

Equilibria for Games with Combined Qualitative and Quantitative Objectives

Julian Gutierrez · Aniello Murano · Giuseppe Perelli · Sasha Rubin · Thomas Steeples · Michael Wooldridge

Received: date / Accepted: date

Abstract The overall aim of our research is to develop techniques to reason about the equilibrium properties of multi-agent systems. We model multi-agent systems as concurrent games, in which each player is a process that is assumed to act independently and strategically in pursuit of personal preferences. In this article, we study these games in the context of finite-memory strategies, and we assume players' preferences are defined by a qualitative

We gratefully acknowledge the financial support of ERC Advanced Investigator grant 291528 at Oxford (J. Gutierrez, G. Perelli, and M. Wooldridge), INdAM research project 2017 "Logica e Autonomi per il Model Checking" at Naples (A. Murano), Marie Curie Fellowship of the Istituto Nazionale di Alta Matematica (S. Rubin), the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems EP/L015897/1 and the Ian Palmer Memorial Scholarship (T. Steeples). A preliminary version of this work appeared in [27].

J. Gutierrez
Monash University
E-mail: julian.gutierrez@monash.edu

A. Murano
University of Naples "Federico II"
E-mail: murano@na.infn.it

G. Perelli
Sapienza University of Rome
E-mail: perelli@diag.uniroma1.it

S. Rubin
University of Sydney
E-mail: sasha.rubin@sydney.edu.au

T. Steeples
University of Oxford
E-mail: thomas.steeples@cs.ox.ac.uk

M. Wooldridge
University of Oxford
E-mail: michael.wooldridge@cs.ox.ac.uk

and a quantitative objective, which are related by a lexicographic order: a player first prefers to satisfy its qualitative objective (given as a formula of Linear Temporal Logic) and then prefers to minimise costs (given by a mean-payoff function). Our main result is that deciding the existence of a strict ϵ Nash equilibrium in such games is 2EXPTIME-complete (and hence decidable), even if players' deviations are implemented as infinite-memory strategies.

Keywords Multi-agent systems · Multi-player games · Nash equilibrium · Linear Temporal logic · Mean-payoff games · Concurrent game structures.

1 Introduction

The last twenty years have seen considerable research directed at the use of game theoretic techniques in the analysis and verification of multi-agent systems [40]. From this standpoint, agents/processes in a multi-agent system can be understood as players in a game played on a directed graph (a transition system), acting strategically and independently in pursuit of their preferences. In this setting, possible behaviours of agents correspond to the strategies of players. One important strand of work in this tradition has been the development of techniques for reasoning about what properties players (or coalitions of players) can bring about (*i.e.*, whether they have “winning strategies” for certain conditions) [4]. Recently, attention has begun to shift from the analysis of strategic ability to the analysis of the *equilibrium properties* of such systems. A typical question in this setting is whether a particular temporal property will hold under the assumption that players select strategies that collectively form a Nash equilibrium [35].

A fundamental question in this work is how the preferences of agents are represented. One widely-adopted answer to this question is to associate with each player a *qualitative* goal (objective), usually given either by a temporal logic formula or else by a winning (acceptance) condition, such as reachability, safety, Büchi, Linear Temporal Logic, etc. [38, 24, 9, 25]. This approach is closely related to the verification of finite-state systems, and the model checking paradigm in particular [17]. However, the preference structures that are induced in this way have a rather simple (dichotomous) structure: a player is simply either satisfied or unsatisfied; no distinction is made between outcomes that satisfy the player's objective, nor is any made between outcomes that do not satisfy the objective. This limits the applicability of such representations for modelling many situations of interest. An alternative setting is given by games where, instead of having a qualitative objective, players have *quantitative* goals — for instance, to minimize a given cost, or to maximise some reward [20, 43]. Yet a third possibility, also the focus in this paper, is to use preference models that *combine* qualitative and quantitative objectives [15, 7, 46].

We consider goals given by a lexicographic order, where a player's primary goal is to satisfy its qualitative objective (given by a formula of Linear Temporal Logic, LTL [18]), and a player's secondary goal is to minimise its costs (where costs are given by a quantitative mean-payoff objective). This approach has several advantages. The qualitative objective can be used to express desirable properties on the states of the system, as is standard practice in the specification and verification of reactive systems [21, 17, 18]. For instance, LTL formulae, which we use to express the qualitative objective of agents, can be used to specify in a natural way that an agent prefers not to enter a given set of states (formally expressed as a safety property) or that an agent desires to eventually visit a given state of the system (for instance to model the termination of a task) [18]. In these cases, we can simply understand agents as “rational processes” within a reactive system. The quantitative objective can be used to

restrict the behaviour of such agents to those that are locally optimal for each agent. Thus, not only do we want an agent to accomplish its goal, but to do so *as efficiently as possible*, that is, such that the cost of performing the task is kept to a minimum. This latter type of preference is very naturally captured by *mean-payoff* specifications, even in cases where infinite behaviour is considered. Moreover, the combination of qualitative and quantitative objectives is natural for situations in which agents aim to satisfy some goal while minimising costs. For example, consider a robot whose task is to deliver packages around a factory environment: the primary goal of the robot is to deliver the packages (a qualitative objective readily expressible in LTL), while the secondary goal is to minimise fuel consumption when achieving this task (a quantitative objective that can be naturally expressed with a mean-payoff function). Scenarios like these are ubiquitous in embedded and cyber-physical systems [3].

The main solution concept we use in this paper is *strict ϵ Nash equilibrium* [40]. The use of Nash equilibrium — where no player in the game can unilaterally change their strategy and be better off as a consequence — is readily justified by the fact that this is the best-known and most widely-used solution concept for non-cooperative games. The use of strict ϵ Nash equilibrium is less common. Informally, by strict we mean that any possible unilateral deviation of a player results in an outcome that is strictly worse for that player. As such, this solution concept is more stable than “ordinary” Nash equilibrium since each player has less incentive to change strategy. Thus, from a stability point of view, strict Nash equilibrium is a desirable feature.¹ However, as expected, a game may have more Nash equilibria than strict Nash equilibria (and every strict Nash equilibrium is already an ordinary Nash equilibrium). In contrast, allowing for an ϵ Nash equilibrium, with $\epsilon > 0$, may lead to games with more equilibria. Informally, in an ϵ Nash equilibrium no player can unilaterally change their strategy and achieve a payoff that is at least as good as ϵ more than the one already obtained in the Nash equilibrium. As a consequence, every ordinary Nash equilibrium is also an ϵ Nash equilibrium (for all $\epsilon > 0$), but the converse may fail. Then, while the strict variant of Nash equilibria can decrease the number of equilibria in a game, the ϵ variant may increase it.

Contributions We consider games in which preferences are defined by a lexicographic order of goals given by an LTL formula (the primary goal) and a mean-payoff condition (the secondary goal). We prove that deciding the existence of a finite-state strict ϵ Nash equilibrium is in 2ExpTime .² This result subsumes the case for ordinary Nash equilibrium in two-player zero-sum games with LTL goals, which is known to be 2ExpTime -hard. Thus the above problem is 2ExpTime -complete. To obtain this result we introduce a reduction to a similar (though doubly exponentially larger) game where, instead of goals given by LTL formulae, goals are given by a parity acceptance condition, and we show how to solve such games in NP .

Our results also show how to solve the rational synthesis problem [22] and the rational verification problem [47,25] within the same complexity class. These problems concern establishing which properties (*e.g.*, temporal, ω -regular, etc.) hold in a game, under the assumption that players in the game choose strategies in equilibrium. More specifically, the questions that we ask are as follows. Given a game as described before, in which each player has a qualitative goal given by an LTL formula, and a quantitative goal given by a mean-payoff

¹ We remark that strict Nash equilibria appear naturally in the study of evolutionary game-theory [41]. Indeed, every strict Nash equilibrium in a symmetric game is evolutionary stable.

² The transition systems on which these games are played are sometimes called “concurrent game structures” and sometimes “arenas”. Note that “concurrent” in this context simply means that players move at the same time, *i.e.*, synchronously.

condition, we ask whether a given LTL formula, say ϕ , is satisfied on some/all (strict ϵ) Nash equilibrium/equilibria, if any, of the game. Since ϕ is not a goal of any of the players, it can be seen as a property that the “designer” of the game wants to see satisfied assuming rational behaviour of the players/agents in the game/system.

The remainder of the paper is structured as follows:

- Section 2 introduces our formal framework and defines the games we study throughout the paper.
- Section 3 defines the solution concepts underlying our main results, and presents the constructions, reductions, and algorithms to solve the main problems considered in the paper.
- Section 4 discusses related work.
- Finally, Section 5 presents a number of concluding remarks and directions for potential future work.

2 Game Structures

In this section we introduce our game model. We use multi-player games played on finite directed graphs (transition systems), rather than games in extensive-form or normal-form [35]. Agents move synchronously (which includes the special sequential case), play deterministic (rather than randomised) and finite-state (instead of simply memoryless or infinite-memory) strategies, in pursuit of their individual preferences, which are given as a lexicographic combination of a qualitative temporal-logic property (intuitively, a goal/objective) and a quantitative long-term average of the rewards of its actions.

We fix some notation. If X is a set, then X^ω is the set of all infinite sequences over X . If α is a sequence and $n \in \mathbb{N}$ (the set of non-negative integers) then α_n represents the $(n + 1)$ st element of α . If X and Y are sets, then X^Y is the set of all functions $\alpha : Y \rightarrow X$. We will often use Greek letters $\alpha, \beta, \kappa, \dots$ to name functions. Also, we will use tuple notation: we write $\alpha_y \in X$ instead of $\alpha(y)$. We write AP for a finite set of *atomic propositions* (or *atoms* for short).

We now define the framework of (propositional) Linear Temporal Logic (LTL), which we use extensively in what follows. Our presentation is complete but is not intended as an introduction to this well-known language: see, e.g., [18] for a detailed overview.

Linear-Temporal Logic (LTL) The *formulae* of LTL (over AP) are generated by the following grammar:

$$\varphi ::= p \mid \varphi \wedge \varphi \mid \neg\varphi \mid X\varphi \mid \varphi \cup \varphi$$

where $p \in AP$. We use the standard classical logic abbreviations, e.g., $\text{false} := p \wedge \neg p$, as well as those for LTL, e.g., $F\varphi := \text{true} \cup \varphi$ and $G\varphi := \neg F\neg\varphi$.

Formulae of LTL are interpreted over infinite words $\alpha \in (2^{AP})^\omega$. Define the *satisfaction* relation \models as follows:

- $(\alpha, n) \models p$ iff $p \in \alpha_n$;
- $(\alpha, n) \models \varphi_1 \wedge \varphi_2$ iff $(\alpha, n) \models \varphi_i$ for $i = 1, 2$;
- $(\alpha, n) \models \neg\varphi$ iff it is not the case that $(\alpha, n) \models \varphi$;
- $(\alpha, n) \models X\varphi$ iff $(\alpha, n + 1) \models \varphi$;
- $(\alpha, n) \models \varphi_1 \cup \varphi_2$ iff there exists $j \geq n$ such that $(\alpha, j) \models \varphi_2$ and for all $n \leq i < j$, $(\alpha, i) \models \varphi_1$.

Finally, define $\alpha \models \varphi$ if $(\alpha, 0) \models \varphi$. The *size* of a formula is simply the number of operators within it.

Arenas and Lexicographic Games An arena is a tuple

$$A = \langle \text{Ag}, \text{Act}, \text{St}, \iota, \tau \rangle$$

where Ag , Act , and St are finite non-empty sets of *agents*, *actions*, and *states*, respectively; $\iota \in \text{St}$ is the *initial state*; and $\tau : \text{St} \times \text{Act}^{\text{Ag}} \rightarrow \text{St}$ is a *transition function* mapping each pair consisting of a state and an action for each agent, namely a *decision* $\delta \in \text{Act}^{\text{Ag}}$, to a successor state.

A Lex(LTL, mp) *game* is a tuple

$$G = \langle A, (\kappa_a)_{a \in \text{Ag}}, \text{AP}, \lambda, (\gamma_a)_{a \in \text{Ag}} \rangle$$

where A is an arena; $\kappa_a : \text{St} \rightarrow \mathbb{Z}$ is a *weight function* for agent $a \in \text{Ag}$ associating an integer *weight* to each state; AP is a finite set of *atomic propositions*; $\lambda : \text{St} \rightarrow 2^{\text{AP}}$ is a *labelling function* assigning a subset of atomic propositions to every state of the arena; and γ_a is an LTL formula over AP , called the LTL *goal* associated with agent a .

In the following, we introduce some basic notions related to games.

Executions A path $\pi = s_0 \delta_0 s_1 \delta_1 \dots$ is an infinite sequence over $\text{St} \times \text{Act}^{\text{Ag}}$ such that $\tau(s_i, \delta_i) = s_{i+1}$ for all i . In particular, $\delta_i(a)$ is the action of agent a in step i .

