

THE COMPLEXITY OF NASH EQUILIBRIA IN STOCHASTIC MULTIPLAYER GAMES*

MICHAEL UMMELS^a AND DOMINIK WOJTCZAK^b

^a RWTH Aachen University, Germany
e-mail address: ummels@logic.rwth-aachen.de

^b CWI Amsterdam, The Netherlands
e-mail address: d.k.wojtczak@cwi.nl

ABSTRACT. We analyse the computational complexity of finding Nash equilibria in turn-based stochastic multiplayer games with ω -regular objectives. We show that restricting the search space to equilibria whose payoffs fall into a certain interval may lead to undecidability. In particular, we prove that the following problem is undecidable: Given a game \mathcal{G} , does there exist a Nash equilibrium of \mathcal{G} where player 0 wins with probability 1? Moreover, this problem remains undecidable when restricted to pure strategies or (pure) strategies with finite memory. One way to obtain a decidable variant of the problem is to restrict the strategies to be positional or stationary. For the complexity of these two problems, we obtain a common lower bound of NP and upper bounds of NP and PSPACE respectively. Finally, we single out a special case of the general problem that, in many cases, admits an efficient solution. In particular, we prove that deciding the existence of an equilibrium in which each player either wins or loses with probability 1 can be done in polynomial time for games where the objective of each player is given by a parity condition with a bounded number of priorities.

1. INTRODUCTION

We study *stochastic games* [53] played by multiple players on a finite, directed graph. Intuitively, a play of such a game evolves by moving a token along edges of the graph: Each vertex of the graph is either controlled by one of the players, or it is *stochastic*. Whenever the token arrives at a non-stochastic vertex, the player who controls this vertex must move the token to a successor vertex; when the token arrives at a stochastic vertex, a fixed probability distribution determines the next vertex. A measurable function maps plays to payoffs. In the simplest case, which we discuss here, the possible payoffs of a single play are 0 and 1 (i.e. each player either wins or loses a given play). However, due to the presence

1998 ACM Subject Classification: F.1.2, G.1.6, G.3.

Key words and phrases: Nash equilibria, Stochastic games, Computational complexity.

* Preliminary versions of parts of this paper appeared in the *Proceedings of the 36th International Colloquium on Automata, Languages and Programming* (ICALP 2009) and the *Proceedings of the 18th Annual Conference of the European Association for Computer Science Logic* (CSL 2009). This work was supported by the DFG Research Training Group 1298 (ALGO SYN) and the ESF Research Networking Programme “Games for Design and Verification” (GAMES).

of stochastic vertices, a player's *expected payoff* (i.e. her probability of winning) can be an arbitrary probability.

Stochastic games with ω -regular objectives have been used as a formal model for the verification and synthesis of reactive systems under the influence of random events [5]. Such a system is usually modelled as a game between the system and its environment, where the environment's objective is the complement of the system's objective: the environment is considered hostile. Therefore, the research in this area has traditionally focused on two-player games where each play is won by precisely one of the two players, so-called *two-player zero-sum games*. However, the system may consist of several components with independent objectives, a situation which is naturally modelled by a multiplayer game.

The most common interpretation of rational behaviour in multiplayer games is captured by the notion of a *Nash equilibrium* [52]. In a Nash equilibrium, no player can improve her payoff by unilaterally switching to a different strategy. Chatterjee et al. [16] gave an algorithm for computing a Nash equilibrium in a stochastic multiplayer game with ω -regular winning conditions. However, it can be shown that their algorithm may compute an equilibrium where all players lose almost surely (i.e. receive expected payoff 0), even when there exist other equilibria where all players win almost surely (i.e. receive expected payoff 1).

In applications, one might look for an equilibrium where as many players as possible win almost surely or where it is guaranteed that the expected payoff of the equilibrium falls into a certain interval. Formulated as a decision problem, we want to know, given a k -player game \mathcal{G} with initial vertex v_0 and two thresholds $\bar{a}, \bar{b} \in [0, 1]^k$, whether (\mathcal{G}, v_0) has a Nash equilibrium with expected payoff at least \bar{a} and at most \bar{b} . This problem, which we call NE for short, is a generalisation of the *quantitative decision problem* for two-player zero-sum games, which asks whether in such a game player 0 has a strategy that ensures to win the game with a probability that exceeds a given threshold.

The problem NE comes in several variants, depending on the type of strategies one considers: On the one hand, strategies may be *randomised* (allowing randomisation over actions) or *pure* (not allowing such randomisation). On the other hand, one can restrict to strategies that use (unbounded or bounded) finite memory or even to *stationary* ones (strategies that do not use any memory at all). For the quantitative decision problem, this distinction is often not meaningful since in a two-player zero-sum simple stochastic game with ω -regular objectives both players have optimal pure strategies with finite memory. Moreover, in many games even positional (i.e. both pure and stationary) strategies suffice for optimality. However, regarding NE this distinction leads to distinct decision problems, which have to be analysed separately.

Our main result is that NE is undecidable if we allow either arbitrary randomised strategies or arbitrary pure strategies. In fact, even the following, presumably simpler, problem is undecidable: Given a game \mathcal{G} , decide whether there exists a Nash equilibrium (in pure strategies) where player 0 wins almost surely. Moreover, the problem remains undecidable if one restricts to randomised or pure strategies with finite memory.

If we restrict to simpler types of strategies like stationary ones, NE becomes decidable. In particular, for positional strategies the problem is typically NP-complete, and for arbitrary stationary strategies it is NP-hard but typically contained in PSPACE. To get a better understanding of the latter problem, we also relate it to the *square root sum problem* (SqrtSum) by providing a polynomial-time reduction from SqrtSum to NE with the restriction to stationary strategies. It is a long-standing open problem whether SqrtSum falls into the polynomial hierarchy; hence, showing that NE for stationary strategies lies inside

the polynomial hierarchy would imply a breakthrough in understanding the complexity of numerical computations.

Finally, we prove decidability for an important fragment of NE, which we call the strictly qualitative fragment. This fragment arises from NE by restricting the two thresholds to be the same binary payoff. Hence, we are only interested in equilibria where each player either wins or loses with probability 1. Formally, the task is to decide, given a k -player game \mathcal{G} with initial vertex v_0 and a binary payoff $\bar{x} \in \{0, 1\}^k$, whether the game has a Nash equilibrium with expected payoff \bar{x} . Apart from proving decidability, we show that, depending on the representation of the objective, this problem is typically complete for one of the complexity classes P, NP, $P^{\text{NP}[\log]}$ and PSPACE, and that the problem is invariant under restricting the search space to equilibria in pure finite-state strategies.

Outline. In Section 2, we introduce the model that underlies this work and survey earlier work on stochastic two-player zero-sum games. In Section 3, we prove that every stochastic multiplayer game has a Nash equilibrium, thereby addressing an inaccuracy in an earlier proof by Chatterjee et al. [16]. In Section 4, we analyse the complexity of the problem NE with respect to the six modes of strategies we consider in this work: positional strategies, stationary strategies, pure finite-state strategies, randomised finite-state strategies, arbitrary pure strategies, and arbitrary randomised strategies. Finally, in Section 5, we prove that the strictly qualitative fragment of NE is decidable and analyse its complexity.

Related Work. Determining the complexity of Nash equilibria has attracted much interest in recent years. In particular, a series of papers culminated in the result that computing a Nash equilibrium of a two-player game in strategic form is complete for the complexity class PPAD [23, 18]. More in the spirit of our work, Conitzer and Sandholm [20] showed that deciding whether there exists a Nash equilibrium in a two-player game in strategic form where player 0 receives payoff at least x and related decision problems are all NP-hard. For non-stochastic infinite games, a qualitative version of the problem NE was studied in [58]. In particular, it was shown that the problem is NP-complete for games with parity winning conditions but in P for games with Büchi winning conditions.

For stochastic games, most results concern the computation of values and optimal strategies; see Section 2 for a survey of the most important results. In the multiplayer case, Chatterjee et al. [16] showed that the problem of deciding whether a (concurrent) stochastic game with reachability objectives has a Nash equilibrium in positional strategies with payoff at least \bar{x} is NP-complete. We sharpen their hardness result by demonstrating that the problem remains NP-hard when it is restricted to games with only three players (as opposed to an unbounded number of players) where payoffs are assigned at terminal vertices only (cf. Theorem 4.4).

A more restricted model of stochastic games, where questions like ours have been studied, are *Markov decision processes* (MDPs) with multiple objectives. These games can be considered as stochastic games where only one player can influence the outcome of the game. For MDPs with multiple ω -regular objectives, Etessami et al. [31] showed that questions like the one we ask are decidable. Their result relies on the fact that, for MDPs with multiple reachability objectives on terminal states, stationary strategies suffice to achieve a payoff that is higher than a given threshold. Unfortunately, this property does not extend to our model: we give an example of a stochastic game with the same kind of objectives

where every Nash equilibrium with payoff 1 for the first player requires infinite memory (see Proposition 4.13).

2. STOCHASTIC GAMES

2.1. Basic definitions. Let us start by giving a formal definition of the game model that underlies this paper. The games we are interested in are played by multiple players taken from a finite set Π of players; we usually refer to them as player 0, player 1, player 2, and so on.

The *arena* of the game is basically a directed, coloured graph. Intuitively, the players take turns to form an infinite path through the arena, a *play*. Additionally, there is an element of chance involved: at some vertices, it is not a player who decides how to proceed but *nature* who chooses a successor vertex according to a probability distribution. To model this scenario, we partition the set V of vertices into sets V_i of vertices *controlled by player* $i \in \Pi$ and a set of *stochastic vertices*, and we extend the edge relation to a transition relation that takes probabilities into account. Formally, an arena for a game with players in Π consists of:

- a countable, non-empty set V of *vertices* or *states*,
- for each player i a set $V_i \subseteq V$ of vertices *controlled* by player i ,
- a *transition relation* $\Delta \subseteq V \times ([0, 1] \cup \{\perp\}) \times V$, and
- a *colouring function* $\chi: V \rightarrow C$ into an arbitrary set C of colours.

We make the assumption that every vertex is controlled by at most one player: $V_i \cap V_j = \emptyset$ if $i \neq j$; vertices that are not controlled by a player are stochastic. For technical reasons, we also assume that for each vertex v the set

$$v\Delta := \{w \in V : \text{there exists } p \in (0, 1] \cup \{\perp\} \text{ such that } (v, p, w) \in \Delta\}$$

of possible successor vertices is finite and non-empty. Moreover, we require that probabilities appear only on transitions originating in stochastic vertices (if $v \in \bigcup_{i \in \Pi} V_i$ and $(v, p, w) \in \Delta$ then $p = \perp$) and that they are unique: for every pair of a stochastic vertex v and an arbitrary vertex w there exists precisely one $p \in [0, 1]$ such that $(v, p, w) \in \Delta$; we denote this probability by $\Delta(w \mid v)$. For computational purposes, we assume that these probabilities are rational numbers. Finally, for each stochastic vertex v the probabilities on outgoing transitions must sum up to 1: $\sum_{w \in V} \Delta(w \mid v) = 1$. Hence, if v is a stochastic vertex, then the mapping $V \rightarrow [0, 1]: w \mapsto \Delta(w \mid v)$ is a discrete probability distribution over V ; we denote the set of all discrete probability distributions over V by $\mathcal{D}(V)$.

The description of a game is completed by specifying an *objective* for each player. On an abstract level, these are just arbitrary sets of infinite sequences of colours, i.e. subsets of C^ω . Since we want to assign a probability to them, we assume that objectives are Borel sets over the usual topology on infinite sequences, if not stated otherwise. Since objectives specify which plays are winning for a player, they are also called *winning conditions*.

In general, we will identify an objective $\text{Win} \subseteq C^\omega$ over colours with the corresponding objective $\chi^{-1}(\text{Win}) := \{\pi \in V^\omega : \chi(\pi) \in \text{Win}\} \subseteq V^\omega$ over vertices (which is also Borel since χ , as a mapping $V^\omega \rightarrow C^\omega$, is continuous). The reason that we allow objectives to refer to a colouring of the vertices is that the number of colours can be much smaller than the number of vertices, and it is possible that an objective can be represented more succinctly as an objective over colours rather than as an objective over vertices.

If Π is a finite set of players, $(V, (V_i)_{i \in \Pi}, \Delta, \chi)$ is an arena and $(\text{Win}_i)_{i \in \Pi}$ is a collection of objectives, we refer to the tuple $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, \Delta, \chi, (\text{Win}_i)_{i \in \Pi})$ as a *stochastic multiplayer game (SMG)*. An SMG is *finite* if the set V of vertices is finite.

A *play* of \mathcal{G} is an infinite path through the arena of \mathcal{G} , i.e. a sequence $\pi = \pi(0)\pi(1)\dots$ of vertices such that for each $k \in \mathbb{N}$ there exists $p \in (0, 1] \cup \{\perp\}$ with $(\pi(k), p, \pi(k+1)) \in \Delta$. Finite prefixes of plays are called *histories*. We say that a play π of \mathcal{G} is *won* by player i if the corresponding sequence of colours fulfils player i 's objective, i.e. $\chi(\pi) \in \text{Win}_i$; the *payoff* of a play π is the vector $\bar{x} \in \{0, 1\}^\Pi$ defined by $x_i = 1$ if and only if $\chi(\pi) \in \text{Win}_i$.

Often, it is convenient to designate an initial vertex $v_0 \in V$; we denote the pair (\mathcal{G}, v_0) an *initialised SMG*. A play or a history of an initialised SMG (\mathcal{G}, v_0) is just a play respectively a history of \mathcal{G} that starts in v_0 . In the following, we will refer to both SMGs and initialised SMGs as SMGs; it should always be clear from the context whether the game is initialised or not.

SMGs generalise various stochastic models, each of them the subject of intensive research. First, there are *Markov chains*, the basic model for stochastic processes, in which no control is possible. These are just SMGs where the set Π of players is empty and (consequently) there are only stochastic vertices. If we extend Markov chains by a single controller, we arrive at the model of a *Markov decision process (MDP)*, a model introduced by Bellman [6] and heavily used in operations research. Formally, an MDP is an SMG where there is only one player (and only one objective). Finally, in a (*perfect-information*) *stochastic two-player zero-sum game (S2G)*, there are only two players, player 0 and player 1, who have opposing objectives: one player wants to fulfil an objective, while the other one wants to prevent her from doing so. Hence, one player's objective is the complement of the other player's objective. Due to their competitive nature, these games are also known as *competitive Markov decision processes* [32].

The SMG model also incorporates several non-stochastic models. In particular, we call an SMG *deterministic* if it contains no stochastic vertices. In the two-player zero-sum setting, the resulting model has found applications in logic and controller synthesis, to name a few.

2.2. Objectives. We have introduced objectives as abstract sets of infinite sequences. In order to be amenable for algorithmic solutions, we need to restrict to a class of objectives representable by finite objects. The objectives we consider for this purpose are standard in logic and verification (see [37]); for all of them, we require that the set C of colours the objective refers to is *finite*. Moreover, whether an infinite sequence α fulfils such an objective only depends on the set $\text{Occ}(\alpha)$ of colours occurring in α or on the set $\text{Inf}(\alpha)$ of colours occurring *infinitely often* in α . In particular, we deal with the following types of objectives:

- A *reachability objective* is given by a set $F \subseteq C$ of *good* colours, and the objective requires that a good colour is seen at least once. The corresponding subset of C^ω is $\text{Reach}(F) := \{\alpha \in C^\omega : \text{Occ}(\alpha) \cap F \neq \emptyset\}$.
- A *Büchi objective* is again given by a set $F \subseteq C$ of good colours, but it requires that a good colour is seen infinitely often. The corresponding subset of C^ω is $\text{Büchi}(F) := \{\alpha \in C^\omega : \text{Inf}(\alpha) \cap F \neq \emptyset\}$.

- A *co-Büchi objective* is also given by a set $F \subseteq C$ of good colours; this time, the objective requires that *from some point onwards* only good colours are seen. The corresponding subset of C^ω is $\text{coBüchi}(F) = \{\alpha \in C^\omega : \text{Inf}(\alpha) \subseteq F\}$.
- A *parity objective* is given by a *priority function* $\Omega: C \rightarrow \{0, \dots, d\}$, where $d \in \mathbb{N}$, which assigns to each colour a certain *priority*. The objective requires that the least priority that occurs infinitely often is even. The corresponding subset of C^ω is $\text{Parity}(\Omega) = \{\alpha \in C^\omega : \min(\text{Inf}(\Omega(\alpha))) \text{ is even}\}$.
- A *Streett objective* is given by a set Ω of *Streett pairs* (F, G) , where $F, G \subseteq C$. The objective requires that, for each of the pairs, if a colour on the left-hand side is seen infinitely often, then so is a colour on the right-hand side. The corresponding subset of C^ω is $\text{Streett}(\Omega) = \{\alpha \in C^\omega : \text{Inf}(\alpha) \cap F = \emptyset \text{ or } \text{Inf}(\alpha) \cap G \neq \emptyset \text{ for all } (F, G) \in \Omega\}$.
- A *Rabin objective* is given by a set Ω of *Rabin pairs* (F, G) , where $F, G \subseteq C$; it requires that for some pair a colour on the left-hand side is seen infinitely often while all colours on the right-hand side are seen only finitely often. The corresponding subset of C^ω is $\text{Rabin}(\Omega) = \{\alpha \in C^\omega : \text{Inf}(\alpha) \cap F \neq \emptyset \text{ and } \text{Inf}(\alpha) \cap G = \emptyset \text{ for some } (F, G) \in \Omega\}$.
- A *Muller objective* is given by a family \mathcal{F} of *accepting sets* $F \subseteq C$, and it requires that the set of colours seen infinitely often equals one of these accepting sets. The corresponding subset of C^ω is $\text{Muller}(\mathcal{F}) = \{\alpha \in C^\omega : \text{Inf}(\alpha) \in \mathcal{F}\}$.

Parity, Streett, Rabin and Muller objectives are of particular relevance because they provide a standard form for arbitrary ω -regular objectives: any game with arbitrary ω -regular objectives can be reduced to one with parity, Streett, Rabin or Muller objectives (over a larger arena) by taking the product of its original arena with a suitable deterministic word automaton for each player's objective [56].

In this work, for reasons that will become clear later, we are particularly attracted to objectives that are invariant under adding and removing finite prefixes; we call such objectives *prefix-independent*. More formally, an objective is prefix-independent if for each $\alpha \in C^\omega$ and $x \in C^*$ the sequence α satisfies the objective if and only if the sequence $x \cdot \alpha$ does. From the objectives listed above, only reachability objectives are, in general, not prefix-independent. However, many of our results (in particular, many of our lower bounds) apply to games with a prefix-independent form of reachability, which we call *terminal reachability*. For these objectives, we assume that each vertex is coloured by itself, i.e. $C = V$, and χ is the identity mapping. The terminal reachability objective for a set $F \subseteq V$ coincides with the reachability objective for F , but we require that each $v \in F$ is a *terminal vertex*: $v\Delta = \{v\}$. For any such set F , we have $\text{Occ}(\pi) \cap F \neq \emptyset$ if and only if $\text{Inf}(\pi) \cap F \neq \emptyset$ for every play π . Hence, terminal reachability objectives can be regarded as prefix-independent objectives.

