_{\chi[s_a \rightarrow \sigma]}^{\mathcal{G}, \rho}(v)$ we need space $O(|k| \cdot |Ac|)$ to store the current strategy and the current maximum value computed. For $\llbracket f(\varphi_1, \dots, \varphi_m) \rrbracket_{\chi}^{\mathcal{G}, \rho}(v)$, by assumption f is computed in polynomial

⁶Note that, as defined in Section 3, we consider only complete assignments. Thus, we can assume that a strategy is assigned for each agent.

space. For $\llbracket X\varphi \rrbracket_{\chi}^{\mathcal{G}, \rho}(v)$, we only need to observe that the next state

in $\text{Out}(v, \chi)$ is computed in constant space. Finally, we detail how $\llbracket \varphi_1 \cup \varphi_2 \rrbracket_{\chi}^{\mathcal{G}, \rho}(v)$ is computed. Let $\pi = \text{Out}(v, \chi)$. Since \mathcal{G} has finitely many states, there exist two indices $g < l$ such that $\pi_g = \pi_l$, and since strategies are bounded by k , the suffix of π starting at index l is equal to the suffix starting at index g . So there exist $\rho_1 = v_0 \dots v_{g-1}$ and $\rho_2 = v_g \dots v_{l-1}$ such that $\pi = \rho_1 \cdot \rho_2^\omega$. It follows that

$$\begin{aligned} \llbracket \varphi_1 \cup \varphi_2 \rrbracket_{\chi}^{\mathcal{G}, \rho}(v) &= \sup_{i \geq 0} \min \left(\llbracket \varphi_2 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_i), \min_{0 \leq j < i} \llbracket \varphi_1 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_j) \right) \\ &= \max_{0 \leq i < l} \min \left(\llbracket \varphi_2 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_i), \min_{0 \leq j < i} \llbracket \varphi_1 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_j) \right) \end{aligned}$$

This can be computed by a while loop that increases i , computes $\llbracket \varphi_2 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_i)$ and $\min_{0 \leq j < i} \llbracket \varphi_1 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_j)$, their minimum, and records the result if it is bigger than the previous maximum. This requires to store the current value of $\min_{0 \leq j < i} \llbracket \varphi_1 \rrbracket_{\chi}^{\mathcal{G}, \rho}(\pi_j)$, the current maximum, and the list of states already visited, which are at most $k \cdot |V|$. Finally, the number of nested recursive calls is at most $|\varphi|$, so the total space needed is bounded by $|\varphi|$ times a polynomial in the size of the input, and is thus polynomial. \square

Since memoryless natural strategies are a special case of natural strategies with recall, we obtain the following result.

Corollary 3. *Assuming that functions in \mathcal{F} can be computed in polynomial space, model checking NatSL[\mathcal{F}] with imperfect information, memoryless natural strategies, and k as parameter of the problem is PSPACE-complete.*

7 CONCLUSION

In this work we have introduced Natural Strategy Logic with quantitative semantics and imperfect information (NatSL[\mathcal{F}]) for reasoning about strategic ability in auctions. NatSL[\mathcal{F}] provides a tool for mechanism design and offers a new perspective for formal verification and design of novel mechanisms and strategies. We demonstrated the usefulness of our approach by modelling and evaluating strategies for repeated keyword auctions.

In terms of technical results, we proved that the model checking problem for NatSL[\mathcal{F}] is PSPACE-complete, that is, no harder than model checking for the much less expressive language of quantitative LTL (LTL[\mathcal{F}]). We also showed that NatSL[\mathcal{F}] has incomparable distinguishing and expressive power to SL[\mathcal{F}]. This means that the characterizations based on simple bounded strategies offer an inherently different view of auctions and mechanism design from characterizations using combinatorial strategies of arbitrary complexity. Amazingly, this aspect has never been studied for natural strategies, not even for the original proposal of NatATL [46].

We consider several directions for future work. First, a probabilistic extension of Strategy Logic [9] would allow handling mechanisms in stochastic settings with mixed strategies. Another direction is to investigate the use of strategies with recall for learning other players' valuations based on their behaviour. Finally, the implementation of a model checker for NatSL[\mathcal{F}] will enable the empirical evaluation of auctions with natural strategies.

ACKNOWLEDGMENTS

This research has been supported by the ANR project AGAPE ANR-18-CE23-0013 and the FNR-CORE/NCBR-POLLUX project STV (POLLUX-VII/1/2019).