A path π induces:

1. the sequence $\lambda(\pi) = \lambda(s_0)\lambda(s_1)\dots$ of sets of atoms, and
2. for each agent a , the sequence $\kappa_a(\pi) = \kappa_a(s_0)\kappa_a(s_1)\dots$ of weights.

An *execution* is a path with $s_0 = \iota$. Let Exec denote the set of all executions.

Qualitative Goals In this work, qualitative goals are represented by LTL formulas. If γ is an LTL formula and π is an execution, we say that π *satisfies* γ , and write $\pi \models \gamma$, if $\lambda(\pi) \models \gamma$.

For an execution $\pi \in \text{Exec}$ and an LTL goal γ_a , define

$$\text{sat}_a(\pi) = \begin{cases} \top & \text{if } \pi \models \gamma_a \\ \perp & \text{otherwise.} \end{cases}$$

Quantitative Goals For a sequence $\alpha \in \mathbb{R}^\omega$, let $\text{mp}(\alpha)$ be the *mean-payoff* of α , that is,

$$\text{mp}(\alpha) = \liminf_{n \rightarrow \infty} \text{avg}_n(\alpha)$$

where, for $n \in \mathbb{N}$, we define

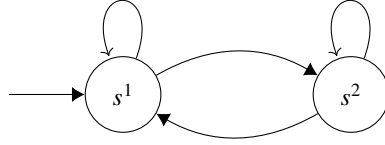
$$\text{avg}_n(\alpha) = \frac{1}{n} \sum_{j=0}^{n-1} \alpha_j.$$

This definition naturally extends to executions, *i.e.*, define $\text{mp}_a(\pi) = \text{mp}(\kappa_a(\pi))$.

Lexicographic Payoffs Let $\Omega = \{\perp, \top\} \times \mathbb{R}$ denote the set of *payoffs*, and define the *payoff function* for agent a to be $\text{pay}_a : \text{Exec} \rightarrow \Omega$ by $\text{pay}_a(\pi) = (\text{sat}_a(\pi), \text{mp}_a(\pi))$. Each agent is trying to maximise its payoff. In other words, agent a 's primary goal is to *satisfy* its LTL formula γ_a , and its secondary goal is to *maximise* its mp-reward $\text{mp}_a(\pi)$. Formally, we define the *lexicographic ordering* on the set Ω of payoffs: $(x, y) <_{lex} (x', y')$ iff, either $(x = \perp$ and $x' = \top)$ or $(x = x'$ and $y < y')$. Note that this ordering is total.

Remark 1 We consider weights as rewards to be maximised. However, one may be concerned with games where the agents have costs they want to minimise. One immediate thought is to take such a cost-game, replace all the weights by their negation and consider the resulting maximisation problem. However, the resulting game will not be strategically identical to the original cost game, since for an arbitrary execution π , we do not have $-\text{mp}(\kappa_a(\pi)) = \text{mp}(-\kappa_a(\pi))$ in general.

One easy way to see this is to consider the arena, A , with two states, s^1, s^2 , (the number of agents and their available actions are not relevant for the most part) with a transition function such that the players can either stay in the same state, or move to the other state. Moreover, we set s^1 to be the start state. Thus, the arena looks like this ³:



Now, for a given player, a , set $\kappa_a(s^1) = 0$ and $\kappa_a(s^2) = 1$. Now define two sequences, a_n, b_n , with,

$$\begin{aligned} a_0 &= 0, \\ b_0 &= 3, \\ a_{n+1} &= 3b_n + 2, \\ b_{n+1} &= 3a_{n+1} + 2, \end{aligned}$$

for all $n \geq 0$. It is easy to see that we have both $a_n < b_n$, as well as $b_n < a_{n+1}$ for all $n \geq 0$. With this, we define an execution π such that $\pi[k] = s^1$ if there exists some n such that $a_n < k < b_n$ and $\pi[k] = s^2$ otherwise. Intuitively, π bounces between s^1 and s^2 , spending three times as long on each state as it did on the previous state.

It is easy to verify that $\text{avg}_{a_n}(\pi) \leq 0.25$ and that $\text{avg}_{b_n}(\pi) \geq 0.75$ for all n . Thus, we have,

$$\begin{aligned} -\text{mp}(\kappa_a(\pi)) &\geq -0.25, \\ \text{mp}(-\kappa_a(\pi)) &\leq -0.75. \end{aligned}$$

Thus, to reason about cost-games, we cannot simply negate the weights and consider the resulting maximisation problem.

Whilst this may seem unsatisfying, there are two ways out here. Firstly, for finite state strategies, these values *do* coincide, as these induce ultimately periodic executions, whose sequence of weights will have a well-defined limit-average. Secondly, whilst the negative

³ Note that the arena is similar to the one of [45, Fig. 3 in Lemma 7] to prove that Multi mean-payoff games require in general infinite memory strategies to be played optimally.

weights may not directly encode the original game, they do generally reflect the strategic nature of it - the resulting game is not *completely* disparate from the game it was based on. Additionally, checking a given strategy profile to see if it is a Nash equilibrium is generally much easier than synthesising one in the first place. As such, we can consider the maximisation game and then try and translate our understanding of it to the original cost game.

Strategies A *history* is a finite, possibly empty, sequence $\delta_0\delta_1 \cdots \delta_{n-1}$ of decisions. The set of all histories is denoted Hst . A *strategy*, σ , for agent $a \in \text{Ag}$ is a function $\text{Hst} \rightarrow \text{Act}$. We emphasize that strategies map finite sequences of decisions (not states) to actions. This differs from the conventional definition in the literature. However, with this alternative definition, the Nash equilibria of the game are invariant under bisimulation [26] – with the conventional definition, they are not. Thus, we use our model so that our algorithm produces equilibria that have the useful property of being invariant under bisimulation.

A *strategy profile* is a function $\vec{\sigma} : \text{Ag} \rightarrow (\text{Hst} \rightarrow \text{Act})$. A strategy profile $\vec{\sigma}$ induces a unique execution $\pi_{\vec{\sigma}}$, *i.e.*, the execution $\pi_{\vec{\sigma}} = s_0\delta_0s_1\delta_1 \cdots$ such that $s_0 = \iota$ and $\delta_i(a) = \vec{\sigma}(a)(\delta_0\delta_1 \cdots \delta_{i-1})$ for $i \geq 0$.

Let $a \in \text{Ag}$ be a player, $\vec{\sigma}$ a strategy profile, and σ'_a be an additional strategy for player a - for convenience, we introduce two associated functions, $\vec{\sigma}_{-a} : \text{Ag} \setminus \{a\} \rightarrow (\text{Hst} \rightarrow \text{Act})$ and $(\vec{\sigma}_{-a}, \sigma'_a) : \text{Ag} \rightarrow (\text{Hst} \rightarrow \text{Act})$. Semantically, we define these as follows: $\vec{\sigma}_{-a}(b) = \vec{\sigma}(b)$ for all $b \in \text{Ag} \setminus \{a\}$, and $(\vec{\sigma}_{-a}, \sigma'_a)(a) = \sigma'_a$ with $(\vec{\sigma}_{-a}, \sigma'_a)(b) = \vec{\sigma}(b)$ for all $b \in \text{Ag}$ with $b \neq a$.

Finite-state strategies A strategy σ is *finite-state* if it is generated by an automaton M with input alphabet $\Sigma = \text{Act}^{\text{Ag}}$ and with output function $\lambda : Q \rightarrow \text{Act}$. That is, on input $h \in \text{Hst}$, the automaton M reaches a state q_h such that $\lambda(q_h) = \sigma(h)$. A strategy profile, $\vec{\sigma}$ is *finite-state* if every strategy $\vec{\sigma}(a)$ is finite-state. Observe that in this case, the unique execution $\pi_{\vec{\sigma}}$ is ultimately periodic.

Remark 2 There are several reasons for considering finite state strategies, rather than strategies with unbounded memory. First, from a modelling perspective (especially within the AI and multi-agent systems communities), this is a desirable representation as it can be used to synthesise models for individual agents in a system. Second, from a more theoretical and computational standpoint, in games with quantitative objectives, finite state strategies can render decidable settings that would be undecidable otherwise [43]. Finally, the use of finite state machines to capture strategies in games played over an infinite number of rounds is standard in the game theory literature [6].

Strict ϵ Nash-equilibria The solution concept we work with is the *strict ϵ Nash-equilibrium*. This is a natural refinement of ϵ Nash-equilibrium [40], and moreover includes strict Nash equilibrium as a special case. For $\epsilon \geq 0$ and $(x, y) \in \Omega$, let $(x, y) + \epsilon$ denote $(x, y + \epsilon) \in \Omega$. A strategy profile $\vec{\sigma}$ is a *strict ϵ Nash-equilibrium* if for every agent $a \in \text{Ag}$, and every strategy $\sigma'_a \neq \vec{\sigma}_a$ for a , we have that $\text{pay}_a(\pi_{(\vec{\sigma}_{-a}, \sigma'_a)}) <_{lex} \text{pay}_a(\pi_{\vec{\sigma}}) + \epsilon$. If $\epsilon = 0$ then we call this a *strict Nash equilibrium*. We remark that an (ordinary) Nash equilibrium uses \leq_{lex} instead of $<_{lex}$. By $\text{FSNE}^\epsilon(G)$ we denote the set of Finite-state Strict ϵ Nash Equilibria in G . We emphasise that, in the definition of a finite-state strict ϵ Nash-equilibrium $\vec{\sigma}$, the deviating strategies $\vec{\sigma}'(a)$ need not be finite-state. Intuitively, this captures worst-case behaviour of the deviators.

Decision Problems The central decision problem of this work, called Rational Synthesis or Rational Verification [22,24,47,32], asks if there exists a FSNE $^\epsilon$ so that the induced play $\pi_{\vec{\sigma}}$ satisfies a given LTL condition Φ . Note that in case $\Phi = \top$ this amounts to the deciding the existence of a FSNE $^\epsilon$.

Formally, for a rational $\epsilon \geq 0$ we consider the following decision problems for the class of Lex(LTL, mp)-games:

- FSNE $^\epsilon$ -emptiness.
Given: game G
Question: Is it the case that FSNE $^\epsilon(G) \neq \emptyset$?
- FSNE $^\epsilon$ -existence.
Given: Game G and LTL formula Φ .
Question: Does there exist a $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$ such that $\pi_{\vec{\sigma}} \models \Phi$?

To help the reader better understand our model and its applicability, we present an example.

Example 1 In an automated warehouse, n robots move around to load items and bring them to the exit. Their objective is to load and unload as efficiently as possible—without crashing into each other.

Assume that the warehouse is represented by a directed edge-labeled graph $\mathcal{G} = (V, E)$ where $E : V \times D \rightarrow V$ for some finite set of *directions* D . For instance, if the warehouse is a grid then we may take $D = \{\text{north}, \text{south}, \text{east}, \text{west}\}$ and an agent in position v executing action $d \in D$ will move to position $E(v, d)$. In addition, we are given particular vertices: for each robot $a \in \text{Ag}$, a vertex r_a representing its initial position, $\text{ex} \in V$ representing the exit and $l_1, \dots, l_n \in V$ representing the loading points for the agents. We assume loading points are different from each other and from the exit.

We can model the setting by means of the arena

$$A = \langle \text{Ag}, \text{Act}, \text{St}, \iota, \tau \rangle$$

where

- the agents $\text{Ag} = \{1, \dots, n\}$ are the robots moving around the warehouse;
- $\text{Act} = D$ are the actions that each robot can take;
- $\text{St} = V^n$ is the set of states of the system, where the a -th component of the tuple denotes the position of robot a ;
- $\iota = (r_1, \dots, r_n)$ is the initial state, denoting the initial position of the robots;
- and $\tau : \text{St} \times \text{Act}^{\text{Ag}} \rightarrow \text{St}$ maps (s, δ) to the state whose a th component is $E(s_a, \delta(a))$.

To capture robot objectives, we define the Lex(LTL, mp) game

$$G = \langle A, (\kappa_a)_{a \in \text{Ag}}, \text{AP}, \lambda, (\gamma_a)_{a \in \text{Ag}} \rangle$$

where A is the arena described above and

- $\kappa_a(v) = 1$ if $v = l_a$, and 0 otherwise (the weight function rewards the agent anytime it hits the loading point);
- $\text{AP} = \{\text{exit}_1, \dots, \text{exit}_n, \text{load}_1, \dots, \text{load}_n, \text{crash}_1, \dots, \text{crash}_n\}$, denoting the agent a is at the exit (exit_a), or the loading vertex (load_a), or crashing with another agent (crash_a);
- For every state $s \in V$ and agent $a \in \text{Ag}$, we have that

- $\text{exit}_a \in \lambda(s)$ iff $s_a = \text{ex}$,
- $\text{load}_a \in \lambda(s)$ iff $s_a = l_a$,
- $\text{crash}_a \in \lambda(s)$ iff $s_a = s_b$ for some $b \neq a$ (i.e., another agent occupies the same position as agent a);
- for every agent a , $\gamma_a = \mathbb{G}(\neg \text{crash}_a) \wedge \mathbb{G}(\text{load}_a \rightarrow \mathbb{X}(\neg \text{load}_a \cup \text{exit}_a))$, i.e., the agent is trying to ensure that it never crashes with another agent and that it visits the exit point between every two occurrences of a visit to the loading point (this captures that the agent successfully loads and unloads an item).