For S2Gs, the distinction between reachability and terminal reachability is not important: every S2G with a reachability objective can easily be transformed into an equivalent S2G with a reachability objective on terminal states. For SMGs, we believe that any such transformation requires exponential time: deciding whether in a deterministic game with terminal reachability objectives there exists a play that fulfils each of the objectives can be done in polynomial time, whereas the same problem is NP-complete for deterministic games with standard reachability objectives [16, 57].

The resulting hierarchy of objectives is depicted in Fig. 1. As explained above, a simple reachability objective can be viewed as a (co-)Büchi objective. Any (co-)Büchi objective is equivalent to a parity objective with only two priorities, and any parity objective is equivalent to both a Streett and a Rabin objective; in fact, the intersection (union) of two parity objectives is equivalent to a Streett (Rabin) objective. Moreover, any Streett or




FIGURE 1. A hierarchy of prefix-independent objectives.

Rabin objective is equivalent to a Muller objective, although the translation from a set of Streett/Rabin pairs to an equivalent family of accepting sets is, in general, exponential. Finally, the complement of a Büchi (Streett) objective is equivalent to a co-Büchi (Rabin) objective, and vice versa, whereas the complement of a parity (Muller) objective is also a parity (Muller) objective. In fact, any objective that is equivalent to both a Streett and a Rabin objective is equivalent to a parity objective [63].

To denote the class of SMGs (S2Gs) with a certain type of objectives, we prefix the name SMG (S2G) with the name(s) of the objective; for instance, we use the term *Streett-Rabin SMG* to denote SMGs where each player has a Streett or a Rabin objective. For S2Gs, we adopt the convention to name the objective of player 0 first; hence, in a Streett-Rabin S2G player 0 has a Streett objective, while player 1 has a Rabin objective. Inspired by Condon [19], we will refer to SMGs with terminal reachability objectives and S2Gs with a (terminal) reachability objective for player 0 as *simple stochastic multiplayer games (SSMGs)* and *simple stochastic two-player zero-sum games (SS2Gs)*, respectively.

Drawing an SMG. When drawing an SMG as a graph, we will use the following conventions: The initial vertex is marked by a dangling incoming edge. Vertices that are controlled by a player are depicted as circles, where the player who controls the vertex is given by the label next to it. Stochastic vertices are depicted as diamonds, where the transition probabilities are given by the labels on its outgoing edges (the default being equal probabilities on all outgoing transitions). Finally, terminal vertices are generally represented by their associated payoff vector. In fact, we allow arbitrary vectors of rational probabilities as payoffs. This does not increase the power of the model since such a payoff vector can easily be realised by an SSMG consisting of stochastic and terminal vertices only.

2.3. Strategies and strategy profiles.

2.3.1. Randomised and pure strategies. The notion of a strategy lies at the heart of game theory. Formally, a (*randomised*) *strategy of player i in an SMG \mathcal{G}* is a mapping $\sigma: V^*V_i \rightarrow \mathcal{D}(V)$ assigning to each sequence $xv \in V^*V_i$ of vertices ending in a vertex controlled by player i a discrete probability distribution over V such that $\sigma(xv)(w) > 0$ only if $(v, \perp, w) \in \Delta$. Instead of $\sigma(xv)(w)$, we usually write $\sigma(w \mid xv)$. We say that a play π of \mathcal{G} is *compatible*

with a strategy σ of player i if $\sigma(\pi(k+1) \mid \pi(0) \dots \pi(k)) > 0$ for all $k \in \mathbb{N}$ with $\pi(k) \in V_i$. Similarly, a history $v_0 \dots v_k$ is compatible with σ if $\sigma(v_{j+1} \mid v_0 \dots v_j) > 0$ for all $0 \leq k < n$.

A (randomised) strategy profile of \mathcal{G} is a tuple $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ where σ_i is a strategy of player i in \mathcal{G} . We say that a play or a history of \mathcal{G} is compatible with a strategy profile $\bar{\sigma}$ if it is compatible with each σ_i . Given a strategy profile $\bar{\sigma} = (\sigma_j)_{j \in \Pi}$ and a strategy τ of player i , we denote by $(\bar{\sigma}_{-i}, \tau)$ the strategy profile resulting from $\bar{\sigma}$ by replacing σ_i with τ .

A strategy σ of player i is called *pure* or *deterministic* if for each $xv \in V^*V_i$ there exists $w \in v\Delta$ with $\sigma(w \mid xv) = 1$; note that a pure strategy of player i can be identified with a function $\sigma: V^*V_i \rightarrow V$. A strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ is called *pure* (or *deterministic*) if each σ_i is pure.

2.3.2. The probability space induced by a strategy profile. Given a game \mathcal{G} and a strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ of \mathcal{G} , the *conditional probability* of $w \in V$ given $xv \in V^*V$ is the number $\sigma_i(w \mid xv)$ if $v \in V_i$ and the probability $\Delta(w \mid v)$ if v is a stochastic vertex; let us denote this probability by $\bar{\sigma}(w \mid xv)$. Given an initial vertex $v_0 \in V$, the probabilities $\bar{\sigma}(w \mid xv)$ give rise to a probability measure: the probability of a basic cylinder set $v_0 \dots v_k \cdot V^\omega$ equals the product $\prod_{j=0}^{k-1} \bar{\sigma}(v_{j+1} \mid v_0 \dots v_j)$; basic cylinder sets that start in a vertex different from v_0 have probability 0. This definition induces a probability measure on the algebra of cylinder sets, which—by Carathéodory’s extension theorem—can be extended to a probability measure on the Borel σ -algebra over V^ω ; we denote the extended measure by $\Pr_{v_0}^{\bar{\sigma}}$. Finally, by viewing the colouring function $\chi: V \rightarrow C$ as a continuous function $V^\omega \rightarrow C^\omega$, we obtain a probability measure on the Borel σ -algebra over C^ω ; we abuse notation and denote this measure also by $\Pr_{v_0}^{\bar{\sigma}}$.

For a strategy profile $\bar{\sigma}$, we are mainly interested in the probabilities $p_i := \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i)$ of winning. We call p_i the (expected) *payoff* of $\bar{\sigma}$ for player i (from v_0) and the vector $(p_i)_{i \in \Pi}$ the (expected) *payoff* of $\bar{\sigma}$ (from v_0). Finally, we say that a history xv of (\mathcal{G}, v_0) is *consistent* with $\bar{\sigma}$ if $\Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) > 0$, i.e. if the basic cylinder induced by this history has positive probability.

In order to apply known results about Markov chains, we can also view the stochastic process induced by a strategy profile $\bar{\sigma}$ as a countable Markov chain $\mathcal{G}^{\bar{\sigma}}$, defined as follows: The set of states of $\mathcal{G}^{\bar{\sigma}}$ is the set V^+ of all non-empty sequences of vertices in \mathcal{G} . The only transitions from a state xv , where $x \in V^*$, $v \in V$, are to states of the form xvw , where $w \in V$, and such a transition occurs with probability $p > 0$ if and only if either v is stochastic and $(v, p, w) \in \Delta$ or $v \in V_i$ and $\sigma_i(w \mid xv) = p$. Finally, the colouring χ of vertices is extended to a colouring of states by setting $\chi(xv) = \chi(v)$ for all $x \in V^*$ and $v \in V$. With this definition, we could equivalently define the payoff of $\bar{\sigma}$ for player i as the probability of the event $\chi^{-1}(\text{Win}_i)$ in $(\mathcal{G}^{\bar{\sigma}}, v_0)$.

For each player i , the Markov decision process $\mathcal{G}^{\bar{\sigma}-i}$ is defined just as $\mathcal{G}^{\bar{\sigma}}$, but states $xv \in V^*V_i$ are controlled by player i (the unique player in $\mathcal{G}^{\bar{\sigma}-i}$), and there is a transition from such a state to each state of the form xvw , where $w \in V$, with $(v, \perp, w) \in \Delta$; player i ’s objective is the same as in \mathcal{G} .

2.3.3. Strategies with memory. A *memory structure* for a game \mathcal{G} with vertices in V is a triple $\mathfrak{M} = (M, \delta, m_0)$ where M is a set of *memory states*, $\delta: M \times V \rightarrow M$ is the *update function*, and $m_0 \in M$ is the *initial memory*. A (randomised) *strategy with memory* \mathfrak{M} of player i is a function $\sigma: M \times V_i \rightarrow \mathcal{D}(V)$ such that $\sigma(m, v)(w) > 0$ only if $w \in vE$. The

strategy σ is a *pure strategy with memory* \mathfrak{M} if additionally the following property holds: for all $m \in M$ and $v \in V$ there exists $w \in V$ such that $\sigma(m, v)(w) = 1$. Hence, a pure strategy with memory \mathfrak{M} can be described by a function $\sigma: M \times V_i \rightarrow V$. Finally, a *(pure) strategy profile with memory* \mathfrak{M} is a tuple $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ such that each σ_i is a (pure) strategy with memory \mathfrak{M} of player i .

A (pure) strategy σ with memory \mathfrak{M} of player i defines a (pure) strategy of player i in the usual sense as follows: Let $\delta^*(x)$ be the memory state after $x \in V^*$, defined inductively by $\delta^*(\varepsilon) = m_0$ and $\delta^*(xv) = \delta(\delta^*(x), v)$ for $x \in V^*$ and $v \in V$. If $v \in V_i$, then the distribution (successor vertex) chosen by the strategy σ for the sequence xv is $\sigma(\delta^*(x), v)$. Vice versa, every strategy (profile) of \mathcal{G} can be viewed as a strategy (profile) with memory $\mathfrak{M} := (V^*, \cdot, \varepsilon)$.

A *finite-state strategy (profile)* is a strategy (profile) with memory \mathfrak{M} for a finite memory structure \mathfrak{M} . Note that a strategy profile is finite-state if and only if each of its strategies is finite-state. If $|M| = 1$, we call a strategy (profile) with memory \mathfrak{M} *stationary*. Moreover, we call a pure stationary strategy (profile) a *positional* strategy (profile). A stationary strategy of player i can be described by a function $\sigma: V_i \rightarrow \mathcal{D}(V)$, and a positional strategy even by a function $\sigma: V_i \rightarrow V$.

If $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ is a strategy profile with memory \mathfrak{M} , we coarsen the Markov chain $\mathcal{G}^{\bar{\sigma}}$ by taking $M \times V$ as its domain. The transition relation is defined as follows: there is a transition from (m, v) to (n, w) with probability $p > 0$ if and only if $\delta(m, v) = n$ and either v is a stochastic vertex of \mathcal{G} and $(v, p, w) \in \Delta$ or $v \in V_i$ and $\sigma_i(m, v)(w) = p$. Finally, a state (m, v) has the same colour as the vertex v in \mathcal{G} . Analogously, we coarsen the Markov decision process $\mathcal{G}^{\bar{\sigma}-i}$ by using $M \times V$ as its domain: vertices $(m, v) \in M \times V_i$ are controlled by player i , and there is a transition from such a vertex (m, v) to $(n, w) \in M \times V$ if and only if $n = \delta(m, v)$ and $(v, \perp, w) \in \Delta$. Note that the arenas of both $\mathcal{G}^{\bar{\sigma}}$ and $\mathcal{G}^{\bar{\sigma}-i}$ are finite if the memory \mathfrak{M} and the original arena of \mathcal{G} are finite.

2.3.4. Residual games and strategies. Given an SMG \mathcal{G} and a sequence $x \in V^*$ (which is usually a history), the *residual game* $\mathcal{G}[x]$ has the same arena as \mathcal{G} but different objectives: if $\text{Win}_i \subseteq C^\omega$ is the objective of player i in \mathcal{G} , then her objective in $\mathcal{G}[x]$ is $\chi(x)^{-1} \text{Win}_i := \{\alpha \in C^\omega : \chi(x) \cdot \alpha \in \text{Win}_i\}$. In particular, if all objectives in \mathcal{G} are prefix-independent, then $\mathcal{G}[x] = \mathcal{G}$.

If player i plays according to a strategy σ in \mathcal{G} , then the corresponding strategy in $\mathcal{G}[x]$ is the *residual strategy* $\sigma[x]$, defined by $\sigma[x](yv) = \sigma(xyv)$. If $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ is a strategy profile, then the *residual strategy profile* $\bar{\sigma}[x]$ is just the profile of the residual strategies $\sigma_i[x]$. The following lemma, taken from [64], shows how to compute probabilities with respect to a residual strategy profile.

Lemma 2.1. Let $\bar{\sigma}$ be a strategy profile of an SMG (\mathcal{G}, v_0) , and let $xv \in V^*V$. If $X \subseteq V^\omega$ is a Borel set, then $\text{Pr}_{v_0}^{\bar{\sigma}}(X \cap xv \cdot V^\omega) = \text{Pr}_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \cdot \text{Pr}_v^{\bar{\sigma}[x]}(x^{-1}X)$.

2.4. Subarenas and end components. Algorithms for stochastic games often employ a *divide-and-conquer* approach and compute a solution for a complex game from the solution of several smaller games. These smaller games are usually obtained from the original game by restricting to a *subarena*. Formally, given an SMG \mathcal{G} , a set $U \subseteq V$ is a subarena if

- $U \neq \emptyset$,

- $v\Delta \cap U \neq \emptyset$ for each $v \in U$, and
- $v\Delta \subseteq U$ for each stochastic vertex $v \in U$.

Clearly, if U is a subarena, then the restriction of \mathcal{G} to vertices in U is again an SMG, which we denote by $\mathcal{G} \upharpoonright U$. Formally,

$$\mathcal{G} \upharpoonright U := (\Pi, U, (V_i \cap U)_{i \in \Pi}, \Delta \cap (U \times ([0, 1] \cup \{\perp\}) \times U), \chi_U, (\text{Win}_i)_{i \in \Pi}),$$

where $\chi_U: U \rightarrow C: u \mapsto \chi(u)$ is the restriction of the colouring function to U .

Of particular interest are the strongly connected subarenas of a game because they can arise as the sets $\text{Inf}(\pi)$ of vertices visited infinitely often in a play; we call these sets *end components*. Formally, a set $U \subseteq V$ is an end component if U is a subarena and every vertex $w \in U$ is reachable from every other vertex $v \in U$, i.e. there exists a sequence $v = v_1, v_2, \dots, v_n = w$ such that $v_{i+1} \in v_i\Delta$ for each $0 < i < n$. An end component U is *maximal* in a set $S \subseteq V$ if there is no end component U' such that $U \subsetneq U' \subseteq S$. For any finite subset $S \subseteq V$, the set of all end components maximal in S can be computed in quadratic time [24].

The theory of end components has been developed by de Alfaro [24, 25] and Courcoubetis and Yannakakis [21, 22]. The central fact about end components in finite SMGs is that, under any strategy profile, the set of vertices visited infinitely often is almost surely an end component.

Lemma 2.2. Let \mathcal{G} be a finite SMG. Then $\Pr_v^{\bar{\sigma}}(\{\pi \in V^\omega : \text{Inf}(\pi) \text{ is an end component}\}) = 1$ for each strategy profile $\bar{\sigma}$ of \mathcal{G} and each $v \in V$.

Moreover, for any end component U , we can construct a stationary strategy profile, or alternatively a pure finite-state strategy profile, that, when started in U , guarantees almost surely to visit all and only the vertices in U infinitely often. In fact, the stationary profile that chooses for each vertex in U a successor in U uniformly at random fulfils this property.

Lemma 2.3. Let \mathcal{G} be a finite SMG and U one of its end components. There exists both a stationary and a pure finite-state strategy profile $\bar{\sigma}$ such that $\Pr_v^{\bar{\sigma}}(\{\pi \in V^\omega : \text{Inf}(\alpha) = U\}) = 1$ for every vertex $v \in U$.

Given an SMG \mathcal{G} with (objectives representable as) Muller objectives given by a family \mathcal{F}_i of accepting sets, we say that an end component U is *winning* for player i if $\chi(U) \in \mathcal{F}_i$; the *payoff* of U is the vector $\bar{z} \in \{0, 1\}^\Pi$, defined by $z_i = 1$ if and only if U is winning for player i .

2.5. Values, determinacy and optimal strategies. Given a strategy τ of player i in \mathcal{G} and a vertex $v \in V$, the *value* of τ from v is the number $\text{val}^\tau(v) := \inf_{\bar{\sigma}} \Pr_v^{\bar{\sigma}-i, \tau}(\text{Win}_i)$, where $\bar{\sigma}$ ranges over all strategy profiles of \mathcal{G} . Moreover, the *value* of \mathcal{G} for player i from v is the supremum of these values: $\text{val}_i^{\mathcal{G}}(v) := \sup_{\tau} \text{val}^\tau(v)$, where τ ranges over all strategies of player i in \mathcal{G} . Intuitively, $\text{val}_i^{\mathcal{G}}(v)$ is the maximal payoff that player i can *ensure* when the game starts from v .

Given an initial vertex $v_0 \in V$, a strategy τ of player i in \mathcal{G} is called (*almost-surely*) *winning* if $\text{val}^\tau(v_0) = 1$. More generally, τ is called *optimal* if $\text{val}^\tau(v_0) = \text{val}_i^{\mathcal{G}}(v_0)$. For $\varepsilon > 0$, it is called ε -*optimal* if $\text{val}^\tau(v_0) \geq \text{val}_i^{\mathcal{G}}(v_0) - \varepsilon$. A *globally* (ε -) *optimal* strategy is a strategy that is (ε -)optimal for every possible initial vertex $v_0 \in V$. Note that optimal strategies need not exist since the supremum in the definition of $\text{val}_i^{\mathcal{G}}$ is not necessarily attained;

in this case, only ε -optimal strategies do exist. Also note that there exists a globally (ε -) optimal strategy whenever there exists an (ε -)optimal strategy for every possible initial vertex. Finally, we say that a strategy τ of player i in (\mathcal{G}, v_0) is *strongly optimal* if the residual strategy $\tau[x]$ is optimal in the residual game $(\mathcal{G}[x], v)$ for every history xv of (\mathcal{G}, v_0) that is compatible with τ . Intuitively, a strategy is strongly optimal if it is also optimal when the other players do not play optimally. Note that, for games with prefix-independent objectives, any globally optimal positional strategy profile is also strongly optimal.