REFERENCES

- [1] T. Ágotnes. 2006. Action and Knowledge in Alternating-time Temporal Logic. *Synthese* 149, 2 (2006), 377–409.
- [2] T. Ágotnes and D. Walther. 2009. A Logic of Strategic Ability Under Bounded Memory. *Journal of Logic, Language and Information* 18, 1 (2009), 55–77.
- [3] N. Alechina, M. Dastani, B. Logan, and J.-J. Ch. Meyer. 2007. A Logic of Agent Programs. In *Proceedings of AAIL*. AAAI Press, Vancouver, 795–800.
- [4] N. Alechina, B. Logan, M. Dastani, and J.-J. Ch. Meyer. 2008. Reasoning about agent execution strategies. In *Proceedings of International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, Estoril, 1455–1458.
- [5] N. Alechina, B. Logan, N.H. Nga, and A. Rakib. 2009. A Logic for Coalitions with Bounded Resources. In *Proc. of International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, Pasadena, 659–664.
- [6] N. Alechina, B. Logan, H.N. Nguyen, and A. Rakib. 2010. Resource-Bounded Alternating-Time Temporal Logic. In *Proceedings of International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, Toronto, 481–488.
- [7] Shaull Almagor, Udi Boker, and Orna Kupferman. 2016. Formally Reasoning about Quality. *J. ACM* 63, 3 (Sept. 2016), 24:1–24:56. <https://doi.org/10.1145/2875421>
- [8] Rajeev Alur, Thomas A Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *Journal of the ACM* 49, 5 (2002), 672–713.
- [9] Benjamin Aminof, Marta Kwiatkowska, Bastien Maubert, Aniello Murano, and Sasha Rubin. 2019. Probabilistic Strategy Logic. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI, Sarit Kraus (Ed.)*. AAAI Press, Macao, 32–38.
- [10] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of Model Checking (Representation and Mind Series)*. The MIT Press, Cambridge.
- [11] M. Barlo, G. Carmona, and H. Sabourian. 2008. Bounded memory with finite action spaces. *Sabanci University, Universidade Nova de Lisboa and University of Cambridge* 1, 1 (2008).
- [12] Gilles Barthe, Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Pierre-Yves Strub. 2016. Computer-Aided Verification for Mechanism Design. In *The 12th Conference on Web and Internet Economics (WINE 2016) (Lecture Notes in Computer Science, Vol. Volume 10123)*. Springer, Montreal, Canada, 279 – 293. https://doi.org/10.1007/978-3-662-54110-4_20
- [13] F. Belardinelli, R. Condurache, C. Dima, W. Jamroga, and A. V. Jones. 2017. Bisimulations for Verifying Strategic Abilities with an Application to ThreeBallot. In *Proc. of AAMAS 17*. Int. Foundation for Autonomous Agents and Multiagent Systems, Richland, 1286–1295.
- [14] Francesco Belardinelli, Alessio Lomuscio, and Vadim Malvone. 2018. Approximating Perfect Recall When Model Checking Strategic Abilities. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, Michael Thielscher, Francesca Toni, and Frank Wolter (Eds.). AAAI Press, Tempe, 435–444. <https://aaai.org/ocs/index.php/KR/KR18/paper/view/18010>
- [15] F. Belardinelli, A. Lomuscio, and F. Patrizi. 2014. Verification of Agent-based Artifact Systems. *Journal of Artificial Intelligence Research* 51 (2014), 333–376.
- [16] R. Bordini, M. Fisher, W. Visser, and M. Wooldridge. 2006. Verifying Multi-Agent Programs by Model Checking. *Autonomous Agents and Multi-Agent Systems* 12, 2 (2006), 239–256.
- [17] R. H. Bordini, M. Wooldridge, and J. F. Hübner. 2007. *Programming Multi-Agent Systems in AgentSpeak using Jason*. John Wiley & Sons, West Sussex.