Intuitively, an agent's primary objective (the LTL formula) is to never crash and never load two items in a row before reaching the exit. Furthermore, an agent's secondary objective (the mean-payoff value) is to ship items as fast as possible, in order to maximise, in the limit-average, the number of times it reaches its loading point.

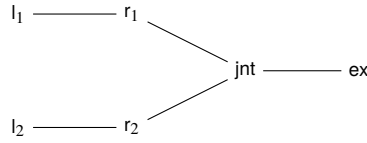


Fig. 1 Representation of an automated warehouse with two operating robots.

What might an equilibrium for this game look like? Observe that agents have essentially two possible behaviours to satisfy the primary goal: *idle* and *cycle*. The *idle* allows the robots a finite number of trips from the loading point to the exit point before keeping them forever away from the loading point, whereas the *cycle* makes them move back and forth between the loading point and the exit forever. In both cases, they additionally have to avoid crashing with each other. The secondary goal makes every agent to prefer the *cycle* behaviour. Indeed, the only way to get a strictly positive mean-payoff reward is to reach the loading point infinitely many times (with bounded delay among two consecutive times).

Consider a warehouse with two operating robots as represented in Figure 1, and assume that the initial positions of robot_1 and robot_2 are r_1 and r_2 , respectively. Moreover, consider the infinite path $\pi = u \cdot (u')^\omega$ with

$$u = (r_1, r_2), (l_1, l_2), (r_1, r_2), (\text{jnt}, r_2), (\text{ex}, r_2)$$

and

$$u' = (\text{jnt}, r_2)(r_1, \text{jnt}), (l_1, \text{ex}), (r_1, \text{jnt}), (\text{jnt}, r_2), (\text{ex}, l_2)$$

where each pair represents the position of the robot at any step.

The path satisfies both γ_1 and γ_2 . Moreover, every robot cycles from the loading point to exit and back in six steps. Therefore, their mean-payoff value on π is given by $\frac{1}{6}$. Observe that it would not be possible to improve such payoff without violating the primary objective. Thus, π achieves optimal value for both robots and it is therefore a ϵ -Nash Equilibrium, for every $\epsilon \geq 0$ (in particular, it is a Nash Equilibrium for $\epsilon = 0$). It is also strict, as any other path of robot_a satisfying γ_a decreases the mean-payoff value.

Observe that for a large enough value of ϵ , every path that satisfies both γ_1 and γ_2 is a strict ϵ Nash Equilibrium. In particular, for $\epsilon = \frac{1}{6}$, the *idle* strategies for the robots producing the outcome $(r_1, r_2)^\omega$ is a strict ϵ Nash Equilibrium, as the primary goals are all achieved, and the secondary goal cannot be improved by strictly more than $\frac{1}{6}$. This is obviously not desirable from the designers point of view, as we might require that none of

the agents remains idle. Thus, we might require an equilibrium to satisfy the LTL formula $\Phi = \bigwedge_a \text{GF load}_a$ which states that each agent is loading a fresh item infinitely often. The existence of such a solution can be checked by solving the FSNE^ϵ -existence problem.

3 FSNE^ϵ -Existence and FSNE^ϵ -Emptiness are 2ExpTime -complete

In this section we establish our main technical result, *i.e.*, that FSNE^ϵ -emptiness is in 2ExpTime . We then show that FSNE^ϵ -existence is 2ExpTime -complete - we show membership by reducing to the FSNE^ϵ -emptiness problem and show hardness using a reduction from LTL games. We remark that some of the assumptions in these results are quite general (*i.e.*, using LTL as a specification language for qualitative properties, mean-payoff for quantitative properties, equilibria as a solution concept, restricting to finite-state strategies) and some are also used to make the proofs practicable (*i.e.*, using strict equilibria rather than ordinary equilibria, see the discussion in the conclusion).

Theorem 1 *The following problem is in 2ExpTime : given a $\text{Lex}(\text{LTL}, \text{mp})$ game G and a rational $\epsilon \geq 0$, decide whether there exists a strategy profile $\vec{\sigma}$ such that $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$.*

We split the proof into four steps, which we now outline. After giving the technicalities of these steps, we show how to put them together in Section 3.5 to establish the theorem.

Replace LTL- by parity-objectives. In Section 3.1 we show that every $\text{Lex}(\text{LTL}, \text{mp})$ game can be converted into a $\text{Lex}(\text{parity}, \text{mp})$ game having the same set of finite-state strict ϵ -Nash Equilibria. Intuitively, we replace each agent's LTL objective by a parity objective. Let P be a finite set of integers - then a sequence of priorities, $p_0 p_1 \dots \in P^\omega$, satisfies the parity condition if the largest priority occurring infinitely often in the sequence is even. Thus, the rest of the proof applies to a $\text{Lex}(\text{parity}, \text{mp})$ -game $G = (A, (\kappa_a)_{a \in \text{Ag}}, (\rho_a)_{a \in \text{Ag}})$ where $\rho_a : \text{St} \rightarrow \mathbb{Z}$ is a priority function for each agent, and agent a 's primary objective is to ensure the sequence $\rho_a(\cdot)$ satisfies the parity condition, and a 's secondary objective is to maximise the mean-payoff of $\kappa_a(\cdot)$. Later, we show that such a reduction results in a $\text{Lex}(\text{parity}, \text{mp})$ game that is at most doubly exponentially larger than the size of the goals γ_a 's.

Two-agent zero-sum games. In Section 3.2 we study two-agent zero-sum games with a $\text{Lex}(\text{parity}, \text{mp})$ objective (played on the same arena as G). We prove that every such game, H , has a minimax value $\text{val}(H)$, and this value is computable in time polynomial in the number of states and edges, and exponential in the number of priorities and weights of the game. Moreover, we show that for every $\epsilon > 0$ there exists a finite-state strategy for the minimizing agent that ensures the maximizing agent's payoff can achieve a payoff of at most $\text{val}(H) + \epsilon$. The proofs in this section make use of Mean-Payoff Parity Games. In particular, it builds from the computation of optimal value of these games as in [15] to derive the optimal value in the $\text{Lex}(\text{parity}, \text{mp})$ setting.

Reducing Equilibrium Finding to Path Finding In Section 3.3 we reduce the problem of FSNE^ϵ -emptiness to the one of finding payoff thresholds $\bar{z} \in \Omega^{\text{Ag}}$ and an ultimately periodic path π in a certain graph (that we call $G[\bar{z}]$) such that $z_a <_{\text{lex}} \text{pay}_a(\pi) + \epsilon$. More precisely, each z_a is a so-called "punishing value", *i.e.*, the value of a two-agent zero-sum game with a $\text{Lex}(\text{LTL}, \text{mp})$ objective played on the same arena as G , starting at some state $s \in \text{St}$, but with a trying to maximise its payoff and the rest of the opponents (viewed as a single player)

trying to minimise a 's payoff. As such, we also show that for each agent a , the value z_a can be taken from the set of values (parameterised over the possible start states of the arena) of the two-player game H considered in the previous step. Thus, the state space of the vector \bar{z} is bounded by the number of states of the game.

Path Finding in Multi-Weighted Graphs with LEX(parity, mp) Payoffs In Section 3.4 we show how to find ultimately periodic paths π such that $z_a <_{lex} \text{pay}_a(\pi) + \epsilon$ in graphs of the form $G[\bar{z}]$. We do this by adapting the linear programming approach for computing zero-cycles in mean-payoff graphs [31].

3.1 Replacing LTL objectives by parity objectives

In this section, we show that every Lex(LTL, mp) game can be converted into a Lex(parity, mp) game having the same set of finite-state strict ϵ -Nash Equilibria. We begin with a definition of the parity condition and Lex(parity, mp)-games.

Parity games A sequence $\alpha \in X^\omega$, where X is a finite non-empty set of integer *priorities* satisfies the *parity condition* if the largest priority occurring infinitely often is even.

A Lex(parity, mp) game G is a tuple

$$\langle A, (\kappa_a)_{a \in \text{Ag}}, (\rho_a)_{a \in \text{Ag}} \rangle$$

where A is an arena, $\kappa_a : \text{St} \rightarrow \mathbb{Z}$ is a weight function for agent a , and $\rho_a : \text{St} \rightarrow \mathbb{Z}$ is a *priority function* for agent a . For an execution $\pi = s_0 \delta_0 s_1 \delta_1 \dots \in \text{Exec}$ let $\rho(\pi) = \rho(s_0) \rho(s_1) \dots \in \mathbb{Z}^\omega$. For an agent $a \in \text{Ag}$ define

$$\text{parity}_a(\pi) = \begin{cases} \top & \text{if } \rho(\pi) \text{ satisfies the parity condition.} \\ \perp & \text{otherwise.} \end{cases}$$

The payoff function $\text{pay}_a : \text{Exec} \rightarrow \Omega$ for agent a is defined by

$$\text{pay}_a(\pi) = (\text{parity}_a(\pi), \text{mp}_a(\pi)).$$

As before, each agent is trying to maximise its payoff under the lexicographic ordering. This completes the definition of parity games.

Deterministic Automata In order to systematically perform the required conversion, we introduce *deterministic parity automata*. An *automaton* is a tuple $\langle \Sigma, Q, q_0, \Delta \rangle$ where Σ is a finite non-empty set called the *input alphabet*, Q is a finite non-empty set of *states*, $q_0 \in Q$ is an *initial state*, and $\Delta : Q \times \Sigma \rightarrow Q$ is a *transition function*.

A *deterministic parity automaton on words (DPW)* is an automaton with a *priority function* $\rho : Q \rightarrow \mathbb{Z}$. The *number of priorities* is the cardinality of the set $\rho(Q)$. An *input word* is an infinite sequence over Σ . A *run* is an infinite sequence over Q . Every input word $w_0 w_1 \dots$ determines a run $s_0 s_1 \dots$, i.e., $s_0 = q_0$ and $\Delta(s_i, w_i) = s_{i+1}$ for every $i \geq 0$. A run is *accepting* if the largest priority occurring infinitely often in $\rho(s_0) \rho(s_1) \dots$ is even. An input word is *accepted* if its run is accepting. The *language* of a DPW \mathcal{A} , denoted $\mathcal{L}(\mathcal{A})$, is the set of input words it accepts.

With this machinery in place, one can effectively compile LTL formulas into DPW:

Theorem 2 [44, 36] *One can effectively transform a given LTL formula φ over AP into a DPW $\mathcal{A} = \langle 2^{\text{AP}}, Q, q^0, \Delta, \rho \rangle$ over the alphabet 2^{AP} such that $\mathcal{L}(\mathcal{A}) = \{\alpha \in (2^{\text{AP}})^\omega : \alpha \models \varphi\}$. Moreover, the number of states of the DPW is at most doubly exponentially larger than the size of φ , and the number of priorities is at most singly exponentially larger than the size of φ .*

Thus, we start by translating every LTL goal γ_a into a deterministic parity word (DPW) automaton $\mathcal{A}_a = \langle 2^{\text{AP}}, Q_a, q_a^0, \Delta_a, \rho_a \rangle$. Then, for a given Lex(LTL, mp) game $G = \langle A, (\kappa_a)_{a \in \text{Ag}}, \text{AP}, \lambda, (\gamma_a)_{a \in \text{Ag}} \rangle$ over the arena $A = \langle \text{Ag}, \text{Act}, \text{St}, \iota, \tau \rangle$, define the arena $A' = \langle \text{Ag}, \text{Act}, \text{St}', \iota', \tau' \rangle$ where

- $\text{St}' = \text{St} \times \prod_{a \in \text{Ag}} Q_a$;
- $\iota' = (\iota, q_{a_1}^0, \dots, q_{a_n}^0)$;
- For each state $(s, q_{a_1}, \dots, q_{a_n})$ in St' and decision $\delta \in \text{Act}^{\text{Ag}}$, let

$$\tau'((s, q_{a_1}, \dots, q_{a_n}), \delta) = (\tau(s, \delta), \Delta_{a_1}(q_{a_1}, \lambda(s)), \dots, \Delta_{a_n}(q_{a_n}, \lambda(s))).$$

Define the Lex(parity, mp) game $G' = \langle A', (\kappa'_a)_{a \in \text{Ag}}, (\rho'_a)_{a \in \text{Ag}} \rangle$ where

$$\begin{aligned} \kappa'_a(s, q_{a_1}, \dots, q_{a_n}) &= \kappa_a(s), \\ \rho'_a(s, q_{a_1}, \dots, q_{a_n}) &= \rho_a(q_a). \end{aligned}$$

Intuitively, the Lex(parity, mp) game G' is the product of the Lex(LTL, mp) game G and the collection of parity automata that recognize the models of each player's LTL goal. Informally, the game executes the original game in parallel with the automata at every step of the game, the arena-component of the product state follows the transition function of the original game G , while the automata-components follow the transition functions of the simulated automata and are updated according to the labelling of the current state of G . As a result, the execution in G' can be recovered from the original execution π in the game G and the unique runs of the (deterministic) automata generated when reading the word $\lambda(\pi)$.