Determining values and finding optimal strategies in SMGs actually reduces to performing the same tasks in S2Gs. Formally, given an SMG \mathcal{G} , define for each player i the *coalition game* \mathcal{G}_i to be the same game as \mathcal{G} but with only two players: player i acting as player 0 and the coalition $\Pi \setminus \{i\}$ acting as player 1. The coalition controls all vertices that in \mathcal{G} are controlled by some player $j \neq i$, and its objective is the complement of player i 's objective in \mathcal{G} . Clearly, \mathcal{G}_i is an S2G, and $\text{val}^{\mathcal{G}_i}(v) = \text{val}_i^{\mathcal{G}}(v)$ for every vertex v . Moreover, any (strongly, ε -) optimal strategy for player i in (\mathcal{G}, v_0) is (strongly, ε -) optimal in (\mathcal{G}_i, v_0) , and vice versa. Hence, when we study values and optimal strategies, we can restrict to S2Gs.

A celebrated theorem due to Martin [48] and Maitra and Sudderth [47] states that S2Gs with Borel objectives are *determined*: $\text{val}_0^{\mathcal{G}} = 1 - \text{val}_1^{\mathcal{G}}$. The number $\text{val}^{\mathcal{G}}(v) := \text{val}_0^{\mathcal{G}}(v)$ is consequently called the *value* of \mathcal{G} from v . In fact, an inspection of the proof shows that for turn-based games both players not only have randomised ε -optimal strategies but pure ε -optimal strategies.

Theorem 2.4 ([48, 47]). Every S2G with Borel objectives is determined; for all $\varepsilon > 0$, both players have ε -optimal pure strategies.

For finite S2Gs with prefix-independent objectives, we can show a stronger result than Theorem 2.4: in these games, both players not only have ε -optimal pure strategies but optimal ones [36]. In fact, the proof reveals the existence of strongly optimal strategies (see also [59]).

Theorem 2.5 ([36]). In any finite S2G with prefix-independent objectives, both players have strongly optimal pure strategies.

For finite S2Gs with ω -regular objectives, more attractive strategies than arbitrary pure strategies suffice for optimality. In particular, in any finite Rabin-Streett S2G there exists a globally optimal positional strategy for player 0 [46, 17].

Theorem 2.6 ([46, 17]). In any finite Rabin-Streett S2G, player 0 has a globally optimal positional strategy.

A consequence of Theorem 2.6 is that the values of a finite Rabin-Streett S2G are rational of polynomial bit complexity in the size of the arena: Given a positional strategy profile $\bar{\sigma}$ of \mathcal{G} , the finite MDP $\mathcal{G}^{\bar{\sigma}-1}$ is not larger than the game \mathcal{G} . Moreover, if σ_0 is globally optimal, then for every vertex v the value of \mathcal{G} from v and the value of $\mathcal{G}^{\bar{\sigma}-1}$ from v sum up to 1. But the values of any Streett MDP form the optimal solution of a linear programme of polynomial size (see [24]) and are therefore rational of small bit complexity.

Of course, it also follows from Theorem 2.6 that finite parity S2Gs are *positionally determined*: both players have globally optimal positional strategies. This result was first proven for deterministic games (even over infinite arenas), independently by Emerson and Jutla [28] and Mostowski [51]. For SS2Gs, the existence of optimal positional strategies follows from a result of Bewley and Kohlberg [7]. Independently, McIver and Morgan [49], Chatterjee et al. [15] and Zielonka [64] extended these results to parity S2Gs.

Corollary 2.7. In any finite parity S2G, both players have globally optimal positional strategies.

Since every finite S2G with ω -regular objectives can be reduced to one with parity objectives, we can conclude from Corollary 2.7 that both players have residually optimal pure finite-state strategies in finite S2Gs with arbitrary ω -regular objectives.

Corollary 2.8. In any finite S2G with ω -regular objectives, both players have strongly optimal pure finite-state strategies.

2.6. Algorithmic problems. For the rest of this section, we only consider finite two-player zero-sum games. The main computational problems for these games are computing the value and optimal strategies for one or both players. Instead of computing the value exactly, we can ask whether the value is greater than some given rational probability p , a problem which we call the *quantitative decision problem*:

Given an S2G \mathcal{G} , a vertex v and a rational number $p \in [0, 1]$, decide whether $\text{val}^{\mathcal{G}}(v) \geq p$.

In many cases, it suffices to know whether the value is 1, i.e. whether player 0 has a strategy to win the game almost surely (asymptotically, at least). We call the resulting decision problem the *qualitative decision problem*.

Clearly, if we can solve the quantitative decision problem, we can approximate the values $\text{val}^{\mathcal{G}}(v)$ up to any desired precision by using binary search. In fact, for parity S2Gs it is well-known that it suffices to solve the decision problems, since the other problems (computing the values and optimal strategies) are polynomial-time equivalent to the quantitative decision problem.

For a Markov decision process whose objective can be represented as a Muller objective, we can compute the values by an analysis of its end components: For a given initial vertex v , the value of the MDP from v equals the maximal probability of reaching a winning end component from v ; this probability can be computed using linear programming.

Even though, the number of end components can be exponential, it is easy to see that the *union* of all winning end components can be computed in polynomial time for MDPs with Rabin or Muller objectives (given by a family of accepting sets). For MDPs with Streett objectives, Chatterjee et al. [17] gave a polynomial-time algorithm for computing this set. Hence, for MDPs with any of these objectives, the quantitative decision problem is solvable in polynomial time.

Theorem 2.9 ([24, 17]). The quantitative decision problem is in P for Streett, Rabin or Muller MDPs.

It follows from Theorems 2.6 and 2.9 that the quantitative decision problem for Rabin-Streett S2Gs is in NP: to decide whether $\text{val}^{\mathcal{G}}(v) \geq p$, it suffices to guess a positional strategy for player 0 and to check whether in the resulting Streett MDP the value from v is not smaller than p . By determinacy, this result implies that the quantitative decision problem is in coNP for Streett-Rabin S2Gs and in $\text{NP} \cap \text{coNP}$ for parity S2Gs.

Corollary 2.10. The quantitative decision problem is

- in NP for Rabin-Streett S2Gs,
- in coNP for Streett-Rabin S2Gs, and

TABLE 1. The complexity of deciding the value in S2Gs.

	Qualitative	Quantitative
SS2Gs	P-complete	$\text{NP} \cap \text{coNP}$
Parity[d]	P-complete	$\text{NP} \cap \text{coNP}$
Parity	$\text{UP} \cap \text{coUP}$	$\text{NP} \cap \text{coNP}$
Rabin-Streett	NP-complete	NP-complete
Streett-Rabin	coNP-complete	coNP-complete
Muller	PSPACE-complete	PSPACE-complete

– in $\text{NP} \cap \text{coNP}$ for parity S2Gs.

A corresponding NP-hardness result for deterministic Rabin-Streett S2Gs was established by Emerson and Jutla [29]. In particular, this hardness result also holds for the qualitative decision problem. Moreover, by determinacy, this result can be turned into a coNP-hardness result for (deterministic) Streett-Rabin S2Gs.

For S2Gs with Muller objectives, Chatterjee [13] showed that the quantitative decision problem falls into PSPACE; for deterministic games, a polynomial-space algorithm had been given earlier by McNaughton [50]. A matching lower bound for deterministic games with Muller objectives was provided by Hunter and Dawar [42].

Theorem 2.11 ([13, 42]). The quantitative and the qualitative decision problem are PSPACE-complete for Muller S2Gs.

Theorem 2.11 does not hold if the Muller objective is given by a family of subsets of *vertices*: Horn [41, 40] showed that the qualitative decision problem for *explicit* Muller S2Gs is in P, and that the quantitative problem is in $\text{NP} \cap \text{coNP}$.

Another class of S2Gs for which the qualitative decision problem is in P is, for each $d \in \mathbb{N}$, the class Parity[d] of all parity S2Gs whose priority function refers to at most d priorities [26]. In particular, the qualitative decision problem for SS2Gs as well as (co-) Büchi S2Gs is in P. For general parity S2Gs, however, the qualitative decision problem is only known to lie in $\text{UP} \cap \text{coUP}$ [44, 14].

Theorem 2.12 ([44, 14, 26]). The qualitative decision problem is in $\text{UP} \cap \text{coUP}$ for parity S2Gs. For each $d \in \mathbb{N}$, the qualitative decision problem is in P for parity S2Gs with at most d priorities.

Table 1 summarises the results about the complexity of the quantitative and the qualitative decision problem for S2Gs. P-hardness (via LOGSPACE-reductions) for all these problems follows from the fact that *and-or graph reachability* is P-complete [43].

The results summarised in Table 1 leave open the possibility that at least one of the following problems is decidable in polynomial time:

- (1) the qualitative decision problem for parity S2Gs,
- (2) the quantitative decision problem for SS2Gs,
- (3) the quantitative decision problem for parity S2Gs.

Note that, given that all of them are contained in both NP and coNP, it is unlikely that one of them is NP-hard or coNP-hard; such a result would imply that $\text{NP} = \text{coNP}$, and the polynomial hierarchy would collapse.

For the first problem, Chatterjee et al. [14] gave a polynomial-time reduction to the qualitative decision problem for *deterministic* two-player zero-sum parity games. Hence, solving the qualitative decision problem for parity S2Gs is not harder than deciding which of the two players has a winning strategy in a deterministic two-player zero-sum parity game. Whether the latter problem is decidable in polynomial time is a long-standing open problem. Several years after Emerson and Jutla [28] put the problem into $\text{NP} \cap \text{coNP}$, Jurdziński [44] improved this bound slightly to $\text{UP} \cap \text{coUP}$. Together with Paterson and Zwick [45], he also gave an algorithm that decides the winner in subexponential time; a *randomised* subexponential algorithm had been given earlier by Björklund et al. [9]. On the other hand, Friedmann [33] recently showed that the most promising candidate for a polynomial-time algorithm for the general case so far, the *discrete strategy improvement* algorithm due to Vöge and Jurdziński [61], requires exponential time in the worst case.

Regarding the second problem, only some progress towards a polynomial-time algorithm has been made since Condon [19] proved membership in $\text{NP} \cap \text{coNP}$; for instance, Björklund and Vorobyov [8] gave a randomised subexponential algorithm for solving SS2Gs, and Gimbert and Horn [35] showed that the quantitative decision problem for SS2Gs is fixed-parameter tractable with respect to the number of stochastic vertices as the parameter.

For the third problem, Andersson and Miltersen [2] recently established a polynomial-time Turing reduction to the second. Hence, there exists a polynomial-time algorithm for (2) if and only if there exists one for (3). In particular, a polynomial-time algorithm for (2) would also give a polynomial-time algorithm for (1). However, to the best of our knowledge, it is plausible that the qualitative decision problem for parity S2Gs is in P while the quantitative decision problem for SS2Gs is not.

3. EXISTENCE OF NASH EQUILIBRIA

To capture rational behaviour of selfish players, Nash [52] introduced the notion of—what is now called—a Nash equilibrium. Formally, given a strategy profile $\bar{\sigma}$ of a game (\mathcal{G}, v_0) , we call a strategy τ of player i in \mathcal{G} a *best response* to $\bar{\sigma}$ if τ maximises the expected payoff of player i : $\Pr_{v_0}^{\bar{\sigma}-i, \tau'}(\text{Win}_i) \leq \Pr_{v_0}^{\bar{\sigma}-i, \tau}(\text{Win}_i)$ for all strategies τ' of player i . A strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ is a *Nash equilibrium* if each σ_i is a best response to $\bar{\sigma}$.

In a Nash equilibrium, no player can improve her payoff by unilaterally switching to a different strategy. In fact, to have a Nash equilibrium, it suffices that no player can gain from switching to a pure strategy.

Proposition 3.1. A strategy profile $\bar{\sigma}$ of a game (\mathcal{G}, v_0) is a Nash equilibrium if and only if, for each player i and for each *pure* strategy τ of player i in \mathcal{G} , $\Pr_{v_0}^{\bar{\sigma}-i, \tau}(\text{Win}_i) \leq \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i)$.

Proof. Clearly, if $\bar{\sigma}$ is a Nash equilibrium, then $\Pr_{v_0}^{\bar{\sigma}-i, \tau}(\text{Win}_i) \leq \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i)$ for each pure strategy τ of player i in \mathcal{G} . Now, assume that $\bar{\sigma}$ is not a Nash equilibrium. Hence, $p := \sup_{\tau} \Pr_{v_0}^{\bar{\sigma}-i, \tau}(\text{Win}_i) = \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i) + \varepsilon$ for some player i and some $\varepsilon > 0$. Consider the Markov decision process $\mathcal{G}^{\bar{\sigma}-i}$. Clearly, the value of $\mathcal{G}^{\bar{\sigma}-i}$ from v_0 equals p . By Theorem 2.4, there exists an $\varepsilon/2$ -optimal pure strategy τ in $(\mathcal{G}^{\bar{\sigma}-i}, v_0)$. Since the arena of $\mathcal{G}^{\bar{\sigma}-i}$ is a forest, we can assume that τ is a positional strategy, which can be viewed as a pure strategy in \mathcal{G} . We have $\Pr_{v_0}^{\bar{\sigma}-i, \tau}(\text{Win}_i) \geq p - \varepsilon/2 > p - \varepsilon = \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i)$. \square

For two-player zero-sum games, a Nash equilibrium is just a pair of optimal strategies.

Proposition 3.2. Let (\mathcal{G}, v_0) be an S2G. A strategy profile (σ, τ) of (\mathcal{G}, v_0) is a Nash equilibrium if and only if both σ and τ are optimal. In particular, every Nash equilibrium of (\mathcal{G}, v_0) has payoff $(\text{val}^{\mathcal{G}}(v_0), 1 - \text{val}^{\mathcal{G}}(v_0))$.

Proof. (\Rightarrow) Assume that both σ and τ are optimal, but that (σ, τ) is not a Nash equilibrium. Hence, one of the players, say player 1, can improve her payoff by playing some strategy τ' . Hence, $\text{val}^{\mathcal{G}}(v_0) = \Pr_{v_0}^{\sigma, \tau}(\text{Win}_0) > \Pr_{v_0}^{\sigma, \tau'}(\text{Win}_0)$. However, since σ is optimal, $\text{val}^{\mathcal{G}}(v_0) \leq \Pr_{v_0}^{\sigma, \tau'}(\text{Win}_0)$, a contradiction. The reasoning in the case that player 0 can improve is analogous.

(\Leftarrow) Let (σ, τ) be a Nash equilibrium of (\mathcal{G}, v_0) , and let us first assume that σ is not optimal, i.e. $\text{val}^{\sigma}(v_0) < \text{val}^{\mathcal{G}}(v_0)$. By the definition of $\text{val}^{\mathcal{G}}$, there exists another strategy σ' of player 0 such that $\text{val}^{\sigma}(v_0) < \text{val}^{\sigma'}(v_0) \leq \text{val}^{\mathcal{G}}(v_0)$. We have

$$\Pr_{v_0}^{\sigma, \tau}(\text{Win}_0) \leq \text{val}^{\sigma}(v_0) < \text{val}^{\sigma'}(v_0) = \inf_{\tau'} \Pr_{v_0}^{\sigma', \tau'}(\text{Win}_0) \leq \Pr_{v_0}^{\sigma', \tau}(\text{Win}_0),$$

where the first inequality follows from the fact that (σ, τ) is a Nash equilibrium. Thus, player 0 can improve her payoff by playing σ' instead of σ , a contradiction to (σ, τ) being a Nash equilibrium. The argumentation in the case that τ is not optimal is analogous. \square

In general, a Nash equilibrium can give a player a higher payoff than her value. However, the payoff a player receives in a Nash equilibrium can never be lower than her value, and this is true for every history that is consistent with the equilibrium. Formally, we say that a strategy profile $\bar{\sigma}$ of a game (\mathcal{G}, v_0) is *favourable* if $\Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) \geq \text{val}_i^{\mathcal{G}[x]}(v)$ for each player i and every history xv that is consistent with $\bar{\sigma}$.

Lemma 3.3. Let (\mathcal{G}, v_0) be an SMG. Every Nash equilibrium of (\mathcal{G}, v_0) is favourable.

Proof. Assume there exists a history xv of (\mathcal{G}, v_0) that is consistent with $\bar{\sigma}$, but $p := \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) < \text{val}_i^{\mathcal{G}[x]}(v)$. By the definition of $\text{val}_i^{\mathcal{G}[x]}$, there exists a strategy τ of player i in $\mathcal{G}[x]$ such that $\text{val}^{\tau}(v) > p$. We define a new strategy σ' for player i in \mathcal{G} as follows: σ' is defined as σ_i for histories that do not begin with xv . For histories of the form xvy , however, we set $\sigma'(xvy) = \tau(vy)$. Clearly, $\Pr_{v_0}^{\bar{\sigma}-i, \sigma'}(xv \cdot V^\omega) = \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega)$. Moreover, it is easy to see that $\Pr_{v_0}^{\bar{\sigma}-i, \sigma'}(X \setminus xv \cdot V^\omega) = \Pr_{v_0}^{\bar{\sigma}}(X \setminus xv \cdot V^\omega)$ for every Borel set $X \subseteq V^\omega$. Using Lemma 2.1, we can conclude that

$$\begin{aligned} & \Pr_{v_0}^{\bar{\sigma}-i, \sigma'}(\text{Win}_i) \\ &= \Pr_{v_0}^{\bar{\sigma}-i, \sigma'}(\text{Win}_i \setminus xv \cdot V^\omega) + \Pr_{v_0}^{\bar{\sigma}-i, \sigma'}(\text{Win}_i \cap xv \cdot V^\omega) \\ &= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus xv \cdot V^\omega) + \Pr_v^{\bar{\sigma}[x]-i, \sigma'[x]}(x^{-1}\text{Win}_i) \cdot \Pr_{v_0}^{\bar{\sigma}-i, \sigma'}(xv \cdot V^\omega) \\ &= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus xv \cdot V^\omega) + \Pr_v^{\bar{\sigma}[x]-i, \tau}(x^{-1}\text{Win}_i) \cdot \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \\ &\geq \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus xv \cdot V^\omega) + \text{val}^{\tau}(v) \cdot \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \\ &> \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus xv \cdot V^\omega) + p \cdot \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \\ &= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus xv \cdot V^\omega) + \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) \cdot \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \\ &= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus xv \cdot V^\omega) + \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \cap xv \cdot V^\omega) \\ &= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i). \end{aligned}$$

Hence, player i can improve her payoff by switching to σ' , a contradiction to $\bar{\sigma}$ being a Nash equilibrium. \square

It follows from Theorem 2.5 and Proposition 3.2 that every finite two-player zero-sum stochastic game with prefix-independent objectives has a Nash equilibrium in pure strategies. Is this still true if the two-player zero-sum assumption is relaxed?

By Lemma 3.3, a pure strategy profile can only be a Nash equilibrium if it is favourable. The next lemma shows that, conversely, we can turn every favourable pure strategy profile into a Nash equilibrium. The proof uses so-called *threat* strategies (or *trigger* strategies), which are added on top of the given strategy profile: each player threatens to change her behaviour when one of the other players deviates from the prescribed strategy profile. Before being applied to stochastic games, this concept proved fruitful in the related area of *repeated games* (see [54, Chapter 8] and [3]).