- [18] L. E. Bourne. 1970. Knowing and using concepts. *Psychol. Rev.* 77 (1970), 546–556. <https://doi.org/10.1037/h0030000>
- [19] Patricia Bouyer, Orna Kupferman, Nicolas Markey, Bastien Maubert, Aniello Murano, and Giuseppe Perelli. 2019. Reasoning about Quality and Fuzziness of Strategic Behaviours. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*. AAAI Press, Macao, 1588–1594.
- [20] N. Bulling and B. Farwer. 2010. Expressing Properties of Resource-Bounded Systems: The Logics RTL* and RTL. In *Proceedings of Computational Logic in Multi-Agent Systems (CLIMA)*. Springer-Verlag, Berlin, 22–45.
- [21] N. Bulling and B. Farwer. 2010. On the (Un-)Decidability of Model Checking Resource-Bounded Agents. In *Proceedings of ECAI (Frontiers in Artificial Intelligence and Applications, Vol. 215)*. IOS Press, Amsterdam, 567–572.
- [22] N. Bulling and W. Jamroga. 2014. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Journal of Autonomous Agents and Multi-Agent Systems* 28, 3 (2014), 474–518.
- [23] Marco B. Caminati, Manfred Kerber, Christoph Lange, and Colin Rowat. 2015. Sound Auction Specification and Implementation. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation (Portland, Oregon, USA) (EC '15)*. Association for Computing Machinery, New York, NY, USA, 547–564. <https://doi.org/10.1145/2764468.2764511>
- [24] Matthew Cary, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. 2007. Greedy Bidding Strategies for Keyword Auctions. In *Proceedings of the 8th ACM Conference on Electronic Commerce (San Diego, California, USA) (EC '07)*. Association for Computing Machinery, New York, NY, USA, 262–271. <https://doi.org/10.1145/1250910.1250949>
- [25] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano. 2014. MCMAS-SLK: A Model Checker for the Verification of Strategy Logic Specifications. In *Proc. of the Int. Conf. on Computer Aided Verification (CAV) (Lecture Notes in Computer Science, Vol. 8559)*. Springer, Berlin, 525–532.
- [26] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano. 2018. Practical verification of multi-agent systems against SLK specifications. *Inf. Comput.* 261 (2018), 588–614.
- [27] Krishnendu Chatterjee, Thomas A Henzinger, and Nir Piterman. 2010. Strategy logic. *Information and Computation* 208, 6 (2010), 677–693.
- [28] M. Dastani and W. Jamroga. 2010. Reasoning about Strategies of Multi-Agent Programs. In *Proceedings of AAMAS*. Int. Foundation for Autonomous Agents and Multiagent Systems, Richland, 625–632.
- [29] E. Davis and G. Marcus. 2015. Commonsense Reasoning. *Commun. ACM* 58, 9 (2015), 92–103. <https://doi.org/10.1145/2701413>
- [30] A. Deutsch, R. Hull, F. Patrizi, and V. Vianu. 2009. Automatic Verification of Data-centric Business Processes. In *Proceedings of the 12th International Conference on Database Theory (ICDT09)*. Association for Computing Machinery, New York, 252–267.
- [31] C. Dima and F.L. Tiplea. 2011. Model-checking ATL under Imperfect Information and Perfect Recall Semantics is Undecidable. *CoRR* abs/1102.4225 (2011), 17.
- [32] C. Dixon, A. Winfield, M. Fisher, and C. Zeng. 2012. Towards temporal verification of swarm robotic systems. *Robotics and Autonomous Systems* 60, 11 (2012), 1429–1441.
- [33] H. Duijf and J.M. Broersen. 2016. Representing Strategies. In *Proc. of Int. Workshop on Strategic Reasoning (SR)*. Open Publishing Association, New York, 15–26. <https://doi.org/10.4204/EPTCS.218.2>