Observe that in the translation from G to its associated G' the set of actions for each player is unchanged. This, in turn, means that the set of strategies in both G and G' is the same; indeed, recall from the definitions that strategies are functions from finite sequences of decisions (not states) to actions. Using this correspondence between strategies in G and strategies in G' , we can prove the following Proposition, which states an invariance result between G and G' with respect to the satisfaction of players' goals.

Proposition 1 (Payoff invariance) *Let G be a Lex(LTL, mp) game and G' its associated Lex(parity, mp) game. Then, for every strategy profile $\vec{\sigma}$ and player a , it is the case that $\text{pay}_a(\pi_{\vec{\sigma}}^G) = \text{pay}_a(\pi_{\vec{\sigma}}^{G'})$, where by $\text{pay}_a(\pi_{\vec{\sigma}}^G)$ we denote the payoff of agent a on the execution $\pi_{\vec{\sigma}}^G$ in G and $\text{pay}_a(\pi_{\vec{\sigma}}^{G'})$ the payoff of agent a on the execution $\pi_{\vec{\sigma}}^{G'}$ in G' .*

Proof We will use the following notation: $\pi = \pi_{\vec{\sigma}}^G$ and $\pi' = \pi_{\vec{\sigma}}^{G'}$. It is sufficient to show, for every agent a , that $\kappa_a(\pi) = \kappa'_a(\pi')$ and that $\pi \models \gamma_a$ iff $\rho'_a(\pi')$ satisfies the parity condition. We use the following fact, which follows by the construction of the game G' : there is an exact two-way correspondence between runs in G , and runs in G' . If $\pi = s_0 \delta_0 s_1 \delta_1 \dots$, then $\pi' = (s_0, \vec{q}^0) \delta_0 (s_1, \vec{q}^1) \delta_1 \dots$, where $\vec{q}^i = (q_{a_1}^i, \dots, q_{a_n}^i) \in \prod_a Q_a$ has the property that $r_a = q_a^0 q_a^1 \dots$ is the unique run of the DPW \mathcal{A}_a for LTL goal γ_a on input $\lambda(s_0) \lambda(s_1) \dots$ (for every agent a). Conversely, if π' is a run in G' , then it is easy to see that π is the unique run in G where the induced sequence of sets of atomic propositions that arises corresponds to the running of the deterministic parity automata.

Now, for the quantitative part of the payoff, note that by definition of κ' , $\kappa_a(\pi) = \kappa'_a(\pi')$, as required. Similarly, for the qualitative part of the payoff, note that $\pi \models \gamma_a$ iff the run r_a is accepting (since this is how \mathcal{A}_a was chosen), i.e., $\rho_a(r_a)$ satisfies the parity condition. But $\rho_a(r_a) = \rho'_a(\pi')$ by definition of ρ' . \square

Proposition 1 allows us to prove that the set of strict ϵ -Nash Equilibria in G exactly corresponds to the set of strict ϵ -Nash Equilibria in G' . The following result holds.

Proposition 2 *Let G be a Lex(LTL, mp) game and G' its associated Lex(parity, mp) game. Then $\text{FSNE}^\epsilon(G) = \text{FSNE}^\epsilon(G')$.*

Proof We show one direction (the other is symmetric). Let $\vec{\sigma}$ be a strategy profile that is not in $\text{FSNE}^\epsilon(G')$, i.e., there is an agent a and a strategy profile $\vec{\sigma}'$ such that $\vec{\sigma}'_a \neq \sigma_a$, $\vec{\sigma}'_b = \sigma_b$ for $b \neq a$, and $\text{pay}_a(\pi_{\vec{\sigma}'}) + \epsilon \leq_{lex} \text{pay}_a(\pi_{\vec{\sigma}})$. Thus, by applying Proposition 1 on both sides of the inequality, we obtain that $\text{pay}_a(\pi_{\vec{\sigma}'}) + \epsilon \leq_{lex} \text{pay}_a(\pi_{\vec{\sigma}})$, which implies that $\vec{\sigma} \notin \text{FSNE}^\epsilon(G)$. \square

Finally, note that the number of states of G' is doubly exponential in the size of the LTL goals of G , and that the number of priorities in G' is singly exponential in the size of the LTL goals of G . This will be important when establishing the complexity of the existence problem later in the paper.

3.2 Two-Agent Zero-Sum Lex(parity,mp)-Games

We begin with a study of two-agent zero-sum Lex(parity, mp) games. To simplify notation, we define these as $H = \langle A, \kappa, \rho \rangle$ where A is an arena with $\text{Ag} = \{1, 2\}$, and $\kappa, \rho : \text{St} \rightarrow \mathbb{Z}$. Define $\text{pay}(\pi) = (\text{parity}(\pi), \text{mp}(\pi))$. Player 1 is called the “maximizer” and player 2 is called the “minimizer”. Thus, intuitively, agent 1 is trying to maximise the value of $\text{pay}(\cdot)$ while agent 2 is trying to minimize it.

Now, a basic question is; ‘what is the highest payoff that player 1 can achieve?’. Formally, we can ask what the following quantity is:

$$\overline{\text{val}}(H) = \sup_{\sigma} \inf_{\zeta} \text{pay}(\pi_{\langle \sigma, \zeta \rangle}).$$

Here, σ ranges over strategies of player 1 (the maximizer), ζ ranges over strategies of player 2 (the minimizer), and $\pi_{\langle \sigma, \zeta \rangle}$ is the unique execution determined by the strategy profile $\langle \sigma, \zeta \rangle$.

Conversely, we can consider the smallest payoff that player 2 can inflict on player 1. That is, the quantity

$$\underline{\text{val}}(H) = \inf_{\zeta} \sup_{\sigma} \text{pay}(\pi_{\langle \sigma, \zeta \rangle}),$$

where σ, ζ and $\pi_{\langle \sigma, \zeta \rangle}$ are as above.

In arbitrary, zero-sum two player games, we’d call these two values the maximin and minimax values, and as in arbitrary games, it is easy to verify that $\underline{\text{val}}(H) \leq \overline{\text{val}}(H)$. Thus, a natural question for such games is under what conditions do these two values coincide – when do we have $\underline{\text{val}}(H) = \overline{\text{val}}(H)$? If they do coincide, then we call this the *value* of the game and denote it by $\text{val}(H)$. We show that for zero-sum Lex(parity, mp)-games, these two values do indeed coincide, and so the value is well-defined. Moreover, we show it can be

computed in TFNP , the class of total function problems which are solvable in nondeterministic polynomial time [34].

To define TFNP formally, suppose we have an alphabet Σ and a left-total binary relation $R : \Sigma^* \times \Sigma^*$ such that for all $x, y \in \Sigma^*$, $R(x, y)$ implies that $|y| \leq p(|x|)$ for some polynomial p . Moreover, suppose that given $x, y \in \Sigma^*$, we can determine in polynomial time whether $R(x, y)$ holds. Then a natural problem is to ask ‘given an $x \in \Sigma^*$, find a $y \in \Sigma^*$ such that $R(x, y)$ ’. The class of all such problems is exactly TFNP .

With this terminology in place, we are in a position to prove the following proposition.

Proposition 3 *Every two-agent zero-sum Lex(parity, mp) game H has a value, denoted $\text{val}(H) \in \Omega$. Moreover, this value can be computed in TFNP .*

Proof W.l.o.g., we can consider H to be turn-based. Indeed, we can ensure this by replacing every transition $s \xrightarrow{(c_1, c_2)} s'$ by two transitions $s \xrightarrow{c_1} s_{c_1} \xrightarrow{c_2} s'$, in a way that all the original states belong to Player 1, while every extra state s_c belongs to Player 2, and has the same weight and priority as s . Note that such a construction depends on the ordering of players, *i.e.*, in order to compute the value for Player 2, we need to employ a construction of a game H'' that replaces $s \xrightarrow{(c_1, c_2)} s'$ by $s \xrightarrow{c_2} s_{c_2} \xrightarrow{c_1} s'$. It is easy to verify that if we have a run π in a non-turn-based game, H , and transform it to a turn-based game, H' , then the corresponding run π' will induce the same payoff as in π .

We compute $\text{val}(H)$ and $\text{val}(H)$ by reducing to solving two-agent turn-based zero-sum games K with mean-payoff parity objectives [15], which are known to be in NP [7]. We show that these two values are equal, and thus, $\text{val}(H)$ is well defined. Moreover, we require two parallel invocations of an NP algorithm, and so the whole process lies within TFNP .

The mean-payoff parity games K are played on the same weighted arenas $\langle A, \kappa, \rho \rangle$ as two-agent zero sum Lex(parity, mp) games. However, the payoff set for K is $\mathbb{R} = \mathbb{R} \cup \{\pm\infty\}$ with its usual ordering $<$, and payoff function $\text{pay}^+ : \text{Exec} \rightarrow \mathbb{R}$ is defined as follows: $\text{pay}^+(\pi)$ equals $-\infty$ if $\text{parity}(\pi) = \perp$, and $\text{mp}(\pi)$ otherwise. Informally, the first player is trying to satisfy the parity condition, and once that holds, maximise its mean-payoff. The value $\text{val}(K)$ is defined to be the maximum payoff that the first player can achieve. It follows from Theorems 2 and 3 of [15] that values exist for these games and can be computed. Moreover, such computation requires nondeterministic polynomial time and lies in NP [7].

To help us calculate the two values of interest, we also introduce an auxillary game, which we call the *dual* of K , which we denote K^* . In K^* , the first player tries to satisfy the parity condition and if this *doesn't* hold, then tries to maximise their mean-payoff. Formally, these are games $K^* = \langle A, \kappa, \rho \rangle$ with payoff function defined as follows: $\text{pay}^*(\pi)$ equals ∞ if $\text{parity}(\pi) = \top$, and $\text{mp}(\pi)$ otherwise.

We argue that these auxillary games also have a well-defined value, $\text{val}(K^*)$. From player 2's perspective, their payoff is $-\infty$ if the parity condition holds and $-\text{mp}(\pi)$ otherwise. Suppose we negated all the weights, and added 1 to each priority. Then in this new game, player 2 is trying to first satisfy the parity condition, then trying to maximise $\limsup_{n \rightarrow \infty} (\text{avg}_n(\kappa(\pi)))$. Thus, we can *almost* view K^* as a mean-payoff parity game, but not quite. Indeed, we certainly can't directly reduce it to a mean-payoff parity game as it stands. However, consider the line of argument in [15] that leads to Theorems 2 and 3 – for every line in this proof, if we use the \limsup rather than the \liminf for player 1, the results still hold. Thus, we can conclude that these auxiliary games also have a value.

With this machinery in place, we compute the value of two-agent zero-sum Lex(parity, mp) games as follows. Let K be the mean-payoff parity game on the same weighted arena as H . Then we are in exactly one of two scenarios: either $\text{val}(K) \neq -\infty$, or $\text{val}(K) = -\infty$.

First suppose that $val(K) \neq -\infty$. Thus, player 1, using some strategy σ , can ensure the parity condition is satisfied, and given this, the greatest payoff they can attain is $val(K)$. This in turn implies that $val(H) = (\top, val(K))$. Additionally, $val(K) \neq -\infty$ also implies that player 2 cannot force the parity condition to not hold, and the lowest payoff they can inflict on player 1 is $val(K)$. This implies that $val(H) = (\top, val(K))$. Putting this together, we get $\overline{val(H)} = val(H)$. Thus, $val(H)$ exists and is equal to $(\top, val(K))$.

Now suppose that $val(K) = -\infty$. If this is the case, then we necessarily have $val(K^*) \neq -\infty$. Thus, player 2 can force that the largest priority occurring infinitely often is odd and can also ensure the smallest payoff player 1 achieves is $val(K^*)$. Thus, we have $val(H) = (\perp, val(K^*))$. Similarly, player 1 cannot force the parity condition to be true, and given this, the greatest payoff they can achieve if $val(K^*)$. Thus, we have $val(H) = (\perp, val(K^*))$. Again, this implies that $val(H)$ exists and is equal to $(\perp, val(K^*))$.