Lemma 3.4. Let (\mathcal{G}, v_0) be a finite SMG with prefix-independent objectives. If $\bar{\sigma}$ is a favourable pure strategy profile of (\mathcal{G}, v_0) , then (\mathcal{G}, v_0) has a pure Nash equilibrium $\bar{\sigma}^*$ with $\Pr_{v_0}^{\bar{\sigma}} = \Pr_{v_0}^{\bar{\sigma}^*}$.

Proof. By Theorem 2.5, for each player i we can fix a globally optimal pure strategy τ_i of the coalition $\Pi \setminus \{i\}$ in the coalition game \mathcal{G}_i ; denote by $\tau_{j,i}$ the corresponding pure strategy of player $j \neq i$ in \mathcal{G} . To simplify notation, we also define $\tau_{i,i}$ to be an arbitrary pure strategy of player i in \mathcal{G} . Player i 's equilibrium strategy σ_i^* is defined as follows: For histories xv that are compatible with $\bar{\sigma}$, we set $\sigma_i^*(xv) = \sigma_i(xv)$. If xv is not compatible with $\bar{\sigma}$, then decompose x into $x = x_1 \cdot x_2$, where x_1 is the longest prefix of x that is compatible with $\bar{\sigma}$, and let j be the player who has deviated, i.e. x_1 ends in V_j ; we set $\sigma_i^*(xv) = \tau_{i,j}(x_2v)$. Intuitively, σ_i^* behaves like σ_i as long as no other player j deviates from playing σ_j , in which case σ_i^* starts to behave like $\tau_{i,j}$.

Note that $\Pr_{v_0}^{\bar{\sigma}^*} = \Pr_{v_0}^{\bar{\sigma}}$. We claim that $\bar{\sigma}^*$ is additionally a Nash equilibrium of (\mathcal{G}, v_0) . Let $i \in \Pi$, and let ρ be a pure strategy of player i in \mathcal{G} ; by Proposition 3.1, it suffices to show that $\Pr_{v_0}^{\bar{\sigma}^* - i, \rho}(\text{Win}_i) \leq \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i)$.

Let us call a history $xv \in V^*V_i$ a *deviation history* if xv is compatible with both $\bar{\sigma}$ and $(\bar{\sigma}_{-i}, \rho)$, but $\sigma_i(xv) \neq \rho(xv)$; we denote the set of all deviation histories consistent with $\bar{\sigma}$ by D . Clearly, $\Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) = \Pr_{v_0}^{\bar{\sigma}^*}(xv \cdot V^\omega) = \Pr_{v_0}^{\bar{\sigma}^* - i, \rho}(xv \cdot V^\omega)$ for all $xv \in D$.

Claim. $\Pr_{v_0}^{\bar{\sigma}^* - i, \rho}(X \setminus D \cdot V^\omega) = \Pr_{v_0}^{\bar{\sigma}}(X \setminus D \cdot V^\omega)$ for every Borel set $X \subseteq V^\omega$.

Proof. This claim can be proved by an induction over the structure of Borel set.

Claim. $\Pr_{v_0}^{\bar{\sigma}^* - i, \rho}(\text{Win}_i \mid xv \cdot V^\omega) \leq \text{val}_i^{\mathcal{G}}(v)$ for every $xv \in D$.

Proof. By the definition of the strategies $\tau_{j,i}$, we have that $\Pr_v^{(\tau_{j,i})_{j \neq i, \rho}}(\text{Win}_i) \leq \text{val}_i^{\mathcal{G}}(v)$ for every vertex $v \in V$ and every strategy ρ of player i . Moreover, if xv is a deviation history, then for each player $j \neq i$ the residual strategy $\sigma_j^*[xv]$ is equal to $\tau_{j,i}$ on histories that start

in $w := \rho(xv)$. Hence, by Lemma 2.1 and since Win_i is prefix-independent,

$$\begin{aligned}
& \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i \mid xv \cdot V^\omega) \\
&= \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i \mid xvw \cdot V^\omega) \\
&= \Pr_w^{\bar{\sigma}^*_{[xv], \rho[xv]}}(\text{Win}_i) \\
&\leq \text{val}_i^{\mathcal{G}}(w) \\
&\leq \text{val}_i^{\mathcal{G}}(v).
\end{aligned}$$

Using the previous two claims, we prove that $\Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i) \leq \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i)$ as follows:

$$\begin{aligned}
& \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i) \\
&= \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i \setminus D \cdot V^\omega) + \sum_{xv \in D} \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i \cap xv \cdot V^\omega) \\
&= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus D \cdot V^\omega) + \sum_{xv \in D} \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i \cap xv \cdot V^\omega) \\
&= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus D \cdot V^\omega) + \sum_{xv \in D} \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i \mid xv \cdot V^\omega) \cdot \Pr_{v_0}^{\bar{\sigma}^*}(xv \cdot V^\omega) \\
&= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus D \cdot V^\omega) + \sum_{xv \in D} \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i \mid xv \cdot V^\omega) \cdot \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \\
&\leq \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus D \cdot V^\omega) + \sum_{xv \in D} \text{val}_i^{\mathcal{G}}(v) \cdot \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \\
&\leq \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus D \cdot V^\omega) + \sum_{xv \in D} \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) \cdot \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) \\
&= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \setminus D \cdot V^\omega) + \sum_{xv \in D} \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \cap xv \cdot V^\omega) \\
&= \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i) \\
&= \Pr_{v_0}^{\bar{\sigma}^*}(\text{Win}_i),
\end{aligned}$$

where the second inequality follows from the assumption that $\bar{\sigma}$ is favourable. \square

A variant of Lemma 3.4 handles games with prefix-independent ω -regular objectives and finite-state strategies.

Lemma 3.5. Let (\mathcal{G}, v_0) be a finite SMG with prefix-independent ω -regular objectives. If $\bar{\sigma}$ is a favourable pure finite-state strategy profile of (\mathcal{G}, v_0) , then (\mathcal{G}, v_0) has a pure finite-state Nash equilibrium $\bar{\sigma}^*$ with $\Pr_{v_0}^{\bar{\sigma}} = \Pr_{v_0}^{\bar{\sigma}^*}$.

Proof. The proof is analogous to the proof of Lemma 3.4. Since, by Corollary 2.8, there exist optimal pure finite-state strategies in every finite SMG with ω -regular objectives, the strategies $\tau_{j,i}$ defined there can be assumed to be pure finite-state strategies. Consequently, the equilibrium profile $\bar{\sigma}^*$ can be implemented using finite-state strategies as well. \square




FIGURE 2. A two-player game with a pair of optimal strategies that cannot be extended to a Nash equilibrium.

Using Lemma 3.4 and Theorem 2.5, we can easily prove the existence of pure Nash equilibria in finite SMGs with prefix-independent objectives.

Theorem 3.6. There exists a pure Nash equilibrium in any finite SMG with prefix-independent objectives.

Proof. Let \mathcal{G} be a finite SMG with prefix-independent objectives and initial vertex v_0 . By Theorem 2.5 and the correspondence between \mathcal{G} and the coalition game \mathcal{G}_i , each player i has a strongly optimal strategy σ_i in \mathcal{G} . Let $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$. For every history xv that is consistent with $\bar{\sigma}$ and each player i , we have $\Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) = \Pr_v^{\bar{\sigma}[x]}(\text{Win}_i) \geq \text{val}_i^{\mathcal{G}}(v)$. Hence, $\bar{\sigma}$ is favourable, and Lemma 3.4 implies the existence of a pure Nash equilibrium. \square

For finite SMGs with ω -regular objectives, we can even show the existence of a pure finite-state equilibrium.

Theorem 3.7. There exists a pure finite-state Nash equilibrium in any finite SMG with ω -regular objectives.

Proof. Since any SMG with ω -regular objectives can be reduced to one with parity objectives using finite memory, it suffices to consider parity SMGs. For these games, the claim follows from Corollary 2.7 and Lemma 3.5 using the same argumentation as in the proof of Theorem 3.6. \square

Theorem 3.7 and a variant of Theorem 3.6 appeared originally in [16]. However, their proof contains an inaccuracy: Essentially, they claim that any profile of optimal strategies can be extended to a Nash equilibrium with the same payoff (by adding threat strategies on top). This is, in general, not true, as the following example demonstrates.

Example 3.8. Consider the deterministic two-player game (\mathcal{G}, v_0) depicted in Fig. 2 and played by players 0 and 1 (with payoffs given in this order). Clearly, the value $\text{val}_0^{\mathcal{G}}(v_0)$ for player 0 from v_0 equals 1, and player 0's optimal strategy σ is to play from v_0 to v_1 . For player 1, the value from v_0 is 0, and both of her positional strategies are optimal. In particular, her strategy τ of playing from v_1 to the terminal vertex with payoff $(1, 0)$ is optimal (albeit not globally optimal). The payoff of the strategy profile (σ, τ) is $(1, 0)$. However, there is no Nash equilibrium of (\mathcal{G}, v_0) with payoff $(1, 0)$: In any Nash equilibrium of (\mathcal{G}, v_0) , player 0 will move from v_0 to v_1 with probability 1. Player 1's best response is to play from v_1 to the terminal vertex with payoff $(1, 1)$ with probability 1. Hence, every Nash equilibrium of this game has payoff $(1, 1)$.

4. COMPLEXITY OF NASH EQUILIBRIA

For the rest of this paper, we consider only *finite* SMGs. Previous research on algorithms for finding Nash equilibria in such games has focused on computing *some* Nash equilibrium [16]. However, a game may have several Nash equilibria with different payoffs, and one might not be interested in *any* Nash equilibrium but in one whose payoff fulfils certain requirements. For example, one might look for a Nash equilibrium where certain players win almost surely while certain others lose almost surely. This idea leads us to the following decision problem, which we call NE:

Given an SMG (\mathcal{G}, v_0) and thresholds $\bar{x}, \bar{y} \in [0, 1]^\Pi$, decide whether there exists a Nash equilibrium of (\mathcal{G}, v_0) with payoff $\geq \bar{x}$ and $\leq \bar{y}$.

To obtain meaningful results, we assume that all transition probabilities in \mathcal{G} as well as the thresholds \bar{x} and \bar{y} are rational numbers (with numerator and denominator given in binary) and that all objectives are ω -regular. A *qualitative* variant of the problem, which omits the thresholds, just asks about a Nash equilibrium where some distinguished player, say player 0, wins with probability 1:

Given an SMG (\mathcal{G}, v_0) , decide whether there exists a Nash equilibrium of (\mathcal{G}, v_0) where player 0 wins almost surely.

Clearly, every instance of the qualitative variant can easily be turned into an instance of NE (by adding the thresholds $\bar{x} = (1, 0, \dots, 0)$ and $\bar{y} = (1, \dots, 1)$). Hence, NE is, a priori, more general than its qualitative variant.

Note that we have so far not put any restriction on the type of strategies that realise the equilibrium. It is natural to restrict the search space to profiles of pure, finite-state, pure finite-state, stationary or positional strategies. We denote the corresponding decision problems by PureNE, FinNE, PureFinNE, StatNE and PosNE, respectively. In the rest of this paper, we are going to prove upper and lower bounds on the complexity of these problems, where all lower bounds even hold for the qualitative variants of these problems.

Our first observation is that neither stationary nor pure strategies are sufficient to implement any Nash equilibrium, even in SSMGs and even if we are only interested in whether a player wins or loses almost surely in the equilibrium. Together with another result from this section (Proposition 4.13), this demonstrates that the problems NE, PureNE, FinNE, PureFinNE, StatNE and PosNE are distinct problems, which have to be analysed separately. This is in sharp contrast to the situation for SS2Gs where all these problems coincide because SS2Gs admit globally optimal positional strategies.

Proposition 4.1. There exists an SSMG with a stationary Nash equilibrium where player 0 wins almost surely, but with no pure Nash equilibrium where player 0 wins with positive probability.

Proof. Consider the SSMG depicted in Fig. 3, played by three players 0, 1 and 2 (with payoffs given in this order). Clearly, the stationary strategy profile where from vertex v_2 player 0 selects both outgoing transitions with probability 1/2 each, player 1 plays from v_0 to v_1 and player 2 plays from v_1 to v_2 is a Nash equilibrium where player 0 wins almost surely. However, for any pure strategy profile where player 0 wins with positive probability, either player 1 or player 2 receives payoff 0 and could improve her payoff by switching her strategy at v_0 or v_1 , respectively. \square




FIGURE 3. An SSMG with no pure Nash equilibrium where player 0 wins with positive probability.




FIGURE 4. An SSMG with no stationary Nash equilibrium where player 0 wins with positive probability.

Proposition 4.2. There exists an SSMG with a pure finite-state Nash equilibrium where player 0 wins almost surely, but with no stationary Nash equilibrium where player 0 wins with positive probability.

Proof. Consider the (deterministic) SSMG \mathcal{G} depicted in Fig. 4, also played by three players 0, 1 and 2. Clearly, the pure finite-state strategy profile that leads to the terminal vertex with payoff $(1, 0, 0)$ and where at v_2 player 0 plays “right” if player 1 has played to v_0 and “left” if player 2 has played to v_0 is a Nash equilibrium of (\mathcal{G}, v_0) . Now consider any stationary equilibrium of (\mathcal{G}, v_0) where player 0 wins with positive probability. If at v_2 the stationary strategy of player 0 prescribes to play “right” with positive probability, then player 2 can improve her payoff by playing to v_2 with probability 1, and otherwise player 1 can improve her payoff by playing to v_2 with probability 1, a contradiction. \square

4.1. Positional equilibria. In this subsection, we analyse the complexity of the (presumably) simplest of the decision problems introduced so far, namely PosNE. Not surprisingly, this problem is decidable; in fact, it is NP-complete for all types of objectives we consider in this paper. Let us start by proving membership to NP. Since terminal reachability, (co-)Büchi and parity objectives can easily be translated to Rabin or Streett objectives, it suffices to consider Streett-Rabin and Muller SMGs.

Theorem 4.3. PosNE is in NP for Streett-Rabin SMGs and Muller SMGs.

Proof. To decide PosNE, on input $\mathcal{G}, v_0, \bar{x}, \bar{y}$, we can guess a positional strategy profile $\bar{\sigma}$, i.e. a mapping $\bigcup_{i \in \Pi} V_i \rightarrow V$; then, we verify whether $\bar{\sigma}$ is a Nash equilibrium with the desired payoff. To do this, we first compute the payoff z_i of $\bar{\sigma}$ for each player i by computing the probability of the event Win_i in the (finite) Markov chain $(\mathcal{G}^{\bar{\sigma}}, v_0)$. Once each z_i is computed, we can easily check whether $x_i \leq z_i \leq y_i$. To verify that $\bar{\sigma}$ is a Nash equilibrium, we additionally compute, for each player i , the value r_i of the (finite) MDP $\mathcal{G}^{\bar{\sigma}-i}$ from v_0 . Clearly, $\bar{\sigma}$ is a Nash equilibrium if and only if $r_i \leq z_i$ for each player i . Since we can compute the value of an MDP (or a Markov chain) with a Streett, Rabin or Muller objective in polynomial time (Theorem 2.9), all these checks can be carried out in polynomial time. \square

To establish NP-completeness, we still need to show NP-hardness. In fact, the reduction we are going to present does not only work for PosNE, but also for StatNE, where we allow arbitrary stationary equilibria.

Theorem 4.4. PosNE and StatNE are NP-hard, even for SSMGs with only two players (three players for the qualitative variants).

Proof. The proof is by reduction from SAT. Given a Boolean formula $\varphi = C_1 \wedge \dots \wedge C_m$ in conjunctive normal form over propositional variables X_1, \dots, X_n , where without loss of generality $m \geq 1$ and each clause is non-empty, we show how to construct a two-player SSMG (\mathcal{G}, v_0) such that the following statements are equivalent:

- (1) φ is satisfiable.
- (2) (\mathcal{G}, v_0) has a positional Nash equilibrium with payoff $(1, \frac{1}{2})$.
- (3) (\mathcal{G}, v_0) has a stationary Nash equilibrium with payoff $(1, \frac{1}{2})$.

Provided that the game can be constructed in polynomial time, these equivalences establish both reductions. The game \mathcal{G} is depicted in Fig. 5. The game proceeds from the initial vertex v_0 to X_i or $\neg X_i$ with probability $1/2^{i+1}$ each, and to vertex φ with probability $1/2^{n+1}$; with the remaining probability of $1/2^{n+1}$ the game proceeds to a terminal vertex with payoff $(1, 0)$. From φ , the game proceeds to each vertex C_j with probability $1/(m+1)$; with the remaining probability of $1/(m+1)$, the game proceeds to a terminal vertex with payoff $(1, 1)$. From vertex C_j (controlled by player 1), there is a transition to a literal L , i.e. $L = X_i$ or $L = \neg X_i$, if and only if L occurs inside the clause C_j . Obviously, the game \mathcal{G} can be constructed from φ in polynomial time. We conclude the proof by showing that (1)–(3) are equivalent.

(1 \Rightarrow 2) Assume that $\alpha: \{X_1, \dots, X_n\} \rightarrow \{\text{true}, \text{false}\}$ is a satisfying assignment of φ . Consider the positional strategy profile where player 0 moves from a literal L to the neighbouring \top -labelled vertex if and only if L is mapped to true by α , and player 1 moves from vertex C_j to a literal L that is contained in C_j and mapped to true by α (which is possible since α is a satisfying assignment); at \top -labelled vertices, player 1 never plays to a terminal vertex. Obviously, player 0 wins almost surely in this strategy profile. In order to compute player 1's payoff, note that for each variable X player 1 either receives payoff 1 from X and payoff 0 from $\neg X$, or she receives payoff 1 from $\neg X$ and payoff 0 from X (because player 0 plays according to a well-defined assignment). Moreover, player 1 wins almost surely from φ since that assignment satisfies φ . Hence, player 1's payoff equals

$$\frac{1}{2^{n+1}} + \sum_{i=1}^n \frac{1}{2^{i+1}} = \frac{1}{2^{n+1}} + \frac{1}{2} \left(\sum_{i=1}^n \frac{1}{2^i} \right) = \frac{1}{2^{n+1}} + \frac{1}{2} \left(1 - \frac{1}{2^n} \right) = \frac{1}{2},$$

Obviously, changing her strategy cannot give player 1 a better payoff. Therefore, we have identified a Nash equilibrium.

(2 \Rightarrow 3) Trivial.

(3 \Rightarrow 1) Let $\bar{\sigma} = (\sigma_0, \sigma_1)$ be a stationary Nash equilibrium of (\mathcal{G}, v_0) with payoff $(1, \frac{1}{2})$. Our first aim is to show that σ_0 is actually a positional strategy. Consider any literal L such that $\sigma_0(L)$ assigns probability $q > 0$ to the neighbouring \top -labelled vertex. Since player 0 wins almost surely, we know that player 1 never plays to a terminal vertex with payoff $(0, 1)$. Hence, the expected payoff for player 1 from L equals q . However, by playing to a terminal vertex with payoff $(0, 1)$, player 1 can get payoff $2q/(1+q)$ from L . Since $\bar{\sigma}$ is a Nash equilibrium, we have $2q/(1+q) \leq q$, which implies that $q = 1$.