- [34] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.
- [35] J. Feldman. 2000. Minimization of Boolean complexity in human concept learning. *Nature* 407 (11 2000), 630–3. <https://doi.org/10.1038/35036586>
- [36] P. Gammie and R. van der Meyden. 2004. MCK: Model Checking the Logic of Knowledge. In *Proc. of 16th Int. Conf. on Computer Aided Verification (CAV) (Lecture Notes in Computer Science, Vol. 3114)*. Springer, Berlin, 479–483.
- [37] M. Ghallab, D. Nau, and P. Traverso. 2004. *Automated Planning: Theory and Practice*. Morgan Kaufmann, San Francisco.
- [38] P. Gonzalez, A. Griesmayer, and A. Lomuscio. 2015. Verification of GSM-based Artifact-centric Systems by Predicate Abstraction. In *Proceedings of the 13th International Conference on Service Oriented Computing (ICSOC15) (Lecture Notes in Computer Science, Vol. 9435)*. Springer, Berlin, 253–268.
- [39] Anshul Gupta, Sven Schewe, and Dominik Wojtczak. 2015. Making the Best of Limited Memory in Multi-Player Discounted Sum Games. *Electronic Proceedings in Theoretical Computer Science* 193 (2015), 16–30. <https://doi.org/10.4204/EPTCS.193.2>
- [40] D. Harel and D. Kozen. 1982. Process Logic: Expressiveness, Decidability, Completeness. *J. Comput. System Sci.* 25, 2 (1982), 144–170.
- [41] A. Herzig, E. Lorini, F. Maffre, and D. Walther. 2014. Alternating-time Temporal Logic with Explicit Programs. In *Proceedings of Workshop on Logical Aspects of Multi-Agent Systems (LAMAS)*. IFAAMAS, Paris, 15.
- [42] K.V. Hindriks, F.S. de Boer, W. van der Hoek, and J.-J. Ch. Meyer. 1999. Agent Programming in 3APL. *Autonomous Agents and Multi-Agent Systems* 2, 4 (1999), 357–401. <https://doi.org/10.1023/A:1010084620690>
- [43] J. Hörner and W. Olszewski. 2009. How robust is the Folk Theorem? *The Quarterly Journal of Economics* 124, 4 (2009), 1773–1814.
- [44] W. Jamroga and T. Ágotnes. 2007. Constructive knowledge: what agents can achieve under imperfect information. *J. Applied Non-Classical Logics* 17, 4 (2007), 423–475.
- [45] Wojciech Jamroga, Damian Kurpiewski, and Vadim Malvone. 2020. Natural Strategic Abilities in Voting Protocols. In *Socio-Technical Aspects in Security and Trust - 10th International Workshop, STAST 2020, Virtual Event, September 14, 2020 (Lecture Notes in Computer Science, Vol. 12812)*, Thomas Groß and Luca Viganò (Eds.). Springer, Berlin, 45–62. https://doi.org/10.1007/978-3-030-79318-0_3
- [46] Wojciech Jamroga, Vadim Malvone, and Aniello Murano. 2019. Natural strategic ability. *Artificial Intelligence* 277 (2019), 103170. <https://doi.org/10.1016/j.artint.2019.103170>
- [47] Wojciech Jamroga, Vadim Malvone, and Aniello Murano. 2019. Natural Strategic Ability under Imperfect Information. In *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. Int. Foundation for Autonomous Agents and Multiagent Systems, Richland, 962–970.
- [48] M. Kerber, C. Lange, and C. Rowat. 2016. An introduction to mechanized reasoning. *Journal of Mathematical Economics* 66 (2016), 26 – 39.
- [49] P. Kouvaros and A. Lomuscio. 2015. A Counter Abstraction Technique for the Verification of Robot Swarms. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI15)*. AAAI Press, Austin, 2081–2088.
- [50] O. Kupferman and M. Y. Vardi. 2000. Synthesis with incomplete informatio. In *Advances in Temporal Logic*. Springer, Berlin, 109–127.