Regarding the complexity, observe that the construction of the games K and K^* is linear in the size of H , and that we employ an TFNP procedure to solve them. Once we have calculated the values of both, this will tell us the value of H . Since the two procedures are independent of one another, we can do both of them sequentially and then compare their values afterwards. Thus, this guarantees that the overall complexity computing $val(H)$ is TFNP. \square

It is worth noting that in mean-payoff parity games, computing the value of the game can be done in time $O(n^d \cdot (m + \text{Parity} + \text{MP}))$ [15]. Here, n is the number of vertices in the game graph, m the number of edges and d the number of priorities. Additionally, Parity and MP are the complexities of solving parity and mean-payoff games respectively. We will use this fact later when we calculate the overall complexity of the FSNE $^\epsilon$ -existence problem.

It is not hard to see that in two-player zero-sum Lex(parity, mp)-games H , a player may need infinite memory to achieve the optimal value $val(H)$ (Cfr. [15, Figure 1]). However, as proven in [7], for every mean-payoff parity game K and every $\epsilon > 0$, there exists a finite-state strategy for the minimizer, ζ (that depends on ϵ) such that for every strategy σ of the maximizer, it holds that $\text{pay}(\pi_{\sigma, \zeta}) \leq val(K) + \epsilon$. Thus, using the same argument as in Proposition 3, we get:

Proposition 4 *For every two-agent zero-sum Lex(parity, mp) game H and every $\epsilon > 0$ there exists a finite-state strategy ζ for the minimizer, such that for every strategy σ of the maximizer (not necessarily finite-state), it holds that $\text{pay}(\pi_{\sigma, \zeta}) \leq_{lex} val(H) + \epsilon$.*

3.3 Reducing Equilibrium Finding to Path Finding

In this section we show that a path π is generated by some $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$ iff π exists in a certain subgraph of the weighted arena of G . To do this, we adapt the proof in [43, Section 6] that shows how to decide the existence of a (not necessarily finite-state) Nash equilibrium for mean-payoff games.

We first need the notion of punishing values and strategies. For $a \in \text{Ag}$ and $s \in \text{St}$ define the *punishing value* $p_a(s)$ to be the \leq_{lex} -largest (x, y) that player a can achieve from state s by playing against the coalition $\text{Ag} \setminus \{a\}$, i.e., by turning the game into a two-player zero-sum game in which the maximizer simulates the moves of player a , the minimizer simulates the moves of the coalition $\text{Ag} \setminus \{a\}$, and the payoff is that of player a . These values can be computed for every player in each state by constructing the appropriate two-agent Lex(parity, mp)-game, H , and invoking Proposition 3. Formally, if $G = \langle A, (\kappa_a), (w_a) \rangle$, and we want

to calculate the punishing value for player i in state s' , then $H = \langle A', \kappa_i, w_i \rangle$, where $\text{Ag}' = \{i, \text{Ag} \setminus \{i\}\}$, $\text{Act}' = \text{Act}$, $\text{St}' = \text{St}$, $t' = s'$ and $\tau'(s, (\delta_1, (\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_{\text{Ag}}))) = \tau(s, (\delta_1, \dots, \delta_{\text{Ag}}))$.

Moreover, for every state $s \in \text{St}$, every player $a \in \text{Ag}$ and every $\epsilon > 0$, fix $\zeta_{s,a}^\epsilon$ to be the strategy of the minimizer described by Proposition 4. We view $\zeta_{s,a}^\epsilon$ as a profile, i.e., $\zeta_{s,a}^\epsilon : \text{Ag} \setminus \{a\} \rightarrow (\text{Hst} \rightarrow \text{Act})$, and call $\zeta_{s,a}^\epsilon(b)$ an ϵ -punishing strategy for agent b . Note that these ϵ -punishing strategies are finite-state.

Definition 1 (Secure values)⁴ For an agent $a \in \text{Ag}$ and $z \in \Omega$, a pair $(s, \delta) \in \text{St} \times \text{Act}^{\text{Ag}}$ is z -secure for a if $p_a(\tau(s, \delta')) \leq_{\text{lex}} z$ for every $\delta' \in \text{Act}^{\text{Ag}}$ that agrees with δ except possibly at a .

With this definition in place, we can now state the following result.

Proposition 5 For every Lex(parity, mp) game G , constant $\epsilon \geq 0$, and ultimately periodic path $\pi = s_0 \delta_0 s_1 \delta_1 \dots$ in G , the following are equivalent:

1. There exists $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$ such that $\pi = \pi_{\vec{\sigma}}$.
2. There exists $\bar{z} \in \Omega^{|\text{Ag}|}$, where $z_a \in \{p_a(s) : s \in \text{St}\}$, $a \in \text{Ag}$, such that for every agent a ,
 - (a) $z_a <_{\text{lex}} \text{pay}_a(\pi) + \epsilon$ and
 - (b) for all $i \in \mathbb{N}$, the pair (s_i, δ_i) is z_a -secure for a .

Proof Fix a game G , constant $\epsilon \geq 0$ and ultimately periodic path $\pi = s_0 \delta_0 s_1 \delta_1 \dots$

For (1) implies (2), suppose there exists $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$ with $\pi = \pi_{\vec{\sigma}}$. Define $\bar{z} \in \Omega^{|\text{Ag}|}$ by $z_a = \max\{p_a(\tau(s_n, \delta'_n)) : n \in \mathbb{N}, \wedge_{b \neq a} \delta'_n(b) = \delta_n(b)\}$, i.e., z_a is the largest value player a can get by deviating from π . For every $n \in \mathbb{N}$, (s_n, δ_n) is z_a -secure for a (by definition of z_a and z_a -secure). Moreover, $z_a <_{\text{lex}} \text{pay}_a(\pi) + \epsilon$: indeed, let n be such that $z_a = p_a(\tau(s_n, \delta'_n))$, and suppose that $\text{pay}_a(\pi) + \epsilon \leq_{\text{lex}} z_a$; then player a would deviate at step n by playing $\delta'_n(a)$ and following a strategy that achieves at least z_a from this point. Note that such a (possibly infinite-state) strategy exists by Proposition 3. But, due to prefix-independence of the payoff function, this is also the payoff of the whole play, contradicting the choice of π as the execution of a strict ϵ Nash-equilibrium.

For (2) implies (1), let $\bar{z} \in \Omega^{|\text{Ag}|}$ be given with the stated properties. We build a strict ϵ Nash-equilibrium $\vec{\sigma}$ such that $\pi_{\vec{\sigma}} = \pi$. For $b \in \text{Ag}$, we define $\vec{\sigma}(b)$ as follows. For every history $h = \delta_0 \dots \delta_n$ (i.e., a decision prefix of π), define $\sigma_b(h) = \delta_n(b)$. Thus, $\vec{\sigma}$ follows π as long as no-one has deviated from π .

For every other history, $h' = \delta'_0 \dots \delta'_n$, let k be the first such integer with $\delta'_k \neq \delta_k$. There are two cases - either δ_k differs in one position, or multiple positions. If δ_k differs in one position, say by player a , then let $s'_{k+1} = \tau(s_k, \delta'_k)$ and then for all other players b , set $\sigma_b(h) = \zeta_{s'_{k+1}, a}^\epsilon(b)(h)$. If δ_k differs in multiple positions, then set $\sigma_b(h)$ arbitrarily.

By construction, we have that $\pi_{\vec{\sigma}} = \pi$. We now aim to show that $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$. Suppose it is not. Then some player a has a strategy, σ_i , such that $\text{pay}_a(\pi_{\vec{\sigma}}) + \epsilon \leq_{\text{lex}} \text{pay}_a(\pi_{(\sigma_{-a}, \sigma_i)})$. By assumption, this implies that $z_a <_{\text{lex}} \text{pay}_a(\pi_{(\sigma_{-a}, \sigma_i)})$. Moreover, we have that $\pi_{(\sigma_{-a}, \sigma_i)} \neq \pi_{\vec{\sigma}}$. Thus, let (s_j, δ'_j) be the first pair from the execution of $\pi_{(\sigma_{-a}, \sigma_i)}$ that differs from the execution of $\pi_{\vec{\sigma}}$ (note that the j^{th} state of both executions is the same). Additionally, let $\tau(s_j, \delta'_j) = s'_{j+1}$.

Now, by assumption, the pair (s_j, δ_j) is z_a -secure for a . This implies that $p_a(s'_{j+1}) \leq_{\text{lex}} z_a$, in turn implying we have $p_a(s'_{j+1}) <_{\text{lex}} \text{pay}_a(\pi_{(\sigma_{-a}, \sigma_i)})$. However, by construction,

⁴ This definition extends the one provided in [43] which considers mean-payoff games.

for all players $b \neq a$, we have $\sigma_b(h) = \sigma_{s'_{j+1}, a}^\epsilon(b)(h)$ for all histories with the prefix $\delta_0 \dots \delta_{j-1} \delta'_j$. By Proposition 4, this implies that $\text{pay}_a(\pi(\sigma_{-a}, \sigma_i)) \prec_{lex} p_a(s'_{j+1})$. But then we can conclude that $p_a(s'_{j+1}) \prec_{lex} p_a(s'_{j+1})$, which is a contradiction. Thus, we may conclude that $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$.

Finally, it is easy to verify that $\vec{\sigma}$ is a finite-state strategy. Since π is ultimately periodic, the ‘main body’ of $\vec{\sigma}$ is finite-state. Moreover, tracking whether any player has deviated yet requires only finite memory. If a player has deviated, then a punishing strategy is used by all remaining players, which is also finite-state. \square

3.4 Path Finding in Multi-Weighted Graphs with LEX(parity,mp) Payoffs

The following theorem, of interest in its own right, will be used to decide the existence of ultimately periodic paths in Proposition 5. A *multi-weighted graph* is a structure of the form $\mathcal{G} = (V, E, (w_a)_{a \in A}, (p_a)_{a \in A})$ where V is a finite set of states, $E \subseteq V^2$ a set of edges, A is a finite index set, and $w_a, p_a : V \rightarrow \mathbb{Z}$ are functions, one for each $a \in A$, mapping states to integers.

Theorem 3 *Given a multi-weighted graph $\mathcal{G} = (V, E, A, (w_a)_{a \in A}, (p_a)_{a \in A})$ over the finite index set A , a starting vertex $\iota \in V$, and a vector of payoffs $f \in \Omega^A$, one can decide in nondeterministic polynomial time whether there exists an ultimately periodic path $\pi = v_0 v_1 \dots$ in the graph with $\iota = v_0$ and, for every $a \in A$, $f_a \prec_{lex} \text{pay}_a(\pi)$.*

Proof W.l.o.g., we may assume that $f_a \in \{\top, \perp\} \times \{0\}$ (to see this, redefine $w_a(s)$ to be $w_a(s) - f_a$ for all $s \in V, a \in A$). Also, we may assume that every state in V is reachable from ι (to see this, restrict V to the states reachable from ι , computable in linear time). Finally, nondeterministically guess a vector $P \in \mathbb{Z}^A$. This vector represents the top priorities visited infinitely often for each index $a \in A$, and thus also determines the exact form of f .

Consider the subgraph, \mathcal{G}' formed by iterating through each $a \in A$ and all the states, and removing those states with a priority higher than P_a for some a . Now, if the original graph \mathcal{G} has some ultimately periodic path π with the top priorities being visited infinitely often given by P , then π is also a path in \mathcal{G}' . Thus, it suffices to form \mathcal{G}' and ask if there is some path π such that $f_a \prec_{lex} \text{pay}_a(\pi)$. In what follows, we simply relabel \mathcal{G}' as \mathcal{G} , on the understanding that the above transformation has taken place.

We now reduce the problem to finding certain cycles in \mathcal{G} . A *cycle* is a finite path C of the form $s_0 s_1 \dots s_n$ (for some $n \geq 1$) such that $s_0 = s_n$ (note that cycles are not necessarily simple). Write $s \in C$ to mean that $s = s_i$ for some $i \leq n$. Define $\text{sum}_a(C) = \sum_{j=1}^n w_a(s_j)$, $\text{avg}_a(C) = \frac{\text{sum}_a(C)}{n}$, and $\text{max}_a(C) = \max_{1 \leq j \leq n} p_a(s_j)$. The stated problem is equivalent to deciding, given W and $h : D \rightarrow \mathbb{Z}$, if there exists a cycle C in W such that i) $\text{max}_a(C) = h(a)$ for every $a \in D$, ii) $\text{avg}_a(C) > 0$ for every $a \in A$. Note that we can replace $\text{avg}_a(C)$ by $\text{sum}_a(C)$ in this problem (since $\text{avg}_a(C) > 0$ iff $\text{sum}_a(C) > 0$). In order to decide the existence of such a cycle, we adapt the proof from [31] that shows how to decide if there is a cycle C such that for every $a \in A$, $\text{sum}_a(C) = 0$.