Now we define a *pseudo assignment* $\alpha: \{X_1, \neg X_1, \dots, X_n, \neg X_n\} \rightarrow \{\text{true}, \text{false}\}$ by setting $\alpha(L) = \text{true}$ if and only if σ_0 prescribes to go from vertex L to the neighbouring \top -labelled vertex. Our next aim is to show that α is actually an assignment: $\alpha(X_i) = \text{true}$




FIGURE 5. Reducing SAT to PosNE and StatNE.

if and only if $\alpha(\neg X_i) = \text{false}$. To see this, note that we can compute player 1's expected payoff from v_0 as follows:

$$\frac{1}{2} = \frac{p}{2^{n+1}} + \sum_{i=1}^n \frac{a_i}{2^{i+1}}, \quad a_i = \begin{cases} 0 & \text{if } \alpha(X_i) = \alpha(\neg X_i) = \text{false}, \\ 1 & \text{if } \alpha(X_i) \neq \alpha(\neg X_i), \\ 2 & \text{if } \alpha(X_i) = \alpha(\neg X_i) = \text{true}, \end{cases}$$

where p is the expected payoff for player 1 from vertex φ . By the construction of \mathcal{G} , we have $p > 0$, and the equality only holds if $p = 1$ and $a_i = 1$ for all $i = 1, \dots, n$, which proves that α is an assignment.

Finally, we claim that α satisfies φ . If this were not the case, there would exist a clause C such that player 1's expected payoff from vertex C equals 0, and therefore $p < 1$. This is a contradiction to $p = 1$, as we have shown above.

To show that the qualitative variants of PosNE and StatNE are also NP-hard, it suffices to modify the game \mathcal{G} as follows: First, we add one new player, player 2, who wins at precisely those terminal vertices where player 1 loses. Second, we add two new vertices v_1 and v_2 . At v_1 , player 1 has the choice to leave the game; if she decides to stay inside the game, the play proceeds to v_2 , where player 2 has the choice to leave the game; if she also decides to stay inside the game, the play proceeds to vertex v_0 from where the game continues normally; if player 1 or player 2 decides to leave the game, then each of them receives payoff $\frac{1}{2}$, but player 0 receives payoff 0. Let us denote the modified game by \mathcal{G}' . It is straightforward to see that the following statements are equivalent:

- (1) (\mathcal{G}', v_1) has a stationary Nash equilibrium where player 0 wins almost surely.
- (2) (\mathcal{G}, v_0) has a stationary Nash equilibrium with payoff $(1, \frac{1}{2})$.
- (3) φ is satisfiable.
- (4) (\mathcal{G}, v_0) has a positional Nash equilibrium with payoff $(1, \frac{1}{2})$.
- (5) (\mathcal{G}', v_1) has a positional Nash equilibrium where player 0 wins almost surely. \square

4.2. Stationary equilibria. To prove the decidability of StatNE, we appeal to results established for the *existential theory of the reals*, the set of all existential first-order sentences (over the appropriate signature) that hold in the ordered field $\mathfrak{R} := (\mathbb{R}, +, \cdot, 0, 1, \leq)$. The best known upper bound for the complexity of the associated decision problem is PSPACE [12], which leads to the following theorem.

Theorem 4.5. StatNE is in PSPACE for Streett-Rabin SMGs and Muller SMGs.

Proof. Since PSPACE = NPSPACE, it suffices to provide a nondeterministic algorithm with polynomial space requirements for deciding StatNE. On input $\mathcal{G}, v_0, \bar{x}, \bar{y}$, where without loss of generality \mathcal{G} is an SMG with Muller objectives given by $\mathcal{F}_i \subseteq \mathcal{P}(C)$, the algorithm starts by guessing the *support* $S \subseteq V \times V$ of a stationary strategy profile $\bar{\sigma}$ of \mathcal{G} , i.e. $S = \{(v, w) \in V \times V : \bar{\sigma}(w | v) > 0\}$. From the set S alone, by standard graph algorithms, one can compute for each player i the following sets in polynomial time (see [4, Chapter 10]):

- (1) the union F_i of all bottom SCCs U of the Markov chain $\mathcal{G}^{\bar{\sigma}}$ with $\chi(U) \in \mathcal{F}_i$,
- (2) the set R_i of vertices v such that $\Pr_v^{\bar{\sigma}}(\text{Reach}(F_i)) > 0$,
- (3) the union T_i of all end components U of the MDP $\mathcal{G}^{\bar{\sigma}-i}$ with $\chi(U) \in \mathcal{F}_i$.

After computing all these sets, the algorithm evaluates an existential first-order sentence ψ , which can be computed in polynomial time from $\mathcal{G}, v_0, \bar{x}, \bar{y}, S, (R_i)_{i \in \Pi}, (F_i)_{i \in \Pi}$ and $(T_i)_{i \in \Pi}$, over \mathfrak{R} and returns the answer to this query.

How does ψ look like? Let $\bar{\alpha} = (\alpha_{vw})_{v,w \in V}$, $\bar{r} = (r_v^i)_{i \in \Pi, v \in V}$ and $\bar{z} = (z_v^i)_{i \in \Pi, v \in V}$ be three sets of variables, and let $V_* = \bigcup_{i \in \Pi} V_i$. The formula

$$\begin{aligned} \varphi(\bar{\alpha}) := & \bigwedge_{v \in V_*} \left(\bigwedge_{w \in v\Delta} \alpha_{vw} \geq 0 \wedge \bigwedge_{w \in V \setminus v\Delta} \alpha_{vw} = 0 \wedge \sum_{w \in v\Delta} \alpha_{vw} = 1 \right) \wedge \\ & \bigwedge_{\substack{v \in V \setminus V_* \\ w \in V}} \alpha_{vw} = \Delta(w | v) \wedge \bigwedge_{(v,w) \in S} \alpha_{vw} > 0 \wedge \bigwedge_{(v,w) \notin S} \alpha_{vw} = 0 \end{aligned}$$

states that the mapping $\bar{\sigma}: V \rightarrow [0, 1]^V$, defined by $\bar{\sigma}(v)(w) = \alpha_{vw}$, constitutes a valid stationary strategy profile of \mathcal{G} whose support is S . Provided that $\varphi(\bar{\alpha})$ holds in \mathfrak{R} , the formula

$$\eta_i(\bar{\alpha}, \bar{z}) := \bigwedge_{v \in F_i} z_v^i = 1 \wedge \bigwedge_{v \in V \setminus R_i} z_v^i = 0 \wedge \bigwedge_{v \in V \setminus F_i} z_v^i = \sum_{w \in v\Delta} \alpha_{vw} \cdot z_w^i$$

states that $z_v^i = \Pr_{\bar{\sigma}}^v(\text{Win}_i)$ for each $v \in V$, where $\bar{\sigma}$ is defined as above. This follows from a well-known result about Markov chains, namely that the vector of the aforementioned probabilities is the unique solution of the given system of equations (see [4, Chapter 10]). Finally, the formula

$$\vartheta_i(\bar{\alpha}, \bar{r}) := \bigwedge_{v \in V} r_v^i \geq 0 \wedge \bigwedge_{v \in T_i} r_v^i = 1 \wedge \bigwedge_{\substack{v \in V_i \\ w \in v\Delta}} r_v^i \geq r_w^i \wedge \bigwedge_{v \in V \setminus V_i} r_v^i = \sum_{w \in v\Delta} \alpha_{vw} \cdot r_w^i$$

states that $r_v^i \geq \sup_{\tau} \Pr_v^{\bar{\sigma}^{-i, \tau}}(\text{Win}_i)$ for all $v \in V$ (see [4, Chapter 10]).

The desired sentence ψ is the existential closure of the conjunction of φ and, for each player i , the formulae η_i and ϑ_i combined with formulae stating that player i cannot improve her payoff and that the expected payoff for player i lies in between the given thresholds:

$$\psi := \exists \bar{\alpha} \exists \bar{r} \exists \bar{z} \left(\varphi(\bar{\alpha}) \wedge \bigwedge_{i \in \Pi} (\eta_i(\bar{\alpha}, \bar{z}) \wedge \vartheta_i(\bar{\alpha}, \bar{r}) \wedge r_{v_0}^i \leq z_{v_0}^i \wedge x_i \leq z_{v_0}^i \leq y_i) \right).$$

Clearly, ψ holds in \mathfrak{R} if and only if (\mathcal{G}, v_0) has a stationary Nash equilibrium with payoff at least \bar{x} and at most \bar{y} whose support is S . Consequently, the algorithm is correct. \square

In Section 4.1, we showed that StatNE is NP-hard, leaving a considerable gap to our upper bound of PSPACE. Towards gaining a better understanding, we relate StatNE to the *square root sum problem* (SqrtSum) of deciding, given numbers $d_1, \dots, d_n, k \in \mathbb{N}$, whether $\sum_{i=1}^n \sqrt{d_i} \geq k$.

Recently, Allender et al. [1] showed that SqrtSum belongs to the fourth level of the *counting hierarchy*, a slight improvement over the previously known PSPACE upper bound. However, it has been an open question since the 1970s as to whether SqrtSum falls into the polynomial hierarchy [34, 30]. We identify a polynomial-time reduction from SqrtSum to StatNE for SSMGs.¹ Hence, StatNE is at least as hard as SqrtSum, and showing that StatNE resides inside the polynomial hierarchy would imply a major breakthrough in understanding the complexity of numerical computation.

Theorem 4.6. SqrtSum is polynomial-time reducible to StatNE, even for 4-player SSMGs.

Before we start with the proof of the theorem, let us first examine the game $\mathcal{G}(p)$, where $0 \leq p \leq 1$, played by players 0, 1, 2 and 3 and depicted in Fig. 7.

Lemma 4.7. The maximal payoff player 3 receives in a stationary Nash equilibrium of $(\mathcal{G}(p), s_1)$ where player 0 wins almost surely equals \sqrt{p} .

Proof. In the following, assume without loss of generality that $0 < p < 1$ (otherwise the statement is trivial), and define $q := 1 - p$. For any stationary strategy profile $\bar{\sigma}$ of $\mathcal{G}(p)$ where player 0 wins almost surely, let $x_1 = \sigma_0(s_2 \mid t_1)$ and $x_2 = \sigma_0(s_1 \mid t_2)$ be the probabilities

¹Some authors define SqrtSum using \leq instead of \geq . With this definition, we would reduce from the complement of SqrtSum instead.




FIGURE 6. Reducing SqrtSum to StatNE.




FIGURE 7. The game $\mathcal{G}(p)$.

that player 0 “stays inside the game” at t_1 , respectively t_2 . Given x_1 and x_2 , for $i = 1, 2$ we can compute the payoff $f_i(x_1, x_2) := \Pr_{s_i}^{\bar{\sigma}}(\text{Win}_i)$ for player i from s_i by

$$f_i(x_1, x_2) = \frac{p/2 + q(1 - x_i)}{1 - q^2 x_1 x_2}.$$

To have a Nash equilibrium, it must be the case that $f_1(x_1, x_2), f_2(x_1, x_2) \geq \frac{1}{2}$ since otherwise player 1 or player 2 would prefer to leave the game at s_1 or s_2 , respectively, which would give the respective player payoff $\frac{1}{2}$ immediately. Vice versa, if $f_1(x_1, x_2), f_2(x_1, x_2) \geq \frac{1}{2}$ then $\bar{\sigma}$ is a Nash equilibrium with expected payoff

$$f(x_1, x_2) := \frac{p + qx_1p}{1 - q^2 x_1 x_2}$$

for player 3. Hence, to determine the maximum payoff for player 3 in a stationary Nash equilibrium where player 0 wins almost surely, we have to maximise $f(x_1, x_2)$ under the constraints $f_1(x_1, x_2), f_2(x_1, x_2) \geq \frac{1}{2}$ and $0 \leq x_1, x_2 \leq 1$. We claim that the maximum is reached only if $x_1 = x_2$. If e.g. $x_1 > x_2$, then we can achieve a higher payoff for player 3 by setting $x'_2 := x_1$, and the constraints are still satisfied:

$$\frac{p/2 + q(1 - x'_2)}{1 - q^2 x_1 x'_2} = \frac{p/2 + q(1 - x_1)}{1 - q^2 x_1^2} \geq \frac{p/2 + q(1 - x_1)}{1 - q^2 x_1 x_2} \geq \frac{1}{2}.$$

Hence, it suffices to maximise $f(x, x)$ subject to $f_1(x, x) \geq \frac{1}{2}$ and $0 \leq x \leq 1$, which is equivalent to maximising $f(x, x)$ subject to $(1 - p)x^2 - 2x + 1 \geq 0$ and $0 \leq x \leq 1$. The roots of the polynomial are $(1 \pm \sqrt{p})/(1 - p)$, but $(1 + \sqrt{p})/(1 - p) > 1$ for $p > 0$. Therefore, any solution x must satisfy $x \leq x_0 := (1 - \sqrt{p})/(1 - p)$. Since $0 \leq x_0 \leq 1$ for $0 < p < 1$ and $f(x, x)$ is strictly increasing on $[0, 1]$, the optimal solution is x_0 , and the maximal payoff for player 3 in a stationary Nash equilibrium of $(\mathcal{G}(p), s_1)$ where player 0 wins almost surely equals indeed

$$f(x_0, x_0) = \frac{p + qx_0p}{1 - q^2 x_0^2} = \frac{p}{1 - qx_0} = \frac{p}{1 - (1 - p)x_0} = \frac{p}{1 - (1 - \sqrt{p})} = \sqrt{p}. \quad \square$$

Proof of Theorem 4.6. Given an instance (d_1, \dots, d_n, k) of SqrtSum, where without loss of generality $n > 0$, $d_i > 0$ for each $i = 1, \dots, n$, and $d := \sum_{i=1}^n d_i$, we construct a 4-player SSMG (\mathcal{G}, v_0) such that (\mathcal{G}, v_0) has a stationary Nash equilibrium where player 0 wins almost surely if and only if $\sum_{i=1}^n \sqrt{d_i} \geq k$. Define $p_i := d_i/d^2$ for $i = 1, \dots, n$. For the reduction, we use n copies of the game $\mathcal{G}(p)$, where in the i th copy we set p to p_i . The complete game \mathcal{G} is depicted in Fig. 6. By Lemma 4.7, the maximal payoff player 3 receives in a stationary Nash equilibrium of $(\mathcal{G}(p_i), s_1)$ where player 0 wins almost surely equals $\sqrt{p_i} = \sqrt{d_i}/d$. Hence, the maximal payoff player 3 receives in a stationary Nash equilibrium of (\mathcal{G}, v_1) where player 0 wins almost surely equals

$$\sum_{i=1}^n \frac{1}{n} \cdot \frac{\sqrt{d_i}}{d} = \frac{1}{dn} \cdot \sum_{i=1}^n \sqrt{d_i}.$$

If $\sum_{i=1}^n \sqrt{d_i} \geq k$, then we can extend such an equilibrium to a stationary Nash equilibrium of (\mathcal{G}, v_0) where player 0 wins almost surely by letting player 3 play from v_0 to v_1 with probability 1. On the other hand, if $\sum_{i=1}^n \sqrt{d_i} < k$, then in any stationary Nash equilibrium of (\mathcal{G}, v_0) player 3 plays to v_1 with probability 0, and player 0 loses almost surely. \square

Remark 4.8. The positive results of Sections 4.1 and 4.2 can easily be extended to equilibria in pure or randomised strategies with a memory of a fixed size $k \in \mathbb{N}$: a nondeterministic algorithm can guess a memory \mathfrak{M} of size k and then look for a positional, respectively stationary, equilibrium in the product of the original game \mathcal{G} with the memory \mathfrak{M} . Hence, for any fixed $k \in \mathbb{N}$, we can decide in PSPACE (NP) the existence of a randomised (pure) equilibrium of size k with payoff $\geq \bar{x}$ and $\leq \bar{y}$.

4.3. Pure and randomised equilibria. In this section, we show that the problems NE and PureNE are undecidable, by exhibiting a reduction from an undecidable problem about *two-counter machines*. Our construction is inspired by a construction used by Brázdil et al. [10] to prove the undecidability of stochastic games with *branching-time objectives* (see Remark 4.12 below).

Let $\Gamma := \{\text{inc}(j), \text{dec}(j), \text{zero}(j) : j = 1, 2\}$ (the set of *instructions*). A two-counter machine is of the form $\mathcal{M} = (Q, q_0, \delta)$, where

- Q is a finite set of *states*,
- $q_0 \in Q$ is the *initial state*, and
- $\delta \subseteq Q \times \Gamma \times Q$ is the *transition relation*.

For $q \in Q$ let $\delta(q) := \{(\gamma, q') \in \Gamma \times Q : (q, \gamma, q') \in \delta\}$. We call \mathcal{M} *deterministic* if for each $q \in Q$ either $\delta(q) = \emptyset$, or $\delta(q) = \{(\text{inc}(j), q')\}$ for some $j \in \{1, 2\}$ and $q' \in Q$, or $\delta(q) = \{(\text{zero}(j), q_1), (\text{dec}(j), q_2)\}$ for some $j \in \{1, 2\}$ and $q_1, q_2 \in Q$.

A configuration of \mathcal{M} is a triple $C = (q, i_1, i_2) \in Q \times \mathbb{N} \times \mathbb{N}$, where q denotes the current state and i_j denotes the current value of counter j . A configuration $C' = (q', i'_1, i'_2)$ is a *successor* of configuration $C = (q, i_1, i_2)$, denoted by $C \vdash C'$, if there exists a “matching” transition $(q, \gamma, q') \in \delta$. For example, $(q, i_1, i_2) \vdash (q', i_1+1, i_2)$ if and only if $(q, \text{inc}(1), q') \in \delta$. The instruction $\text{zero}(j)$ performs a *zero test*: $(q, i_1, i_2) \vdash (q', i_1, i_2)$ if and only if $i_1 = 0$ and $(q, \text{zero}(1), q') \in \delta$, or $i_2 = 0$ and $(q, \text{zero}(2), q') \in \delta$.

A *partial computation* of \mathcal{M} is a finite or infinite sequence $\rho = \rho(0)\rho(1)\dots$ of configurations such that $\rho(0) \vdash \rho(1) \vdash \dots$ and $\rho(0) = (q_0, 0, 0)$ (the *initial configuration*). A partial computation of \mathcal{M} is a *computation* of \mathcal{M} if it is infinite or ends in a configuration C for

which there for which there exists no successor configuration. Note that each deterministic two-counter machine has a unique computation.

The *halting problem* is to decide, given a machine \mathcal{M} , whether \mathcal{M} has a finite computation. It is well-known that deterministic two-counter machines are Turing powerful, which makes the halting problem and its dual, the *non-halting problem*, undecidable, even when restricted to deterministic two-counter machines. In fact, the non-halting problem for deterministic two-counter machines is not recursively enumerable.