- [51] A. Lomuscio, H. Qu, and F. Raimondi. 2015. MCMAS: A Model Checker for the Verification of Multi-Agent Systems. *Software Tools for Technology Transfer* 19 (2015), 9–30. <https://doi.org/10.1007/s10009-015-0378-x> <http://dx.doi.org/10.1007/s10009-015-0378-x>.
- [52] Yoshihiro Maruyama. 2021. A Reasoning System for Fuzzy Distributed Knowledge Representation in Multi-Agent Systems. In *2021 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*. IEEE, New York, 1–6.
- [53] Bastien Maubert, Munyque Mittelmann, Aniello Murano, and Laurent Perrussel. 2021. Strategic Reasoning in Automated Mechanism Design. In *Proc. of the Eighteen Conference on Principles of Knowledge Representation and Reasoning*. IJCAI Organization, Online event, 10.
- [54] F. Mogavero, A. Murano, G. Perelli, and M. Vardi. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Trans. Comput. Log.* 15, 4 (2014), 47.
- [55] Marco Montali, Diego Calvanese, and Giuseppe De Giacomo. 2014. Verification of data-aware commitment-based multiagent system. In *Proc. of the 14th International Conference on Autonomous Agents and Multi-Agent systems (AAMAS14)*. Int. Foundation for Autonomous Agents and Multiagent Systems, Richland, 157–164.
- [56] R. Morris and G. Ward. 2014. *The Cognitive Psychology of Planning*. Psychology Press, New York.
- [57] J. Nielsen. 1994. *Usability Engineering*. Morgan Kaufmann, San Francisco.
- [58] N.J. Nilsson. 1994. Telemo-Reactive Programs for Agent Control. *Journal of Artificial Intelligence Research* 1 (1994), 139–158. <https://doi.org/10.1613/jair.30>
- [59] P. Novák and W. Jamroga. 2009. Code Patterns for Agent Oriented Programming. In *Proceedings of AAMAS'09*. Int. Foundation for Autonomous Agents and Multiagent Systems, Richland, 105–112.
- [60] M. Pauly and Rohit Parikh. 2003. Game logic-an overview. *Studia Logica* 75, 2 (2003), 165–182.
- [61] M. Pauly and M. Wooldridge. 2003. Logic for mechanism design—a manifesto. In *Proc. of the 2003 Workshop on Game Theory and Decision Theory in Agent Systems*. ACM, Melbourne.
- [62] R. Pfeifer and C. Scheier. 1999. *Understanding Intelligence*. MIT Press, Cambridge.
- [63] A. Pnueli and R. Rosner. 1989. On the Synthesis of a Reactive Module.. In *Symposium on the Principles of Programming Languages (POPL 89)*. ACM, New York, 179–190.
- [64] J. H. Reif. 1984. The Complexity of Two-Player Games of Incomplete Information. *J. Comput. Syst. Sci.* 29, 2 (1984), 274–301.
- [65] Tim Roughgarden. 2010. Algorithmic game theory. *Commun. ACM* 53, 7 (2010), 78–86.
- [66] S. Russel and P. Norvig. 1995. *Artificial Intelligence: a Modern Approach*. Prentice Hall, Upper Saddle River.
- [67] F.P. Santos. 2018. *Dynamics of Reputation and the Self-organization of Cooperation*. Ph.D. Dissertation. University of Lisbon.
- [68] F.P. Santos, F.C. Santos, and J.M. Pacheco. 2018. Social norm complexity and past reputations in the evolution of cooperation. *Nature* 555 (2018), 242–245.
- [69] N. Troquard, W. van der Hoek, and M. Wooldridge. 2011. Reasoning about Social Choice Functions. *Journal of Philosophical Logic* 40, 4 (2011), 473–498.
- [70] W. van der Hoek, W. Jamroga, and M. Wooldridge. 2005. A Logic for Strategic Reasoning. In *Proceedings of AAMAS'05*. Association for Computing Machinery, New York, 157–164.
- [71] Hal R Varian. 2007. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.
- [72] S. Vester. 2013. Alternating-time temporal logic with finite-memory strategies. In *Proceedings of GandALF (EPTCS)*. Open Publishing Association, Borca di Cadore, 194–207. <https://doi.org/10.4204/EPTCS.119.17>
- [73] D. Walther, W. van der Hoek, and M. Wooldridge. 2007. Alternating-time Temporal Logic with Explicit Strategies. In *Proceedings TARK XI*. Association for Computing Machinery, New York, 269–278.
- [74] Y. Wang and F. Dechesne. 2009. On expressive power and class invariance. *CoRR* abs/0905.4332 (2009).
- [75] M. Wooldridge, T. Agotnes, P. Dunne, and W. Van der Hoek. 2007. Logic for automated mechanism design—a progress report. In *Proc. of AAAI*. AAAI Press, Vancouver, 9–17.
- [76] N. Yadav and S. Sardiña. 2012. Reasoning about Agent Programs Using ATL-Like Logics. In *Proceedings of JELIA*. Springer Berlin Heidelberg, Berlin, 437–449.
- [77] Yong Yuan, Fei-Yue Wang, and Daniel Zeng. 2017. Competitive Analysis of Bidding Behavior on Sponsored Search Advertising Markets. *IEEE Transactions on Computational Social Systems* 4, 3 (2017), 179–190. <https://doi.org/10.1109/TCSS.2017.2730925>
- [78] Yunhong Zhou and Rajan Lukose. 2007. Vindictive bidding in keyword auctions. *ACM International Conference Proceeding Series* 258 (2007), 141–146. <https://doi.org/10.1145/1282100.1282130>