A *multicycle* \mathcal{M} is a non-empty multiset of cycles. Thus a cycle is a multicycle \mathcal{M} with $|\mathcal{M}| = 1$. Extend sum_a and max_a to multicycles as follows: $\text{sum}_a(\mathcal{M}) = \sum_{C \in \mathcal{M}} \text{sum}_a(C)$ and $\text{max}_a(\mathcal{M}) = \max\{\text{max}_a(C) : C \in \mathcal{M}\}$. An η -*multicycle* is a multicycle \mathcal{M} such that for all $a \in A$, we have i) $\text{max}_a(\mathcal{M}) = P_a$, and ii) $\text{sum}_a(\mathcal{M}) > 0$. Additionally, an η -*cycle* is simply an η -multicycle, \mathcal{M} , with $|\mathcal{M}| = 1$ - that is, it consist of a single cycle.

Thus, deciding the problem started in the theorem is equivalent to deciding if there is an η -cycle \mathcal{M} . We now show that it is sufficient to decide if there is an η -multicycle.

Define a relation on V : $v \equiv w$ iff $v = w$ or there exists an η -multicycle \mathcal{M} and $C \in \mathcal{M}$ such that $v, w \in C$. Note that \equiv is an equivalence relation: indeed, if $u, v \in C$ for $C \in \mathcal{M}$, and $v, w \in C'$ for $C' \in \mathcal{M}'$ then $u, w \in C''$ for $C'' \in \mathcal{M}''$ where C'' is formed by tracing C from v to v and then tracing C' from v to v , and \mathcal{M}'' is $(\mathcal{M} \cup \mathcal{M}' \cup \{C''\}) \setminus \{C, C'\}$. Note that \mathcal{M}'' is an η -multicycle because $\text{sum}_a(C'') = \text{sum}_a(C) + \text{sum}_a(C')$ and $\max_a(C'') = \max\{\max_a(C), \max_a(C')\}$.

Suppose \equiv has index 1, *i.e.*, for all $v, w \in V$, $v \equiv w$. There are two cases - $|V| = 1$ and $|V| > 1$. First suppose that $|V| = 1$ with $V = \{v\}$. Then either v has a self-loop or it doesn't. If it does not, then there can be no η -cycle. If it does, and the weight of the v is strictly positive and its priority even, then it has an η -cycle. Otherwise, it does not.

Now suppose that $|V| > 1$. We claim that there exists an η -cycle. Indeed: for every $v, v' \in V$ let $\mathcal{M}_{v,v'}$ be an η -multicycle containing a cycle C that visits v and v' . Then $\mathcal{M} = \cup_{v,v' \in V} \mathcal{M}_{v,v'}$ is an η -multicycle such that (*): for every $v, v' \in V$ there exists $C \in \mathcal{M}$ such that $v, v' \in C$. We now define two transformations of multicycles $\mathcal{M} \mapsto \mathcal{M}'$ that maintain the following invariants: a) $\text{sum}_a(\mathcal{M}) = \text{sum}_a(\mathcal{M}')$ for $a \in A$, b) $\max_a(\mathcal{M}) = \max_a(\mathcal{M}')$ for $a \in D$, c) if \mathcal{M} satisfies (*) then so does \mathcal{M}' , d) $|\mathcal{M}'| < |\mathcal{M}|$ (*i.e.*, the number of cycles decreases). Thus, repeatedly applying these transformation results in an η -cycle.

First, if C occurs more than once in \mathcal{M} , say n times, then remove all occurrences of C from \mathcal{M} and add the single cycle formed by tracing C n -many times. Thus, we have that \mathcal{M} is a *set* of cycles (*i.e.*, not a proper multiset). Second, if \mathcal{M} is not a single cycle, take $C \neq C' \in \mathcal{M}$, $v \in C, v' \in C'$ and by (*) pick $D \in \mathcal{M}$ such that $v, v' \in D$. There are three cases: if $D \neq C, D \neq C'$ then form the cycle F by tracing C from v to v , then tracing “half” of D from v to v' , then tracing C' from v' to v' , and then tracing the “other half” of D from v' to v and let \mathcal{M}' be $\mathcal{M} \cup \{F\} \setminus \{C, C', D\}$; if $D = C$ (the case $D = C'$ is symmetric), then $v' \in C$ and thus form F by tracing C from v' to v' and then tracing C' from v' to v' , and let \mathcal{M}' be $(\mathcal{M} \cup \{F\}) \setminus \{C, C'\}$. Both transformations satisfy the invariants.

Thus, the following algorithm decides if there is an η -cycle (assuming one can decide if there exists an η -multicycle): if $|V| = 1$, check if the single node in V has a self-loop, a strictly positive weight and an even priority. If it does, output “yes”, otherwise, output “no”. If $|V| > 1$, compute \equiv for V ; if it has index 1 then output “yes”; else, for each equivalence class $X \in V/\equiv$, recurse on the subgraph induced by X . The algorithm is clearly sound, *i.e.*, if it outputs “yes” then there is indeed an η -cycle. To see that it is complete, note if that C is an η -cycle, then for all $v, w \in C$, $v \equiv w$; and thus C is contained in an \equiv -class.

Finally, we show how to decide if there exists an η -multicycle using linear programming. We temporarily make the assumption that all the nodes of the graph lie in the same strongly connected component (SCC). For every edge e introduce a variable x_e . Informally, the value x_e is the number of times that the edge e is used on an η -multicycle. Formally, let $\text{src}(e) = \{v \in V : \exists w e = (v, w) \in E\}$; $\text{trg}(e) = \{v \in V : \exists w e = (w, v) \in E\}$; $\text{out}(v) = \{e \in E : \text{src}(e) = v\}$ and $\text{in}(v) = \{e \in E : \text{trg}(e) = v\}$.

The linear program LP has the following inequalities and equations:

- Eq1: $x_e \geq 0$ for each edge e — this is a basic consistency criterion;
- Eq2: $\sum_{e \in E} x_e \geq 1$ — this ensures that at least one edge is chosen;
- Eq3: for each $a \in A$, $\sum\{w_a(\text{src}(e)) \cdot x_e : e \in E\} > 0$ — this enforces that the total sum is positive;
- Eq4: for each $a \in D$, $\sum\{x_e : p_a(\text{src}(e)) = P_a\} \geq 1$ — this ensures that the largest appearing priority for agent a is P_a ;

Eq5: for each $v \in V$, $\sum_{e \in \text{out}(v)} x_e = \sum_{e \in \text{in}(v)} x_e$ — this “preservation” condition says that the number of times one enters a vertex is equal to the number of times one leaves that vertex.

We now prove that there exists a η -multicycle if and only if the LP has an integer solution. Indeed, from left to right let \mathcal{M} a (non-empty) η -multicycle and let x_e be the number of occurrences of the edge e in any of the cycles in \mathcal{M} . Clearly, Eq1 and Eq2 are satisfied. Moreover, since \mathcal{M} a η -multicycle, it holds that both $\sum_a(\mathcal{M}) > 0$ and $\max_a(\mathcal{M}) = P_a$ for each $a \in A$, which implies Eq3 and Eq4, respectively. Finally, observe that every vertex in a cycle is entered and left an equal number of times. Therefore, being \mathcal{M} a multicycle implies that Eq5 is satisfied and so the LP has an integer solution.

From right to left, assume the LP has an integer solution $\{x_e\}_{e \in E}$. From Eq1, Eq2 we obtain that the solution provides that each edge e is counted x_e times. Moreover, by Euler’s Theorem [5, Theorem 1.6.3], Eq5 implies that the graph has an Eulerian cycle. Thus, the edges selected by the solution can be arranged in a multicycle \mathcal{M} . Finally, Eq3 and Eq4 guarantee that $\sum_a(\mathcal{M}) > 0$ and $\max_a(\mathcal{M}) = P_a$ for each $a \in A$, thus implying that \mathcal{M} is a η -multicycle.

Now, from [33], we obtain that the program LP has a solution in the reals iff it has a solution in the rationals [33]. Moreover, if $(x_e)_{e \in E}$ is a solution to LP and $k \in \mathbb{N} \setminus \{0\}$, then $(kx_e)_{e \in E}$ is also a solution to LP. Thus, the LP has a solution iff it has an integer solution. Thus, the LP has a solution iff the graph has an η -multicycle.

With this in place, we can now relax the assumption of the graph consisting of a strongly connected component. First, we can use Tarjan’s algorithm [42] to obtain the strongly connected components (SCCs) of the graph in linear time. Then for every non-trivial SCC that is reachable from the start node, we can use the above procedure to determine if it contains an η -cycle \mathcal{M} . As payoffs are prefix-independent with respect to paths, this implies that \mathcal{M} is an η -cycle for the whole graph.

Now, in terms of complexity, we need to non-deterministically guess a vector $P \in \mathbb{Z}^A$. Then we recursively compute the equivalence relation \equiv , which can be done with an application of Tarjan’s algorithm, in time $O(|V| + |E|)$, following by the application of $|E|$ linear programs (Cfr. [31]). Finally, observe that the size of the linear program is polynomial in the size of the graph. By an identical argument to [31], we can conclude that the problem can be determined in non-deterministic polynomial time. \square

3.5 Putting the Steps Together

We can now finish the proof of Theorem 1. Consider a rational $\epsilon \geq 0$ and a $\text{Lex}(\text{LTL}, \text{mp})$ game G . Throughout, let $|\gamma|$ and W denote the size of the largest goal and weight respectively.

First, by Proposition 2 we can transform G into a $\text{Lex}(\text{parity}, \text{mp})$ -game, G' . In this new game, the number of states, n' is at most doubly-exponential in the size of the goals, $|\gamma|$ and the number of priorities, $|\rho|$ is at most singly-exponential in the size of the goals.

Second, for every agent $a \in \text{Ag}$ and a state $s \in \text{St}$, we compute the punishing value $p_a(s)$ by means of Proposition 3. Constructing the corresponding mean-payoff parity games needed to do this can be done in time linear of the size of the game. Moreover, we can determine the value of these games in time $O(n^d \cdot (m + \text{MP} + \text{Parity}))$, where n is the number of vertices in the game graph, m the number of edges, d is the number of priorities, and MP and Parity denote the complexity of solving mean-payoff and parity games respectively. Solving mean-payoff games can be done in time $O(|V|^3 \cdot |E| \cdot W)$, where V is the set of

vertices of the game, E is the set of edges, and W is the largest weight [48]. Moreover, parity games can be solved by reducing them to mean-payoff games [30]. The reduction implies that these games can be solved in time $O(|V|^3 \cdot |E| \cdot |V|^{|\rho|})$, where V and E are the vertices and edges of the game graph, and $|\rho|$ is the size of the largest priority. Putting all this together, we see that calculating the punishment values of all players in all states can be done in time,

$$\begin{aligned} & |\text{Ag}| \cdot n' \cdot O\left(n'^{|\rho|} \cdot \left(n'^2 + n'^3 \cdot n'^2 \cdot W + n'^3 \cdot n'^2 \cdot n'^{|\rho|}\right)\right) \\ &= O\left(|\text{Ag}| \cdot W \cdot 2^{2^{q(|\gamma|)}}\right), \end{aligned}$$

where q is some appropriately chosen polynomial. Thus, the above step can be done in time doubly exponential in the size of the goals of the game and linear in the number of agents and the size of the largest weight.

Third, thanks to the characterization of finite-state strict ϵ -Nash Equilibria provided in Proposition 5, there is a $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$ if, and only if, there exists a tuple $\vec{z} \in \mathbb{Z}^{\text{Ag}}$ of values with $z_a \in \{p_a(s) : s \in \text{St}\}$, and a path π in G such that $z_a <_{\text{lex}} \text{pay}_a(\pi) + \epsilon$ for all $a \in \text{Ag}$, and for all $i \in \mathbb{N}$, $\pi_i = (s_i, \delta_i)$ is z_a -secure for a . Now, let $G[\vec{z}]$ denote the multi-weighted graph $(\text{St}, E, (\kappa_a)_{a \in \text{Ag}}, (\rho_a)_{a \in \text{Ag}})$ such that $(s, s') \in E$ iff there exists δ such that $\tau(s, \delta) = s'$ and (s, δ) is \vec{z} -secure for all $a \in \text{Ag}$. Observe that (s_i, δ_i) being z_a -secure for a , for every $a \in \text{Ag}$ and $i \in \mathbb{N}$, is equivalent to the fact that π is contained in $G[\vec{z}]$. Therefore, we reduced the problem of deciding whether $\text{FSNE}^\epsilon(G) \neq \emptyset$ to deciding whether there exists an ultimately periodic path $\pi \in G[\vec{z}]$ such that $z_a <_{\text{lex}} \text{pay}_a(\pi) + \epsilon$, for every $a \in \text{Ag}$. We solve this problem by guessing the top priorities visited by each player, and then employing the linear programming approach described in Theorem 3 on the multi-weighted graph $G[\vec{z}]$ and vector of payoffs $\vec{z} + \epsilon = (z_{a_1} + \epsilon, \dots, z_{a_n} + \epsilon)$. In terms of complexity, we need to iterate through all possible vectors of punishment values, of which there are $n'^{|\text{Ag}|}$, iterate through all possible priorities for each player, of which there are $|\rho|^{|\text{Ag}|}$, then compute the equivalence relation recursively by solving the corresponding linear programs. Letting q' be some appropriate polynomial, capturing the complexity of computing the equivalence relation, we see that this can all be done in time $O\left(n'^{|\text{Ag}|} \cdot |\rho|^{|\text{Ag}|} \cdot q'(n')\right)$. Expanding this out, and letting q'' be appropriately chosen polynomials, we can conclude that our path finding algorithm can be done in time,

$$O\left(2^{|\text{Ag}| \cdot 2^{q''(|\gamma|)}}\right)$$

Thus, in total, our algorithm to determine if a given game has a finite state strict ϵ Nash equilibrium can be done in time linear in the number of weights, exponential in the number of agents and doubly exponential in the size of the goals. As such, our algorithm lies in 2EXPTIME.