Theorem 4.9. NE and PureNE are not recursively enumerable, even for 10-player SSMGs.

To prove Theorem 4.9, we give a reduction from the non-halting problem for deterministic two-counter machines. Our aim is thus to compute from a machine \mathcal{M} a 10-player SSMG (\mathcal{G}, v_0) such that the computation of \mathcal{M} is infinite if and only if (\mathcal{G}, v_0) has a (pure) Nash equilibrium in which player 0 wins almost surely. Without loss of generality, we assume that in \mathcal{M} there is no zero test that is followed by another zero test: if $(\text{zero}(j), q') \in \delta(q)$, then $|\delta(q')| \leq 1$.

The game \mathcal{G} is played by players 0, 1 and eight other players A_j^t and B_j^t , indexed by $j \in \{1, 2\}$ and $t \in \{0, 1\}$. Intuitively, player 0 and player 1 build up the computation of \mathcal{M} : player 0 updates the counters, and player 1 chooses transitions. The other players make sure that player 0 updates the counters correctly: players A_j^0 and A_j^1 ensure that, in each step, the value of counter j is not too high, and players B_j^0 and B_j^1 ensure that, in each step, the value of counter j is not too low. More precisely, A_j^0 and B_j^0 monitor the odd steps of the computation, while A_j^1 and B_j^1 monitor the even steps.

Let $\Gamma' := \Gamma \cup \{\text{init}\}$. For each $q \in Q$, each $\gamma \in \Gamma'$, each $j \in \{1, 2\}$ and each $t \in \{0, 1\}$, the game \mathcal{G} contains the gadgets $S_{\gamma,q}^t$, I_q^t and $C_{\gamma,j}^t$, which are depicted in Fig. 8. For better readability, terminal vertices are depicted as squares; the label indicates which players win. The initial vertex of \mathcal{G} is $v_0 := v_{\text{init},q_0}^0$. Note that in the gadget $S_{\gamma,q}^t$, each of the players A_j^t and B_j^t may *quit* the game, which gives her a payoff of $\frac{1}{3}$ or $\frac{1}{6}$, respectively, but payoff 0 to players 0 and 1.

It will turn out that player 1 will play a pure strategy in any Nash equilibrium of (\mathcal{G}, v_0) where player 0 wins almost surely, except possibly for histories that are not consistent with the equilibrium. Formally, we say that a strategy profile $\bar{\sigma}$ of (\mathcal{G}, v_0) is *safe* if for all histories xv consistent with $\bar{\sigma}$ and ending in a vertex $v \in I_q^t$ there exists $w \in V$ with $\sigma_1(w \mid xv) = 1$.

For any safe strategy profile $\bar{\sigma}$ of \mathcal{G} where player 0 wins almost surely, let $x_0v_0 \prec x_1v_1 \prec x_2v_2 \prec \dots$ (where $x_i \in V^*$, $v_i \in V$ and $x_0 = \varepsilon$) be the unique sequence containing all histories xv of (\mathcal{G}, v_0) that are consistent with $\bar{\sigma}$ and end in a vertex v of the form $v = v_{\gamma,q}^t$. This sequence is infinite because player 0 wins almost surely. Additionally, let q_0, q_1, \dots and $\gamma_0, \gamma_1, \dots$ be the corresponding sequences of states and instructions, respectively, i.e. $v_n = v_{\gamma_n, q_n}^0$ or $v_n = v_{\gamma_n, q_n}^1$ for all $n \in \mathbb{N}$. For each $j \in \{1, 2\}$ and $n \in \mathbb{N}$, we set:

$$\begin{aligned} a_j^n &:= \Pr_{v_0}^{\bar{\sigma}}(\text{player } A_j^{n \bmod 2} \text{ wins} \mid x_n v_n \cdot V^\omega); \\ b_j^n &:= \Pr_{v_0}^{\bar{\sigma}}(\text{player } B_j^{n \bmod 2} \text{ wins} \mid x_n v_n \cdot V^\omega). \end{aligned}$$

Note that at every terminal vertex of the counter gadgets $C_{\gamma,j}^t$ and $C_{\gamma,j}^{1-t}$ either player A_j^t or player B_j^t wins. For each j , the conditional probability that, given the history $x_n v_n$, we reach such a vertex is $\sum_{k \in \mathbb{N}} 1/2^k \cdot \frac{1}{4} = \frac{1}{2}$. Hence, $a_j^n = \frac{1}{2} - b_j^n$ for all $n \in \mathbb{N}$. We say that $\bar{\sigma}$ is *stable* if $a_j^n = \frac{1}{3}$ or, equivalently, $b_j^n = \frac{1}{6}$ for each $j \in \{1, 2\}$ and for all $n \in \mathbb{N}$.




FIGURE 8. Simulating a two-counter machine.

Finally, for each $j \in \{1, 2\}$ and $n \in \mathbb{N}$, we define a number $c_j^n \in [0, 1]$ as follows: After the history $x_n v_n$, with probability $\frac{1}{4}$ the play enters the counter gadget $C_{\gamma_n, j}^{n \bmod 2}$. The number c_j^n is defined as the probability of subsequently reaching a grey-coloured vertex. Note that, by the construction of \mathcal{G} , it holds that $c_j^n = 1$ if $\gamma_n = \text{zero}(j)$ or $\gamma_n = \text{init}$; in particular, $c_1^0 = c_2^0 = 1$.

Lemma 4.10. Let $\bar{\sigma}$ be a safe strategy profile of (\mathcal{G}, v_0) in which player 0 wins almost surely. Then $\bar{\sigma}$ is stable if and only if

$$c_j^{n+1} = \begin{cases} \frac{1}{2} \cdot c_j^n & \text{if } \gamma_{n+1} = \text{inc}(j), \\ 2 \cdot c_j^n & \text{if } \gamma_{n+1} = \text{dec}(j), \\ c_j^n = 1 & \text{if } \gamma_{n+1} = \text{zero}(j), \\ c_j^n & \text{otherwise,} \end{cases} \quad (4.1)$$

for each $j \in \{1, 2\}$ and $n \in \mathbb{N}$.

To prove the lemma, consider a safe strategy profile $\bar{\sigma}$ of (\mathcal{G}, v_0) in which player 0 wins almost surely. For each $j \in \{1, 2\}$ and $n \in \mathbb{N}$, set

$$p_j^n := \Pr_{v_0}^{\bar{\sigma}}(\text{player } A_j^{n \bmod 2} \text{ wins} \mid x_n v_n \cdot V^\omega \setminus x_{n+2} v_{n+2} \cdot V^\omega).$$

The following claim relates the numbers a_j^n and p_j^n .

Claim. Let $j \in \{1, 2\}$. Then $a_j^n = \frac{1}{3}$ for all $n \in \mathbb{N}$ if and only if $p_j^n = \frac{1}{4}$ for all $n \in \mathbb{N}$.

Proof. (\Rightarrow) Assume that $a_j^n = \frac{1}{3}$ for all $n \in \mathbb{N}$. We have $a_j^n = p_j^n + \frac{1}{4} \cdot a_j^{n+2}$ and therefore $\frac{1}{3} = p_j^n + \frac{1}{12}$ for all $n \in \mathbb{N}$. Hence, $p_j^n = \frac{1}{4}$ for all $n \in \mathbb{N}$.

(\Leftarrow) Assume that $p_j^n = \frac{1}{4}$ for all $n \in \mathbb{N}$. Since $a_j^n = p_j^n + \frac{1}{4} \cdot a_j^{n+2}$ for all $n \in \mathbb{N}$, the numbers a_j^n must satisfy the following recurrence: $a_j^{n+2} = 4a_j^n - 1$. Since all the numbers a_j^n are probabilities, $0 \leq a_j^n \leq 1$ for all $n \in \mathbb{N}$. It is easy to see that the only values for a_j^0 and a_j^1 such that $0 \leq a_j^n \leq 1$ for all $n \in \mathbb{N}$ are $a_j^0 = a_j^1 = \frac{1}{3}$. But this implies that $a_j^n = \frac{1}{3}$ for all $n \in \mathbb{N}$. \square

Proof of Lemma 4.10. By the previous claim, we only need to show that $p_j^n = \frac{1}{4}$ if and only if (4.1) holds. Let $j \in \{1, 2\}$, $n \in \mathbb{N}$ and $t = n \bmod 2$. The probability p_j^n can be expressed as the sum of the probability that the play reaches a terminal vertex that is winning for player A_j^t inside $C_{\gamma_n, j}^t$ and the probability that the play reaches such a vertex inside $C_{\gamma_{n+1}, j}^{1-t}$. The first probability does not depend on γ_n , but the second depends on γ_{n+1} . Let us consider the case that $\gamma_{n+1} = \text{inc}(j)$. In this case,

$$p_j^n = \frac{1}{4} \cdot (1 - \frac{1}{4} \cdot c_j^n) + \frac{1}{8} \cdot c_j^{n+1} = \frac{1}{4} - \frac{1}{16} \cdot c_j^n + \frac{1}{8} \cdot c_j^{n+1}.$$

Obviously, this sum is equal to $\frac{1}{4}$ if and only if $c_j^{n+1} = \frac{1}{2} \cdot c_j^n$. For any other value of γ_{n+1} , the argumentation is similar. \square

To establish the reduction, we need to show that the following statements are equivalent:

- (1) the computation of \mathcal{M} is infinite;
- (2) (\mathcal{G}, v_0) has a pure Nash equilibrium in which player 0 wins almost surely;
- (3) (\mathcal{G}, v_0) has a Nash equilibrium in which player 0 wins almost surely.

(1 \Rightarrow 2) Assume that the computation $\rho = \rho(0)\rho(1)\dots$ of \mathcal{M} is infinite. We define a pure strategy profile $\bar{\sigma}$ as follows: (1) For a history that ends at the unique vertex $v \in C_{\gamma,j}^t$ controlled by player 0 after visiting a vertex of the form $v_{\gamma',q}^t$ or $v_{\gamma',q}^{1-t}$ exactly $n > 0$ times and v exactly $k \geq 0$ times, player 0 plays to the grey-coloured successor vertex if k is greater than or equal to the value of counter j in configuration $\rho(n-1)$; otherwise, player 0 plays to the other successor vertex. (2) For a history that ends in one of the instruction gadgets I_q^t for $\delta(q) = \{(\text{zero}(j), q_1), (\text{dec}(j), q_2)\}$ after visiting a vertex of the form $v_{\gamma,q'}^t$ or $v_{\gamma,q'}^{1-t}$ exactly $n > 0$ times, player 1 plays to $S_{\text{zero}(j),q_1}^{1-t}$ if the value of counter j in configuration $\rho(n-1)$ is zero and to $S_{\text{dec}(j),q_2}^{1-t}$ if the value of counter j in configuration $\rho(n-1)$ is not zero. (3) Any other player's pure strategy is defined as follows: after a history ending in $S_{\gamma,q}^t$, the strategy prescribes to quit the game if and only if the history is not compatible with ρ (i.e. if the corresponding sequence of instructions does not match ρ).

Note that the resulting strategy profile $\bar{\sigma}$ is safe. Moreover, since the players follows the computation of \mathcal{M} , a terminal vertex inside one of the counter gadgets $C_{\gamma,j}^t$ is reached with probability 1 in $\bar{\sigma}$. Hence, player 0 wins almost surely. Moreover, by the definition of $\bar{\sigma}$, (4.1) holds, and we can conclude from Lemma 4.10 that $\bar{\sigma}$ is stable. We claim that $\bar{\sigma}$ is, in fact, a Nash equilibrium of (\mathcal{G}, v_0) : It is obvious that player 0 cannot improve her payoff. If player 1 deviates, we reach a history that is not compatible with ρ . Hence, player A_1^0 or A_2^0 will quit the game, which ensures that player 1 will not receive a higher payoff. Finally, since $\bar{\sigma}$ is stable, none of the players A_j^t or B_j^t can improve her payoff.

(2 \Rightarrow 3) Trivial.

(3 \Rightarrow 1) Assume that $\bar{\sigma}$ is a Nash equilibrium of (\mathcal{G}, v_0) in which player 0 wins almost surely. In order to apply Lemma 4.10, we first prove that $\bar{\sigma}$ is safe. By contradiction, assume that there exists a history xv consistent with $\bar{\sigma}$ and ending in a vertex $v \in I_q^t$ such that $\sigma_1(xv)$ assigns positive probability to two distinct successor vertices. Hence, $\delta(q) = \{(\text{zero}(j), q_1), (\text{dec}(j), q_2)\}$ for some $j \in \{1, 2\}$ and $q_1, q_2 \in Q$. By our assumption that there are no consecutive zero tests and since player 0 wins almost surely,

$$\Pr_{v_0}^{\bar{\sigma}}(\text{player 1 wins} \mid xv \cdot v_{\text{zero}(j),q_1}^{1-t} \cdot V^\omega) \geq \frac{1}{4},$$

but

$$\Pr_{v_0}^{\bar{\sigma}}(\text{player 1 wins} \mid xv \cdot v_{\text{dec}(j),q_2}^{1-t} \cdot V^\omega) \leq \frac{1}{6}.$$

Hence, player 1 could improve her payoff by playing to $v_{\text{zero}(j),q_1}^{1-t}$ with probability 1, a contradiction to $\bar{\sigma}$ being a Nash equilibrium.

To apply Lemma 4.10 and obtain (4.1), it remains to be shown that $\bar{\sigma}$ is stable. In order to derive a contradiction, assume that there exists $j \in \{1, 2\}$ and $n \in \mathbb{N}$ such that either $a_j^n < \frac{1}{3}$ or $a_j^n > \frac{1}{3}$ (i.e. $b_j^n < \frac{1}{6}$). In the former case, player $A_j^{n \bmod 2}$ could improve her payoff by quitting the game after history $x_n v_n$, while in the latter case, player $B_j^{n \bmod 2}$ could improve her payoff by quitting the game, again a contradiction to $\bar{\sigma}$ being a Nash equilibrium.

From $c_j^0 = 1$ and (4.1), it follows that each c_j^n is of the form $c_j^n = 1/2^i$ where $i \in \mathbb{N}$. We denote by i_j^n the unique number i such that $c_j^n = 1/2^i$ and set $\rho(n) = (q_n, i_1^n, i_2^n)$ for each $n \in \mathbb{N}$. We claim that $\rho := \rho(0)\rho(1)\dots$ is in fact the computation of \mathcal{M} . In particular, this computation is infinite. It suffices to verify the following two properties:

– $\rho(0) = (q_0, 0, 0)$;

– $\rho(n) \vdash \rho(n+1)$ for all $n \in \mathbb{N}$.

The first property is immediate. To prove the second property, let $\rho(n) = (q, i_1, i_2)$ and $\rho(n+1) = (q', i'_1, i'_2)$. Hence, v_n lies inside $S_{\gamma, q}^t$, and v_{n+1} lies inside $S_{\gamma', q'}^{1-t}$ for suitable γ, γ' and $t = n \bmod 2$. We only prove the claim for $\delta(q) = \{(\text{zero}(1), q_1), (\text{dec}(1), q_2)\}$; the other cases are similar. Note that, by the construction of the gadget I_q^t , it must be the case that either $q' = q_1$ and $\gamma' = \text{zero}(1)$, or $q' = q_2$ and $\gamma' = \text{dec}(1)$. By (4.1), if $\gamma' = \text{zero}(1)$, then $i'_1 = i_1 = 0$ and $i'_2 = i_2$, and if $\gamma' = \text{dec}(1)$, then $i'_1 = i_1 - 1$ and $i'_2 = i_2$. This implies $\rho(n) \vdash \rho(n+1)$: on the one hand, if $i_1 = 0$, then $i'_1 \neq i_1 - 1$, which implies $\gamma' \neq \text{dec}(1)$ and thus $\gamma' = \text{zero}(1)$, $q' = q_1$ and $i'_1 = i_1 = 0$; on the other hand, if $i_1 > 0$, then $\gamma' \neq \text{zero}(1)$ and thus $\gamma' = \text{dec}(1)$, $q' = q_2$ and $i'_1 = i_1 - 1$. \square

Remark 4.11. For the problem PureNE, we can strengthen Theorem 4.9 slightly by showing undecidability already for 9-player SSMGs. This can be achieved by merging player 0 and player 1 in the game described in the proof of Theorem 4.9.

Remark 4.12. The proof of Theorem 4.9 can also be viewed as a proof for the undecidability of a problem about the logic PCTL (*probabilistic computation tree logic*), introduced by Hansson and Jonsson [38]. PCTL is evaluated over labelled Markov chains and replaces the universal and existential path quantifiers of CTL by a family of probabilistic quantifiers $P^{\sim x}$, where \sim is a comparison operator and $x \in [0, 1]$ is a rational probability. For example, the formula $P^{=1/2}F Q$ holds in state v if (and only if) the probability of reaching a state labelled with Q from v equals $\frac{1}{2}$.

Brázdil et al. [10] proved the undecidability of the following problem: given a labelled Markov decision process (\mathcal{G}, v_0) and a PCTL formula φ , decide whether the controller has a strategy σ such that the Markov chain $(\mathcal{G}^\sigma, v_0)$ is a model of φ . We can prove a stronger result, namely that there exists a *fixed* PCTL formula φ , which only contains the quantifiers $P^{=x}F$ and $P^{=x}G$, for which the problem is undecidable. It suffices to add propositions $A_1^0, A_1^1, A_2^0, A_2^1, Q, Q_1, Q_2, T, Z_0$ and Z_1 according to the following rules:

- (1) if v is a terminal vertex that is winning for player $A \in \{A_1^0, A_1^1, A_2^0, A_2^1\}$, then label v with A ;
- (2) If v is controlled by player 0 and $|v\Delta| = 2$, then label v with Q and label one of its successors with Q_1 and the other with Q_2 .
- (3) if v is a terminal vertex that is winning for player 0, then label v with T ;
- (4) if $v = v_{\gamma, q}^0$, then label v with Z_0 ; if $v = v_{\gamma, q}^1$, then label v with Z_1 .

To obtain an MDP, we make all non-stochastic vertices controlled by player 0. Finally, the PCTL formula for which we prove undecidability is

$$P^{=1}F T \wedge \bigwedge_{t=0,1} P^{=1}G \left(Z_t \rightarrow \bigwedge_{j=1,2} P^{=1/3}F A_j^t \right) \wedge P^{=1}G \left(Q \rightarrow \bigvee_{i=0,1} P^{=1}F Q_i \right).$$

The first part of the formula states that player 0 wins almost surely, the second part requires the strategy to be stable, and the last part of the formula requires the strategy to be safe.

4.4. Finite-state equilibria. We can use the construction in the proof of Theorem 4.9 to show that Nash equilibria may require infinite memory, even if we are only interested in whether a player wins with probability 0 or 1.

Proposition 4.13. There exists a finite SSMG that has a pure Nash equilibrium where player 0 wins almost surely, but that has no finite-state Nash equilibrium where player 0 wins with positive probability.