With this, we can now prove the main result of this work:

Theorem 4 *There is a 2EXPTIME algorithm that, given a rational $\epsilon \geq 0$, a $\text{Lex}(\text{LTL}, \text{mp})$ game $G = \langle A, (\kappa_a)_{a \in \text{Ag}}, \text{AP}, \lambda, (\gamma_a)_{a \in \text{Ag}} \rangle$ on arena $A = \langle \text{Ag}, \text{Act}, \text{St}, \iota, \tau \rangle$ and an LTL-formula Φ , decides whether there is $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$ such that $\pi_{\vec{\sigma}} \models \Phi$.*

Proof We show that the problem can be reduced to deciding whether there exists a Nash Equilibrium in a game G' defined over the arena $A' = \langle \text{Ag} \cup \{a_1, a_2\}, \text{Act}, \text{St} \times \{0, 1\}, (\iota, 0), \tau' \rangle$ in which a_1, a_2 are two new fresh agents and the transition function τ' is defined as follows.

$$\tau'((s, \varpi), \delta) = \begin{cases} (\tau(s, \delta_{\uparrow \text{Ag}}), 0) & \text{if } \delta(a_1) = \delta(a_2) \\ (\tau(s, \delta_{\uparrow \text{Ag}}), 1) & \text{otherwise} \end{cases}.$$

Then, consider a fresh atomic proposition p and define

$$G' = \langle A', (\kappa'_a)_{a \in \text{Ag}'}, \text{AP} \cup \{p\}, \lambda', (\gamma'_a)_{a \in \text{Ag}'} \rangle$$

such that $\kappa'_a = \kappa_a$ for every $a \in \text{Ag}$ and $\kappa'_{a_1}(c) = \kappa'_{a_2}(c) = 0$ for every action $c \in \text{Act}$, and $\lambda'(s, 0) = \lambda(s)$ and $\lambda'(s, 1) = \lambda(s) \cup \{p\}$, for every state $s \in \text{St}$. Finally, define $\gamma'_a = \gamma_a$, for every $a \in \text{Ag}$ and $\gamma'_{a_1} = \Phi \vee X p$ and $\gamma'_{a_2} = \Phi \vee X \neg p$.

Intuitively, the game G' results from pairing G with a two-player game played by agents a_1 and a_2 that are triggered to play against each other in case the formula Φ is not satisfied along the path.

Now, on one hand, let $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$ such that $\pi_{\vec{\sigma}} \models \Phi$ and consider a strategy profile $\vec{\sigma}'$ in G' such that $\vec{\sigma}'_{\uparrow \text{Ag}} = \vec{\sigma}$ ⁵. Clearly, each agent $a \in \text{Ag}$ takes the same sequence of actions in both strategy profiles, implying that $\text{mp}_a(\pi_{\vec{\sigma}}) = \text{mp}_a(\pi_{\vec{\sigma}'})$. Moreover, it is easy to see that $\lambda'(\pi_{\vec{\sigma}'})_{\uparrow \text{AP}} = \lambda(\pi_{\vec{\sigma}})$ ⁶, and so that $\pi_{\vec{\sigma}'} \models \gamma_a$ if and only if $\pi_{\vec{\sigma}} \models \gamma_a$ for every $a \in \text{Ag}$. This means that $\text{sat}_a(\pi_{\vec{\sigma}'}) = \text{sat}_a(\pi_{\vec{\sigma}})$ and so that $\text{pay}_a(\pi_{\vec{\sigma}'}) = \text{pay}_a(\pi_{\vec{\sigma}})$. Moreover, since $\pi_{\vec{\sigma}} \models \Phi$, it holds that $\pi_{\vec{\sigma}'} \models \gamma_{a_1}$ and $\pi_{\vec{\sigma}'} \models \gamma_{a_2}$ and that $\text{pay}_{a_1}(\pi_{\vec{\sigma}'}) = \text{pay}_{a_2}(\pi_{\vec{\sigma}'}) = (\top, 0)$.

Now, assume by contradiction that $\vec{\sigma}' \notin \text{FSNE}^\epsilon(G')$. First observe that, $(\top, 0)$ is the maximum payoff that players a_1 and a_2 can achieve in G' and so neither of them has an incentive to deviate from $\vec{\sigma}'$. Then, assume there is an agent $a \in \text{Ag}$ and a strategy σ''_a such that $\text{pay}_a(\pi_{\vec{\sigma}'}) + \epsilon \leq_{\text{lex}} \text{pay}_a(\pi(\vec{\sigma}'_{-a}, \sigma''_a))$. Then, we would also have that $\text{pay}_a(\pi_{\vec{\sigma}}) + \epsilon \leq_{\text{lex}} \text{pay}_a(\pi(\vec{\sigma}'_{-a}, \sigma''_a))$, in contradiction with the fact that $\vec{\sigma} \in \text{FSNE}^\epsilon(G)$.

On the other hand, assume that $\vec{\sigma}' \in \text{FSNE}^\epsilon(G')$. Then, by a symmetrical reasoning, we obtain that $\vec{\sigma} = \vec{\sigma}'_{\uparrow \text{Ag}} \in \text{FSNE}^\epsilon(G)$. Moreover, note that $\pi_{\vec{\sigma}'} \models \Phi$, otherwise, either $\text{pay}_{a_1}(\pi_{\vec{\sigma}'}) = (\perp, 0)$ or $\text{pay}_{a_2}(\pi_{\vec{\sigma}'}) = (\perp, 0)$, and there would exist a beneficial deviation for one of them, contradicting the fact that $\vec{\sigma}' \in \text{FSNE}^\epsilon(G')$. Hence, from $\pi_{\vec{\sigma}'} \models \Phi$ we obtain that $\pi_{\vec{\sigma}} \models \Phi$, which concludes the proof. \square

Theorem 5 *Deciding whether there exists a finite-state strict ϵ Nash Equilibrium in a given $\text{Lex}(\text{LTL}, \text{mp})$ game is 2EXPTIME-HARD .*

Proof We show a reduction from the problem of finding finite-state Nash Equilibria in (simple) LTL games, whose complexity is 2EXPTIME-COMplete (Cfr., see [24]). For an LTL game $G = (A, \text{AP}, \lambda, (\gamma_a)_{a \in \text{Ag}})$ on the arena $A = \langle \text{Ag}, \text{Act}, \text{St}, \iota, \tau \rangle$, consider the $\text{Lex}(\text{LTL}, \text{mp})$ game $G' = (A, (\kappa_a)_{a \in \text{Ag}}, \text{AP}, \lambda, (\gamma_a)_{a \in \text{Ag}})$ over the same arena and such that $\kappa_a(s) = 0$ for every $s \in \text{St}$. Intuitively, G' is the same as G but with a vacuous weighting added to match the game type of $\text{Lex}(\text{LTL}, \text{mp})$. In particular, note that every strategy σ_a in G for player a is also a strategy in G' for the same player, and vice-versa. At this point, for a non negative $\epsilon > 0$, and denoting the set of finite-state Nash Equilibria of an LTL game, G , by $\text{FNE}(G)$, we claim that $\text{FNE}(G) = \text{FSNE}^\epsilon(G')$. The proof is by double inclusion.

On one hand, let $\vec{\sigma} \in \text{FNE}(G)$ and assume, by contradiction that $\vec{\sigma} \notin \text{FSNE}^\epsilon(G')$. Then, there exists an agent $a \in \text{Ag}$ and a strategy σ'_a such that $\text{pay}_a(\pi_{\vec{\sigma}}) + \epsilon \leq_{\text{lex}} \text{pay}_a(\pi(\vec{\sigma}_{-a}, \sigma'_a))$.

⁵ This is an abuse of notation, as the transition functions of the two strategies are defined on a different set of decisions. Here we mean that the transition functions of $\vec{\sigma}'$ simply copy the ones in $\vec{\sigma}$ by ignoring the components expressed by agents a_1 and a_2

⁶ With another abuse of notation, we here mean the restriction of sequences in $\text{AP} \cup \{p\}$ to sequences in AP .

Now, observe that, since the mp value in G' for every player is always null, the above inequality can only apply when $\text{pay}_a(\pi_{\tilde{\sigma}}) = (\perp, 0)$ and $\text{pay}_a(\pi_{(\tilde{\sigma}_{-a}, \sigma'_a)}) = (\top, 0)$. Hence, we obtain that $\pi_{\tilde{\sigma}} \not\models \gamma_a$ and $\pi_{(\tilde{\sigma}_{-a}, \sigma'_a)} \models \gamma_a$ which contradicts the fact that $\tilde{\sigma} \in \text{FNE}(G)$.

On the other hand, let $\tilde{\sigma} \in \text{FSNE}^\epsilon(G')$ and assume, by contradiction, that $\tilde{\sigma} \notin \text{FNE}(G)$. Then, there exists an agent $a \in \text{Ag}$ and a strategy σ'_a such that $\pi_{\tilde{\sigma}} \not\models \gamma_a$ and $\pi_{(\tilde{\sigma}_{-a}, \sigma'_a)} \models \gamma_a$. Hence, in G' , we have that $\text{pay}_a(\pi_{\tilde{\sigma}}) + \epsilon = (\perp, \epsilon) \leq_{lex} (\top, 0) = \text{pay}_a(\pi_{(\tilde{\sigma}_{-a}, \sigma'_a)})$, in contradiction with the fact that $\tilde{\sigma} \in \text{FSNE}^\epsilon(G')$. \square

4 Related Work

Our work has its origins in several threads of research in AI and mainstream computer science. In computer science, the problem of checking whether a (typically finite state) system satisfies a specification expressed as a temporal logic formula has been a major research area since the 1970s [37]. In the 1980s, the introduction of the model checking paradigm meant that verification based on temporal logic became a practical possibility, leading to a rapid growth of interest in model checking both in the verification community and beyond [17]. In the late 1990s, attention began to move from the verification of *closed* systems to *open* systems. A key problem in the verification of open systems is that of determining whether a particular system or system component can *force* a property to hold: that is, whether there exists a strategy such that, by following that strategy, the system component can ensure that the property will be *guaranteed* to hold, irrespective of the behaviour of other system components. A logic called Alternating-time Temporal Logic (ATL) was introduced to explicitly support such reasoning about strategic ability [4]. ATL proved to be extremely influential, and was widely taken up within the multi-agent systems community. Although ATL embodies an important game theoretic concept—the notion of a winning strategy—it provides no mechanism for capturing preferences, and so its ability to capture game theoretic concepts beyond strategic ability—such as Nash equilibrium—is inherently limited. For this reason, research then began to shift to formalisms that could capture properties such as Nash equilibrium. Of these, Strategy Logic is currently the best-known example of such a formalism [16].

In the context of concurrent games and multi-agent systems, the main decision problem in this work concerns the existence of an equilibrium satisfying a system property; this problem is called “rational synthesis” [22] or “rational verification” (cf., E-NASH in [47, 25]) and includes equilibrium-emptiness as a special case. In this article, we studied this problem, specifically, for Lex(LTL, mp)-games and finite-state strict ϵ Nash equilibria (via a reduction Lex(parity, mp)-games), which includes, as a special case, strict Nash equilibria as well as Nash equilibria in games with LTL goals. All such problems can be solved in 2EXPTIME regardless of the particular setting. Most other work in rational synthesis concerns ordinary (i.e., not necessarily finite-state, nor strict) NE-emptiness. In particular, NE-emptiness has been studied for other objectives, notably mean-payoff (NP-complete) [43], Büchi (PTIME-complete) [9], and lexicographic order on Büchi objectives (in NP, but not known to be NP-complete) [9].

E-NASH (similar to what we call equilibrium-existence) for finite-state strategies has been studied on iterated Boolean Games (a simple form of infinite-duration multiplayer concurrent games) as follows: with LTL objectives, E-NASH is 2EXPTIME-complete [24]; with objective-LTL, i.e., each agent has to optimise a reward based on the truth value of a finite number of fixed LTL formulae, E-NASH is 2EXPTIME-complete [32]. A special case of objective-LTL is the lexicographic order on a finite number of components, each consisting of an LTL formula, also 2EXPTIME-complete.

Actually, these lower-bounds are inherited from the fact that solving two-player zero-sum games with LTL objectives is already $2\text{EXP}\text{TIME}$ -complete [39].

We remark that all these works (except objective-LTL and LTL[F], discussed below) concern equilibria concepts in multiplayer games with either qualitative or quantitative objectives, but not a combination, as we do.

Most of the work that combines more than one objective has mainly been studied for the restricted setting of two-player games. There work that considers a combination of quantitative objectives such as multi-mean-payoff and multi-energy games [45]. In the non-zero-sum case, secure-equilibria (in which each player tries to maximise their own payoff and then minimise their opponent's payoff) has been studied for a host of quantitative objectives, including mean-payoff [12].

In terms of work that has studied combinations of qualitative and quantitative objectives, as we do, there are a number of results. Again, such combinations have mainly been studied in the setting of two-player games. In the zero-sum case notable works combine the parity condition with mean-payoff objectives [15, 7] (we draw on results of [15] in our proofs) or with energy objectives [14]. On the other hand, there has been some work in the multi-agent setting, but in different settings. [2] studies the rational synthesis problem for an extension of LTL by quality operators. Contrary to our work, such a logical extension does not account qualitative and quantitative combination of objectives, but rather aims at measuring the degree of satisfaction of a given qualitative objective. In addition, it is not expressive enough to capture mean-payoff objectives. Objective-LTL combines Boolean objectives (given as LTL formulae) in a weak way, *i.e.*, there are only finitely many possible payoffs. In contrast, $\text{Lex}(\text{LTL}, \text{mp})$ combines qualitative objectives (given as LTL formulae) with quantitative objectives (given as mean-payoff objectives), and thus result in infinitely many possible payoffs.

Our work is somewhat related to the paradigm of *Boolean games* [28, 8, 46, 23, 19]. A Boolean game is a non-cooperative game played over a set of Boolean variables. Each player in such a game desires the satisfaction of a goal, specified as a logical formula over the overall set of variables, and is assumed to control a subset of the variables: the choices available to a player correspond to the assignments that can be made to the variables controlled by that player. Players simultaneously choose valuations for the variables they control, and a player is satisfied if their goal is made true by the resulting overall valuation. In addition to being an interesting game-theoretic model in their own right, Boolean games are a natural abstract model for studying strategic behaviour in multi-agent systems. Specifically related to our setting is the generalisation of Boolean games known as *Iterated Boolean Games*, in which goals are specified as LTL formulae, and the game takes place over an infinite number of rounds [24]. As with our approach, most work in Boolean games assumes pure strategies, although some attention has recently been given to mixed strategies [29].

Finally, it is also worth mentioning work on multi-agent planning models such as the multi-agent STRIPS model [11, 10]. These frameworks, model systems where each agent is an individual planning system, attempting to achieve an individual goal. The relationship between these frameworks and concurrent games was investigated by Gutierrez *et al.* [25].

5 Conclusion

In the last twenty years, significant effort was devoted to analyze qualitative aspects of multi-agent systems, and more recently, also quantitative aspects. However, the two settings are often investigated independently. As this is not appropriate in many natural scenarios,

researchers began to look at the combination of these two worlds. The achievements to date in this direction, however, are far from satisfactory, either because the settings are too weak, (*e.g.*, they cannot model important solution concepts such as Nash Equilibrium [13]), or because they are too expensive in terms of complexity, (*e.g.*, between EXPTIME-hard and undecidable [1]).

In this paper we introduce a model of multi-agent systems in which each agent’s payoff is a lexicographic combination of qualitative (LTL) and quantitative (mean-payoff) payoffs. We call these Lex(LTL, mp) games. The solution concept we focus on is finite-state strict ϵ Nash equilibria (for $\epsilon \geq 0$). In this setting, we proved that the rational synthesis problem (a generalisation of the equilibrium existence problem) is decidable, and moreover is in NP when the qualitative goals are represented as parity conditions and is 2EXPTIME-complete when they are given by LTL formulae. The proof characterises the equilibrium executions as certain ultimately periodic paths in a multi-weighted graph. To compute this graph we solve two-player zero-sum games with lexicographic objectives, and to find paths in such graphs we use Linear Programming. The question as to whether the main result also holds for ordinary equilibria (*i.e.*, non-strict) in games with LTL goals has been left open. Indeed, our choice of using strict equilibria allowed us to supply the characterisation of such equilibria in Proposition 5, and a first-step towards handling ordinary equilibria would be to supply an analogous characterisation. Another interesting question is about the strategic ability of the agents. Here, we analysed the case of finite-state strategies: this is motivated by the intention to provide implementable solutions for AI and multi-agent systems communities. However, from the theoretical perspective it is interesting to consider the infinite-state case. We can already deduce that this case is not equivalent to the finite-state one, as our setting includes the one of two-player parity mean-payoff games, for which such an inequivalence is proved in [15]. In addition to this, a new technique would be required for addressing that case. As a matter of fact, infinite-state strategies do not guarantee an ultimately periodic path outcome. Therefore, the multicycle approach misses the non-periodic solutions.

For future work we foresee a number of possibilities. Some immediate questions include investigating whether our results extend to non-strict Nash equilibria (the question we left open) and whether they still hold in variations of the game, for instance, in games with more complex preference relations or in games played in different arenas where the main complexity results may change. Given the 2EXPTIME nature of the main problem, we can also investigate to what extent the main problem becomes easier in restricted settings. For instance, typical cases of study include games with memoryless strategies or with simpler classes of goals, such as properties that can be described in fragments of LTL. Many questions can also be asked with respect to the quantitative component of our games, *e.g.*, whether there is a Nash equilibrium in which each player, or some designated set of players more generally, can be ensured a mean-payoff within a certain interval. All of these questions constitute important avenues for further work.

6 Acknowledgments

The authors would like to thank the reviewers for their careful reading of preliminary versions of this manuscript. This allowed us to significantly improve the quality and solidity of our results.

References

1. N. Alechina, B. Logan, N.H. Nga, and F. Raimondi. Model-Checking for Resource-Bounded ATL with Production and Consumption of Resources. *CoRR*, abs/1504.06766, 2015.
2. Shaull Almagor, Orna Kupferman, and Giuseppe Perelli. Synthesis of controllable Nash equilibria in quantitative objective game. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 35–41, 2018.
3. R. Alur. *Principles of Cyber-Physical Systems*. The MIT Press: Cambridge, MA, 2015.
4. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *Journal of the ACM*, 49(5):672–713, 2002.
5. J. Bang-Jensen and G. Gutin. *Digraphs: theory, algorithms and applications*. Springer, 2008.
6. K. Binmore. *Fun and Games: A Text on Game Theory*. D. C. Heath and Company: Lexington, MA, 1992.
7. R. Bloem, K. Chatterjee, T. A. Henzinger, and B. Jobstmann. Better Quality in Synthesis through Quantitative Objectives. In *CAV'09*, pages 140–156, 2009.
8. E. Bonzon, M.-C. Lagasquie, J. Lang, and B. Zanuttini. Boolean games revisited. In *Proceedings of the Seventeenth European Conference on Artificial Intelligence (ECAI-2006)*, Riva del Garda, Italy, 2006.
9. P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Pure Nash Equilibria in Concurrent Deterministic Games. *Logical Methods in Computer Science*, 11(2), 2015.
10. R. Brafman and C. Domshlak. On the complexity of planning for agent teams and its implications for single agent planning. *Artificial Intelligence*, 198:52–71, 2013.
11. R. Brafman, C. Domshlak, Y. Engel, and M. Tennenholtz. Planning games. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, 2009.
12. V. Bruyère, N. Meunier, and J.-F. Raskin. Secure equilibria in weighted games. In *CSL-LICS'14*, pages 26:1–26:26, 2014.
13. N. Bulling and V. Goranko. How to Be Both Rich and Happy: Combining Quantitative and Qualitative Strategic Reasoning about Multi-Player Games (Extended Abstract). In *SR'13*, pages 33–41, 2013.
14. K. Chatterjee and L. Doyen. Energy parity games. *Theor. Comput. Sci.*, 458:49–60, 2012.
15. K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Mean-Payoff Parity Games. In *LICS'05*, pages 178–187, 2005.
16. K. Chatterjee, T.A. Henzinger, and N. Piterman. Strategy logic. *Information and Computation*, 208(6):677–693, June 2010.
17. E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. The MIT Press: Cambridge, MA, 2000.
18. S. Demri, V. Goranko, and M. Lange. *Temporal Logics in Computer Science*. Cambridge University Press: Cambridge, England, 2017.
19. P. E. Dunne, S. Kraus, W. van der Hoek, and M. Wooldridge. Cooperative Boolean games. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2008)*, Estoril, Portugal, 2008.
20. A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.
21. E. A. Emerson. Temporal and modal logic. In *Handbook of Theoretical Computer Science Volume B: Formal Models and Semantics*, pages 996–1072. Elsevier Science Publishers B.V.: Amsterdam, The Netherlands, 1990.
22. D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *TACAS'10*, pages 190–204. Springer, 2010.
23. J. Grant, S. Kraus, M. Wooldridge, and I. Zuckerman. Manipulating Boolean games through communication. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, Barcelona, Catalonia, Spain, 2011.
24. J. Gutierrez, P. Harrenstein, and M. Wooldridge. Iterated Boolean games. *Inf. Comput.*, 242:53–79, 2015.
25. J. Gutierrez, P. Harrenstein, and M. Wooldridge. From model checking to equilibrium checking: Reactive modules for rational verification. *Artificial Intelligence*, 248:123 – 157, 2017.
26. Julian Gutierrez, Paul Harrenstein, Giuseppe Perelli, and Michael J. Wooldridge. Nash equilibrium and bisimulation invariance. *Logical Methods in Computer Science*, 15(3), 2019.
27. Julian Gutierrez, Aniello Murano, Giuseppe Perelli, Sasha Rubin, and Michael J. Wooldridge. Nash equilibria in concurrent games with lexicographic preferences. In *IJCAI*, pages 1067–1073. ijcai.org, 2017.
28. P. Harrenstein, W. van der Hoek, J.-J.Ch. Meyer, and C. Witteveen. Boolean games. In J. van Benthem, editor, *Proceeding of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 287–298, Siena, Italy, 2001.
29. E. Iarovski and L. Ong. The complexity of decision problems about equilibria in two-player Boolean games. *Artificial Intelligence*, 261:1–15, 2018.

30. Marcin Jurdzinski. Deciding the winner in parity games is in $UP \cap co-UP$. *Inf. Process. Lett.*, 68(3):119–124, 1998.
31. S. R. Kosaraju and G. Sullivan. Detecting Cycles in Dynamic Graphs in Polynomial Time. In *STOC'88*, pages 398–406. ACM, 1988.
32. O. Kupferman, G. Perelli, and M.Y. Vardi. Synthesis with Rational Environments. *Ann. Math. Artif. Intell.*, 78(1):3–20, 2016.
33. J. Matousek and B. Gärtner. *Understanding and Using Linear Programming*. Springer Science & Business Media, 2007.
34. Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theor. Comput. Sci.*, 81(2):317–324, 1991.
35. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.
36. N. Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3(3), 2007.
37. A. Pnueli. The temporal logic of programs. In *Proceedings of the Eighteenth IEEE Symposium on the Foundations of Computer Science*, pages 46–57, 1977.
38. A. Pnueli and R. Rosner. On the Synthesis of a Reactive Module. In *POPL'89*, pages 179–190. ACM, 1989.
39. R. Rosner. *Modular Synthesis of Reactive Systems*. PhD thesis, Weizmann, 1991.
40. Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press: Cambridge, England, 2008.
41. John Maynard Smith. *Evolution and the Theory of Games*. Cambridge university press, 1982.
42. Robert Endre Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.
43. M. Ummels and D. Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games. In *CONCUR'11*, pages 482–496, 2011.
44. M. Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Logics for Concurrency - Structure versus Automata (8th Banff Higher Order Workshop, August 27 - September 3, 1995, Proceedings)*, pages 238–266, 1995.
45. Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Moshe Rabinovich, and Jean-François Raskin. The complexity of multi-mean-payoff and multi-energy games. *Inf. Comput.*, 241:177–196, 2015.
46. M. Wooldridge, U. Endriss, S. Kraus, and J. Lang. Incentive engineering for Boolean games. *Artificial Intelligence*, 195:418–439, 2013.
47. M. Wooldridge, J. Gutierrez, P. Harrenstein, E. Marchioni, G. Perelli, and A. Toumi. Rational Verification: From Model Checking to Equilibrium Checking. In *AAAI'16*, pages 4184–4191, 2016.
48. U. Zwick and M. Paterson. The Complexity of Mean Payoff Games on Graphs. *TCS*, 158(1&2):343–359, 1996.