Proof. Consider the game (\mathcal{G}, v_0) constructed in the proof of Theorem 4.9 for the machine \mathcal{M} with the single transition $(q_0, \text{inc}(1), q_0)$. We modify this game by adding a new initial vertex v_1 which is controlled by a new player, player 2, and from where she can either move to v_0 or to a new terminal vertex where she receives payoff 1 and every other player receives payoff 0. Additionally, player 2 wins at every terminal vertex of the game \mathcal{G} that is winning for player 0. Let us denote the modified game by \mathcal{G}' .

Since the computation of \mathcal{M} is infinite, the game (\mathcal{G}, v_0) has a pure Nash equilibrium where player 0 wins almost surely. This equilibrium induces a pure Nash equilibrium of (\mathcal{G}', v_1) where both player 0 and player 2 win almost surely.

Now assume that there exists a finite-state Nash equilibrium of (\mathcal{G}', v_1) where player 0 wins with positive probability. Such an equilibrium induces a finite-state Nash equilibrium $\bar{\sigma}$ of (\mathcal{G}, v_0) where player 2, and thus also player 0, wins almost surely; otherwise, player 2 would prefer to play from v_1 to the new terminal vertex. Using the same notation as in the proof of Theorem 4.9, it follows from Lemma 4.10 that $c_1^n = 1/2^n$ for each $n \in \mathbb{N}$. But this is impossible if $\bar{\sigma}$ is a finite-state strategy profile. \square

Propositions 4.1 and 4.13 imply that the decision problems NE, FinNE, PureNE and PureFinNE are pairwise distinct. Another way to see that PureNE and PureFinNE are distinct is to observe that PureFinNE is recursively enumerable: to decide whether an SMG (\mathcal{G}, v_0) has a pure finite-state Nash equilibrium with payoff $\geq \bar{x}$ and $\leq \bar{y}$, one can just enumerate all possible pure finite-state profiles $\bar{\sigma}$ and check for each of them whether it constitutes a Nash equilibrium with the desired properties by analysing the finite Markov chain $\mathcal{G}^{\bar{\sigma}}$ and the finite MDPs $\mathcal{G}^{\bar{\sigma}-i}$. Hence, to prove that PureFinNE is undecidable, we cannot reduce from the non-halting problem. Instead, we reduce from the halting problem (which is recursively enumerable itself). The same reduction proves that FinNE is undecidable.

Theorem 4.14. FinNE and PureFinNE are undecidable, even for 14-player SSMGs.

Proof sketch. The construction is similar to the one for proving the undecidability of NE. Given a two-counter machine \mathcal{M} , we modify the SSMG \mathcal{G} constructed in the proof of Theorem 4.9 by adding another counter (together with four more players for checking whether the counter is updated correctly) that has to be incremented in each step. Moreover, the gadget $I_{\gamma, q}$ for $\delta(q) = \emptyset$ is replaced by the gadget shown in Fig. 9, and a new instruction halt is added, together with a suitable gadget $C_{\text{halt}, j}^t$, also depicted in Fig. 9. Let us denote the new game by \mathcal{G}' . If \mathcal{M} does not halt, any Nash equilibrium of (\mathcal{G}', v_0) where player 0 wins with probability 1 needs infinite memory: to win almost surely, player 0 must follow the computation of \mathcal{M} and increment the new counter at each step, which requires infinite memory. On the other hand, if \mathcal{M} halts, there exists a pure finite-state Nash equilibrium of (\mathcal{G}', v_0) in which player 0 wins almost surely. (The arguments for the existence of such an equilibrium are the same as in the proof of Theorem 4.9; since \mathcal{M} halts, the equilibrium can be implemented with finite memory). \square

Remark 4.15. With the same reasoning as for PureNE, we can eliminate one player in the reduction for PureFinNE. Hence, this problem is already undecidable for SSMGs with 13 players.

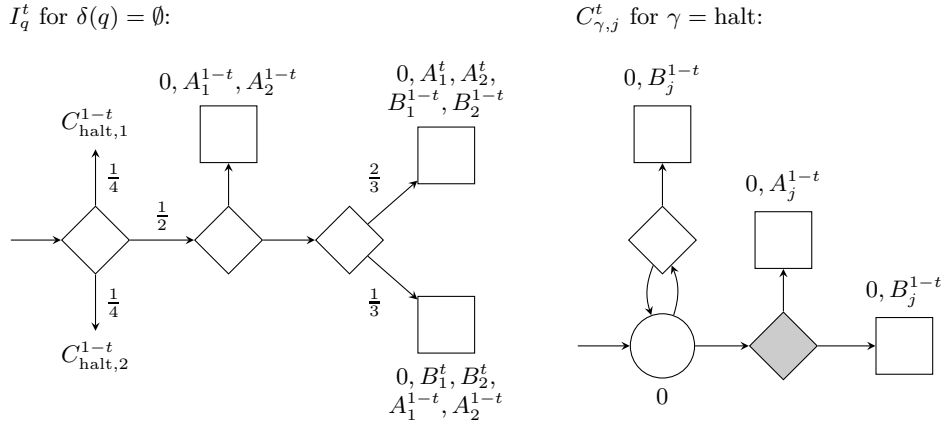


FIGURE 9. Reducing from the halting problem.

5. THE STRICTLY QUALITATIVE FRAGMENT

In this section, we prove that the fragment of NE that arises from restricting the thresholds to be the same binary payoff (i.e. each entry is either 0 or 1) is decidable for games with ω -regular objectives; we denote this problem by StrQualNE. Formally, StrQualNE is defined as follows:

Given an SMG (\mathcal{G}, v_0) and $\bar{x} \in \{0, 1\}^{\Pi}$, decide whether (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} .

To prove decidability, we first characterise the existence of a Nash equilibrium with a binary payoff in games with prefix-independent objectives.

5.1. Characterisation of existence. Given an SMG \mathcal{G} and a player i , we denote by W_i the set of all vertices $v \in V$ such that $\text{val}_i^{\mathcal{G}}(v) > 0$.

Proposition 5.1. Let (\mathcal{G}, v_0) be any SMG with prefix-independent objectives, and let $\bar{x} \in \{0, 1\}^{\Pi}$. Then the following statements are equivalent:

- (1) (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} ;
- (2) there exists a strategy profile $\bar{\sigma}$ of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$ for each player i with $x_i = 0$;
- (3) there exists a pure strategy profile $\bar{\sigma}$ of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$ for each player i with $x_i = 0$;
- (4) (\mathcal{G}, v_0) has a pure Nash equilibrium with payoff \bar{x} .

If additionally all objectives are ω -regular, then each of the above statements is equivalent to each of the following statements:

- (5) There exists a pure finite-state strategy profile $\bar{\sigma}$ with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$ for each player i with $x_i = 0$.
- (6) (\mathcal{G}, v_0) has a pure finite-state Nash equilibrium with payoff \bar{x} .

Proof. (1 \Rightarrow 2) Let $\bar{\sigma}$ be a Nash equilibrium of (\mathcal{G}, v_0) with payoff \bar{x} . We claim that $\bar{\sigma}$ is already the strategy profile we are looking for: $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$ for each player i with

$x_i = 0$. Let $i \in \Pi$ be a player with $x_i = 0$. By Lemma 3.3 and since Win_i is prefix-independent, we have $0 = \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) \geq \text{val}_i^{\mathcal{G}}(v)$ for all histories xv that are consistent with $\bar{\sigma}$. Hence, $v \in V \setminus W_i$ for all such histories xv , and $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$.

(2 \Rightarrow 3) Let $\bar{\sigma}$ be a strategy profile of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$ for each player i with $x_i = 0$. Consider the MDP \mathcal{M} that is obtained from \mathcal{G} by removing all vertices $v \in V$ such that $v \in W_i$ for some player i with $x_i = 0$, merging all players into one, and imposing the objective

$$\text{Win} = \bigcap_{\substack{i \in \Pi \\ x_i = 1}} \text{Win}_i \cap \bigcap_{\substack{i \in \Pi \\ x_i = 0}} C^\omega \setminus \text{Win}_i.$$

The MDP \mathcal{M} is well-defined since its domain is a subarena of \mathcal{G} . Moreover, the value $\text{val}^{\mathcal{M}}(v_0)$ of \mathcal{M} from v_0 equals 1 because the strategy profile $\bar{\sigma}$ induces a strategy σ in \mathcal{M} satisfying $\Pr_{v_0}^{\sigma}(\text{Win}) = 1$. Since each of the objectives Win_i is prefix-independent, so is the objective Win . Hence, by Theorem 2.5, (\mathcal{M}, v_0) admits an optimal pure strategy τ . Since $\text{val}^{\mathcal{M}}(v_0) = 1$, we have $\Pr_{v_0}^{\tau}(\text{Win}) = 1$, and τ induces a pure strategy profile of (\mathcal{G}, v_0) with the desired properties.

(3 \Rightarrow 4) Consider any pure strategy profile $\bar{\sigma}$ of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$ for each player i with $x_i = 0$. We show that $\bar{\sigma}$ is favourable: $\Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) \geq \text{val}_i^{\mathcal{G}}(v)$ for each player i and each history xv of (\mathcal{G}, v_0) that is consistent with $\bar{\sigma}$. There are two cases: If $x_i = 1$, then $\Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i \mid xv \cdot V^\omega) = 1$ for all histories xv consistent with $\bar{\sigma}$, and the inequality holds. Otherwise, $x_i = 0$ and $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(W_i)) = 0$. Hence, $\text{val}_i^{\mathcal{G}}(v) = 0$ for all histories xv consistent with $\bar{\sigma}$, and the inequality holds as well. Now, by Lemma 3.4, we can extend $\bar{\sigma}$ to a pure Nash equilibrium with payoff \bar{x} .

(4 \Rightarrow 1) Trivial.

Under the additional assumption that all objectives are ω -regular, the implications (2 \Rightarrow 5) and (5 \Rightarrow 6) are proven analogously (using Lemma 3.5 instead of Lemma 3.4); the implication (6 \Rightarrow 1) is trivial. \square

As an immediate consequence of Proposition 5.1, we can conclude that pure finite-state strategies are as powerful as arbitrary randomised strategies as far as the existence of Nash equilibria with binary payoffs in finite SMGs with ω -regular objectives is concerned.

Corollary 5.2. Let (\mathcal{G}, v_0) be a finite SMG with ω -regular objectives, and let $\bar{x} \in \{0, 1\}^\Pi$. There exists a Nash equilibrium of (\mathcal{G}, v_0) with payoff \bar{x} if and only if there exists a pure finite-state Nash equilibrium of (\mathcal{G}, v_0) with payoff \bar{x} .

Proof. The claim follows from Proposition 5.1 and the fact that every SMG with ω -regular objectives can be reduced to one with parity objectives (using finite memory). \square

5.2. Computational Complexity. We can now give an algorithm that decides StrQualNE for SMGs with Muller objectives. The algorithm relies on Proposition 5.1, which allows us to reduce StrQualNE to an MDP problem.

Formally, given a Muller SMG $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, \Delta, \chi, (\mathcal{F}_i)_{i \in \Pi})$ and a binary payoff $\bar{x} = (x_i)_{i \in \Pi}$, we define the Markov decision process $\mathcal{G}(\bar{x})$ as follows: Let $Z \subseteq V$ be the set of all vertices v such that $\text{val}_i^{\mathcal{G}}(v) = 0$ for each player i with $x_i = 0$; the set of vertices of $\mathcal{G}(\bar{x})$ is precisely the set Z , with the set of vertices controlled by player 0 being $Z_0 := \bigcup_{i \in \Pi} (V_i \cap Z)$; if $Z = \emptyset$, we define $\mathcal{G}(\bar{x})$ to be a trivial MDP with the empty set as its objective. The transition

relation of $\mathcal{G}(\bar{x})$ is the restriction of Δ to transitions between Z -states. Note that the transition relation of $\mathcal{G}(\bar{x})$ is well-defined since Z is a subarena of \mathcal{G} . Finally, the single objective in $\mathcal{G}(\bar{x})$ is $\text{Reach}(T)$ where $T \subseteq Z$ is the union of all end components $U \subseteq Z$ with payoff \bar{x} .

Lemma 5.3. Let (\mathcal{G}, v_0) be a finite Muller SMG, and let $\bar{x} \in \{0, 1\}^\Pi$. Then (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} if and only if $\text{val}^{\mathcal{G}(\bar{x})}(v_0) = 1$.

Proof. (\Rightarrow) Assume that (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} . By Proposition 5.1, there exists a strategy profile $\bar{\sigma}$ of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V \setminus Z)) = 0$. We claim that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(T)) = 1$. Otherwise, by Lemma 2.2, there would exist an end component $U \subseteq Z$ such that $\Pr_{v_0}^{\bar{\sigma}}(\{\pi \in V^\omega : \text{Inf}(\pi) = U\}) > 0$, and U is either not winning for some player i with $x_i = 1$ or it is winning for some player i with $x_i = 0$. But then $\bar{\sigma}$ cannot have payoff \bar{x} , a contradiction. Now, since $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V \setminus Z)) = 0$, the strategy profile $\bar{\sigma}$ induces a strategy σ in $\mathcal{G}(\bar{x})$ such that $\Pr_{v_0}^\sigma(X) = \Pr_{v_0}^{\bar{\sigma}}(X)$ for every Borel set $X \subseteq Z^\omega$. In particular, $\Pr_{v_0}^\sigma(\text{Reach}(T)) = 1$ and hence $\text{val}^{\mathcal{G}(\bar{x})}(v_0) = 1$.

(\Leftarrow) Assume that $\text{val}^{\mathcal{G}(\bar{x})}(v_0) = 1$ (in particular, $v_0 \in Z$), and let σ be an optimal strategy in $(\mathcal{G}(\bar{x}), v_0)$. From σ , using Lemma 2.3, we can devise a strategy σ' such that $\Pr_{v_0}^{\sigma'}(\{\pi \in V^\omega : \text{Inf}(\pi) \text{ has payoff } \bar{x}\}) = 1$. Finally, σ' can be extended to a strategy profile $\bar{\sigma}$ of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V \setminus Z)) = 0$. By Proposition 5.1, this implies that (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} . \square

Since the values of an MDP with a reachability objective can be computed in polynomial time, the difficult part lies in computing the MDP $\mathcal{G}(\bar{x})$ from \mathcal{G} and \bar{x} (i.e. its domain Z and the target set T). For Muller SMGs, polynomial space suffices to achieve this. In fact, StrQualNE is PSPACE-complete for these games.

Theorem 5.4. StrQualNE is PSPACE-complete for Muller SMGs.

Proof. Hardness follows from Theorem 2.11. To prove membership in PSPACE, we describe a polynomial-space algorithm for deciding StrQualNE on Muller SMGs: On input $\mathcal{G}, v_0, \bar{x}$, the algorithm starts by computing for each player i with $x_i = 0$ the set of vertices v such that $\text{val}_i^{\mathcal{G}}(v) = 0$, which can be done in polynomial space by Theorem 2.11. The intersection of these sets is the domain Z of the Markov decision process $\mathcal{G}(\bar{x})$. If v_0 is not contained in this intersection, the algorithm immediately rejects. Otherwise, the algorithm determines the union T of all end components with payoff \bar{x} contained in Z by enumerating all subsets of Z , one at a time, and checking which ones are end components with payoff \bar{x} . Finally, the algorithm computes (in polynomial time) the value $\text{val}^{\mathcal{G}(\bar{x})}(v_0)$ of the MDP $\mathcal{G}(\bar{x})$ from v_0 and accepts if this value is 1. In all other cases, the algorithm rejects. The correctness of the algorithm follows immediately from Lemma 5.3. \square

For games with Streett objectives, StrQualNE becomes NP-complete; we start by proving the upper bound.

Theorem 5.5. StrQualNE is in NP for Streett SMGs.

Proof. We describe a nondeterministic polynomial-time algorithm for solving StrQualNE: On input $\mathcal{G}, v_0, \bar{x}$, the algorithm starts by guessing a subarena $Z' \subseteq V$ and for each player i with $x_i = 0$ a positional strategy τ_i of the coalition $\Pi \setminus \{i\}$ in the coalition game \mathcal{G}_i . In the next step, the algorithm checks (in polynomial time) whether $\text{val}^{\tau_i}(v) = 1$ for each




FIGURE 10. Reducing SAT to StrQualNE for games with Streott objectives.

vertex $v \in Z'$ and each player i with $x_i = 0$. If not, the algorithm rejects immediately. Otherwise, the algorithm proceeds by guessing (at most) $n := |V|$ subsets $U_1, \dots, U_n \subseteq Z'$ and checks whether they are end components with payoff \bar{x} (which can be done in polynomial time). If yes, the algorithm sets $T' := \bigcup_{j=1}^n U_j$ and computes (in polynomial time) the value $\text{val}^{\mathcal{G}(\bar{x})}(v_0)$ of the MDP $\mathcal{G}(\bar{x})$ from v_0 with Z' substituted for Z and T' substituted for T . If this value equals 1, the algorithm accepts; otherwise, it rejects.

It remains to be shown that the algorithm is correct: On the one hand, if (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} , then the run of the algorithm where it guesses $Z' = Z$, globally optimal positional strategies τ_i (which exist by Theorem 2.6) and end components U_i such that $T' = T$ will be accepting since then, by Lemma 5.3, $\text{val}^{\mathcal{G}(\bar{x})}(v_0) = 1$. On the other hand, in any accepting run of the algorithm we have $Z' \subseteq Z$ and $T' \subseteq T$, and the computed value cannot be higher than $\text{val}^{\mathcal{G}(\bar{x})}(v_0)$; hence, $\text{val}^{\mathcal{G}(\bar{x})}(v_0) = 1$, and Lemma 5.3 guarantees the existence of a Nash equilibrium with payoff \bar{x} . \square

The matching lower bound does even hold for deterministic two-player Streott games and was established in [58].

Theorem 5.6. StrQualNE is NP-hard for deterministic two-player Streott games.

Proof. The proof is accomplished by a variant of the proof for NP-hardness of the qualitative decision problem for deterministic two-player zero-sum Rabin-Streott games [29] and by a reduction from SAT. Given a Boolean formula $\varphi = C_1 \wedge \dots \wedge C_m$ in conjunctive normal form, where without loss of generality $m \geq 1$ and each clause is nonempty, we construct a deterministic two-player Streott game \mathcal{G} as follows: For each clause C , the game \mathcal{G} has a vertex C , which is controlled by player 0, and for each literal L occurring in φ , there is a vertex L , which is controlled by player 1. There are edges from a clause to each literal that occurs in this clause, and from a literal to each clause occurring in φ . The structure of the game is depicted in Fig. 10. Player 0's objective is given by the empty Streott objective, i.e. she wins every play of the game, whereas player 1's objective consists of all Streott pairs of the form $(\{X\}, \{\neg X\})$ or $(\{\neg X\}, \{X\})$, i.e. she wins if, for each variable X , either X and $\neg X$ are both visited infinitely often or neither of them is.

Clearly, \mathcal{G} can be constructed from φ in polynomial time. We claim that φ is satisfiable if and only if (\mathcal{G}, C_1) has a Nash equilibrium with payoff $(1, 0)$.

(\Rightarrow) Assume that φ is satisfiable, and consider the following positional strategy σ_0 of player 0: whenever the play reaches a clause, then σ_0 plays to a literal that is mapped to

true by the satisfying assignment. This strategy ensures that for each variable X at most one of the literals X or $\neg X$ is visited infinitely often. Hence, (σ_0, σ_1) is a Nash equilibrium of (\mathcal{G}, C_1) with payoff $(1, 0)$ for every strategy σ_1 of player 1.

(\Leftarrow) Let (σ_0, σ_1) be a Nash equilibrium of (\mathcal{G}, C_1) with payoff $(1, 0)$, and assume that φ is not satisfiable. Consider the two-player zero-sum Rabin-Streett game $\tilde{\mathcal{G}}$, which is derived from \mathcal{G} by setting player 0's objective to the complement of player 1's objective. We claim that player 1 has a winning strategy in $(\tilde{\mathcal{G}}, C_1)$, which she could use to improve her payoff in (\mathcal{G}, C_1) , a contradiction to (σ_0, σ_1) being a Nash equilibrium. By determinacy, we only need to show that player 0 does not have a winning strategy. Let τ be an optimal positional strategy of player 0 in $(\tilde{\mathcal{G}}, C_1)$ (which exists by Theorem 2.6). Since φ is unsatisfiable, there must exist a variable X and clauses C and C' such that $\tau(C) = X$ and $\tau(C') = \neg X$. But player 1 can counter this strategy by playing from X to C' and from any other literal to C . Hence, τ is not winning. \square

For games with Rabin objectives, the situation is more delicate. One might think that, because of the duality of Rabin and Streett objectives, StrQualNE is in coNP for SMGs with Rabin objectives.² However, as we will see later, this is rather unlikely, and we can only show that the problem lies in the class $\text{P}^{\text{NP}[\log]}$ of problems solvable by a deterministic polynomial-time algorithm that may perform a logarithmic number of queries to an NP oracle. In fact, the same upper bound holds for games with a Streett *or* a Rabin objective for each player.

Theorem 5.7. StrQualNE is in $\text{P}^{\text{NP}[\log]}$ for Streett-Rabin SMGs.

Proof. Let us describe a polynomial-time algorithm performing a logarithmic number of queries to an NP oracle for the problem. On input $\mathcal{G}, v_0, \bar{x}$, the algorithm starts by determining for each vertex v and each Rabin player i with $x_i = 0$ whether $\text{val}_i^{\mathcal{G}}(v) = 0$. Naively implemented, this requires a super-logarithmic number of queries to the oracle. To reduce the number of queries, we use a neat trick, due to Hemachandra [39]. Let us denote by R and S the set of players $i \in \Pi$ with $x_i = 0$ who have a Rabin, respectively a Streett objective. Instead of looping through all pairs of a vertex and a player, we start by determining the number r of all pairs (v, i) such that $i \in R$ and $\text{val}_i^{\mathcal{G}}(v) = 0$. It is not difficult to see that this number can be computed using binary search by performing only a logarithmic number of queries to an NP oracle, which we can use for deciding whether $\text{val}_i^{\mathcal{G}}(v) > 0$ (Corollary 2.10). Then we perform one more query; we ask whether for each player $i \in R \cup S$ there exists a set $Z_i \subseteq V$ as well as sets $U_1, \dots, U_{|V|} \subseteq V$ and positional strategies $(\sigma_i)_{i \in R}$ and $(\tau_i)_{i \in S}$, where σ_i is a strategy of player i and τ_i is a strategy of the coalition $\Pi \setminus \{i\}$ in the coalition game \mathcal{G}_i , with the following properties:

- (1) $Z := \bigcap_{i \in R \cup S} Z_i$ is a subarena of \mathcal{G} with $v_0 \in Z$, and $\sum_{i \in R} |Z_i| = r$;
- (2) $\text{val}^{\sigma_i}(v) > 0$ for each player $i \in R$ and each $v \in V \setminus Z_i$;
- (3) $\text{val}^{\tau_i}(v) = 1$ for each player $i \in S$ and each $v \in Z_i$;
- (4) each U_j is an end component of $\mathcal{G} \upharpoonright Z$ with payoff \bar{x} ;
- (5) the value from v_0 of the MDP that is obtained from \mathcal{G} by restricting to vertices inside Z and imposing the objective $\text{Reach}(\bigcup\{U_1, \dots, U_{|V|}\})$ equals 1.

This query can be decided by an NP oracle by guessing suitable sets and strategies and verifying (1)–(5) in polynomial time. If the answer to the query is *yes*, the algorithm accepts; otherwise it rejects.

²In fact, Ummels and Wojtczak [60] claimed that the problem is in coNP.

Obviously, the algorithm runs in polynomial time. To see that the algorithm is correct, first note that for each player $i \in R$ the set Z_i does not only include all $v \in V$ such that $\text{val}_i^{\mathcal{G}}(v) = 0$, but also excludes all other vertices. Otherwise, there would exist a vertex $v \in Z_i$ with $\text{val}_i^{\mathcal{G}}(v) > 0$. But then the number of pairs (v, i) with $i \in R$ and $\text{val}_i^{\mathcal{G}}(v) = 0$ would be strictly less than r , a contradiction. Now, the correctness of the algorithm follows with the same reasoning as in the proof of Theorem 5.5. \square

Remark 5.8. For a bounded number of players, StrQualNE is in coNP for SMGs with Rabin objectives.

Regarding lower bounds for StrQualNE in SMGs with Rabin objectives, we start by proving that the problem is coNP-hard, even for deterministic two-player games.

Theorem 5.9. StrQualNE is coNP-hard for deterministic two-player Rabin games.

Proof. The proof is similar to the proof of Theorem 5.6 and is accomplished by a reduction from the unsatisfiability problem for Boolean formulae in conjunctive normal form. Given a Boolean formula $\varphi = C_1 \wedge \dots \wedge C_m$ in conjunctive normal form, where without loss of generality $m \geq 1$ and each clause is nonempty, we construct a deterministic two-player Rabin game \mathcal{G} as follows. The arena of \mathcal{G} is the same as in the proof of Theorem 5.6, depicted in Fig. 10. However, this time player 1 wins every play of the game (her objective consists of the single Rabin pair (V, \emptyset)), and player 0's objective consists of all Rabin pairs of the form $(\{X\}, \{\neg X\})$ or $(\{\neg X\}, \{X\})$.

Clearly, \mathcal{G} can be constructed from φ in polynomial time. We claim that the φ is unsatisfiable if and only if (\mathcal{G}, C_1) has a Nash equilibrium with payoff $(0, 1)$.

(\Rightarrow) Assume that φ is unsatisfiable, and consider the two-player zero-sum Rabin-Streett game $\tilde{\mathcal{G}}$, which is derived from \mathcal{G} by setting player 1's objective to the complement of player 0's objective. Let σ_1 be a globally optimal strategy for player 1 in this game. We claim that σ_1 is winning in (\mathcal{G}_0, C_1) . Consequently, (σ_0, σ_1) is a Nash equilibrium of (\mathcal{G}, C_1) with payoff $(0, 1)$ for every strategy σ_0 of player 0. Otherwise, player 0 would have a positional winning strategy in $(\tilde{\mathcal{G}}, C_1)$. But a positional strategy τ of player 0 picks for each clause a literal contained in this clause. Since φ is unsatisfiable, there must exist a variable X and clauses C and C' such that $\tau(C) = X$ and $\tau(C') = \neg X$. Player 1 could counter this strategy by playing from X to C' and from any other literal to C , a contradiction.

(\Leftarrow) Let (σ_0, σ_1) be a Nash equilibrium of (\mathcal{G}, C_1) with payoff $(0, 1)$, and assume that φ is satisfiable. Consider the following positional strategy τ of player 0: whenever the play reaches a clause, then τ plays to a literal that is mapped to true by the satisfying assignment. This strategy ensures that for each variable X at most one of the literals X or $\neg X$ is visited infinitely often. Since the construction of \mathcal{G} ensures that, under any strategy profile, at least one literal is visited infinitely often, τ ensures a winning play for player 0. Hence, player 0 can improve her payoff by playing τ instead of σ_0 , a contradiction to the fact that (σ_0, σ_1) is a Nash equilibrium. \square

The next result shows that StrQualNE is not only coNP-hard for Rabin games, but also NP-hard. In fact, it is even NP-hard to decide whether in a deterministic Rabin game there exists a play that fulfils the objective of each player.

Proposition 5.10. The problem of deciding, given a deterministic Rabin game, whether there exists a play that is won by each player is NP-hard.




FIGURE 11. Reducing SAT to deciding the existence of a play winning for all players in a deterministic Rabin game.

Proof. We reduce from SAT: given a Boolean formula $\varphi = C_1 \wedge \dots \wedge C_m$ in conjunctive normal form over propositional variables X_1, \dots, X_n , where without loss of generality $m \geq 1$ and each clause is nonempty, we show how to construct in polynomial time a deterministic $(n + 1)$ -player Rabin game \mathcal{G} such that φ is satisfiable if and only if there exists a play of \mathcal{G} that is won by each player. The game has vertices C_1, \dots, C_m and, for each clause C and each literal L that occurs in C , a vertex (C, L) . All vertices are controlled by player 0. There are edges from a clause C_j to each vertex (C_j, L) such that L occurs in C_j and from there to $C_{(j \bmod m)+1}$. The arena of \mathcal{G} is schematically depicted in Fig. 11. The Rabin objectives are defined as follows:

- player 0 wins every play of \mathcal{G} ;
- player $i \neq 0$ wins if each vertex of the form (C, X_i) is visited only finitely often *or* each vertex of the form $(C, \neg X_i)$ is visited only finitely often.

Clearly, \mathcal{G} can be constructed from φ in polynomial time. To establish the reduction, we need to show that φ is satisfiable if and only if there exists a play of \mathcal{G} that is won by each player.

(\Rightarrow) Assume that $\alpha: \{X_1, \dots, X_n\} \rightarrow \{\text{true}, \text{false}\}$ is a satisfying assignment of φ . Clearly, the positional strategy of player 0 where from each clause C she plays to a fixed vertex (C, L) such that L is mapped to true by α induces a play that is won by each player.

(\Leftarrow) Assume that there exists a play π of \mathcal{G} that is won by each player. Obviously, it is not possible that both a vertex (C, X_i) and a vertex $(C', \neg X_i)$ are visited infinitely often in π since this would violate player i 's objective. Consider the variable assignment that maps X to true if some vertex (C, X) is visited infinitely often in π . This assignment satisfies the formula because, by the construction of \mathcal{G} , for each clause C there exists a literal L in C such that the vertex (C, L) is visited infinitely often in π . \square

It follows from Theorem 5.9 and Proposition 5.10 that, unless $\text{NP} = \text{coNP}$, StrQualNE is not contained in $\text{NP} \cup \text{coNP}$, even for deterministic Rabin games. With a little more effort, one can show that StrQualNE is DP-hard for deterministic Rabin games (see [59]). Finally, for stochastic Rabin games, we can show that StrQualNE is $\text{P}^{\text{NP}[\log]}$ -complete.

Theorem 5.11. StrQualNE is $\text{P}^{\text{NP}[\log]}$ -hard for Rabin SMGs.

Proof. Wagner [62] and, independently, Buss and Hay [11] showed that $P^{\text{NP}[\log]}$ is the closure of NP with respect to *polynomial-time Boolean formula reducibility*. The canonical complete problem for this class is to decide, given a Boolean combination α of statements of the form “ φ is satisfiable”, where φ ranges over all Boolean formulae, whether α evaluates to true. We claim that for every such statement α we can construct in polynomial time a Rabin SMG (\mathcal{G}, v_0) such that α evaluates to true if and only if (\mathcal{G}, v_0) has a Nash equilibrium with payoff $(0, 1, \dots, 1)$. The game \mathcal{G} is constructed by induction on the complexity of α ; without loss of generality, we assume that negations are only applied to atoms. If α is of the form “ φ is satisfiable” or “ φ is not satisfiable”, then the existence of a suitable game \mathcal{G} follows from Proposition 5.10 or Theorem 5.9, respectively.

Now, let $\alpha = \alpha_1 \wedge \alpha_2$, and assume that we already have constructed suitable games (\mathcal{G}_1, v_1) and (\mathcal{G}_2, v_2) , played by the same players $0, 1, \dots, n$. The game \mathcal{G} is the disjoint union of \mathcal{G}_1 and \mathcal{G}_2 combined with one new stochastic vertex v_0 . From v_0 , the game moves with probability $\frac{1}{2}$ each to v_1 or v_2 . Obviously, (\mathcal{G}, v_0) has a Nash equilibrium with payoff $(0, 1, \dots, 1)$ if and only if both (\mathcal{G}_1, v_1) and (\mathcal{G}_2, v_2) have such an equilibrium.

Finally, let $\alpha = \alpha_1 \vee \alpha_2$, and assume that we already have constructed suitable games (\mathcal{G}_1, v_1) and (\mathcal{G}_2, v_2) , again played by the same players $0, 1, \dots, n$. As in the previous case, the game \mathcal{G} is the disjoint union of \mathcal{G}_1 and \mathcal{G}_2 combined with one new vertex v_0 , which has transitions to both v_1 and v_2 . However, this time v_0 is controlled by player 1. Obviously, (\mathcal{G}, v_0) has a Nash equilibrium with payoff $(0, 1, \dots, 1)$ if and only if at least one of the games (\mathcal{G}_1, v_1) and (\mathcal{G}_2, v_2) has such an equilibrium. \square

Our next aim is to prove that StrQualNE is in $\text{UP} \cap \text{coUP}$ for parity SMGs. We will make use of Algorithm 5.1, which computes for a game \mathcal{G} with priority functions $(\Omega_i)_{i \in \Pi}$ and $\bar{x} \in \{0, 1\}^\Pi$ the union of all end components with payoff \bar{x} . The algorithm is a straightforward adaptation of the algorithm for computing the union of all winning end components in a Streett MDP [17]. At the heart of the algorithm lies the procedure FindEC that returns on input $X \subseteq V$ the union of all end components with payoff \bar{x} that are contained in X . The procedure starts by computing all end components maximal in X . If such an end component U has payoff \bar{x} , all vertices in U can be added to the result of the procedure. Otherwise, there exists a player i such that either $x_i = 0$ and the least priority for player i in U is odd or $x_i = 1$ and the least priority for player i in U is even. Each end component with payoff \bar{x} inside U must exclude all vertices with this least priority. Hence, we call the procedure recursively on the subset of U that results from removing these vertices.

Note that on input X , the total number of recursive calls to the procedure FindEC is bounded by $|X|$. Since, additionally, the set of all end components maximal in a set X can be computed in polynomial time, this proves that Algorithm 5.1 runs in polynomial time.

Theorem 5.12. StrQualNE is in $\text{UP} \cap \text{coUP}$ for parity SMGs.

Proof. A UP algorithm that decides StrQualNE for parity SMGs works as follows: On input $\mathcal{G}, v_0, \bar{x}$, the algorithm starts by guessing, for each player i with $x_i = 0$, the set Z_i of vertices v with $\text{val}_i^{\mathcal{G}}(v) = 0$. Then, for each $v \in V$, the guess whether $v \in Z_i$ or $v \notin Z_i$ is verified by running the UP algorithm for the respective problem. If some guess was not correct, the algorithm rejects immediately. Otherwise, it constructs the subarena $Z := \bigcap_{i \in \Pi: x_i = 0} Z_i$ and uses Algorithm 5.1 to determine the union T of all end components with payoff \bar{x} . If $v_0 \notin Z$, the algorithm rejects immediately. Otherwise, it computes in polynomial time the value $\text{val}^{\mathcal{G}(\bar{x})}(v_0)$ of the MDP $\mathcal{G}(\bar{x})$ from v_0 . If this value equals 1, the algorithm accepts;

Algorithm 5.1. Finding end components in parity SMGs.

Input: parity SMG $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, \Delta, \chi, (\Omega_i)_{i \in \Pi})$, $\bar{x} = (x_i)_{i \in \Pi} \in \{0, 1\}^\Pi$

Output: $\bigcup \{U \subseteq V : U \text{ is an end component of } \mathcal{G} \text{ with payoff } \bar{x}\}$

output FindEC(V)

procedure FindEC(X)

$Z := \emptyset$

compute all end components of \mathcal{G} maximal in X

for each such end component U **do**

$P := \{i \in \Pi : \min \Omega_i(\chi(U)) \equiv x_i \pmod{2}\}$

if $P = \emptyset$ **then**

(* U is an end component with payoff \bar{x} *)

$Z := Z \cup U$

else

(* U has the wrong payoff *)

$Y := \bigcap_{i \in P} \{v \in U : \Omega_i(\chi(v)) > \min \Omega_i(\chi(U))\}$

$Z := Z \cup \text{FindEC}(Y)$

end if

end for

return Z

end procedure

otherwise, it rejects. Analogously, an algorithm for the complement of StrQualNE accepts if and only if $v_0 \notin Z$ or $\text{val}^{\mathcal{G}(\bar{x})}(v_0) < 1$.

Obviously, both algorithms run in polynomial time. Moreover, on each input there exists at most one accepting run because the algorithms only accept if each of the sets Z_i has been guessed correctly. Finally, their correctness follows from Lemma 5.3. \square

Recall from Section 2.6 that it is an open question whether the qualitative decision problem for parity S2Gs admits a polynomial-time algorithm. Such an algorithm would allow us compute the domain of the MDP $\mathcal{G}(\bar{x})$ efficiently, which would imply that StrQualNE is in P for parity SMGs. In fact, given a class \mathcal{C} of parity S2Gs for which the qualitative decision problem is in P, we can easily derive a class of parity SMGs for which StrQualNE is in P, namely the class \mathcal{C}^* of all parity SMGs such that for each player i the coalition game \mathcal{G}_i is in \mathcal{C} .

Theorem 5.13. Let \mathcal{C} be a class of finite parity S2Gs such that the qualitative decision problem is decidable in P for games in \mathcal{C} . Then StrQualNE is in P for games in \mathcal{C}^* .

Proof. Consider the algorithm given in the proof of Theorem 5.12. For each player i , the set Z_i can be computed in polynomial time if $\mathcal{G}_i \in \mathcal{C}$, and there is no need to guess this set. The resulting deterministic algorithm still runs in polynomial time. \square

By Theorem 2.12, for each $d \in \mathbb{N}$, the qualitative decision problem for parity S2Gs with at most d priorities belongs to P. Hence, it follows from Theorem 5.13 that StrQualNE

is decidable in polynomial time for parity SMGs with at most d priorities. In particular, StrQualNE is in P for (co-)Büchi SMGs.

Corollary 5.14. For each $d \in \mathbb{N}$, StrQualNE is decidable in polynomial time for parity SMGs with at most d priorities.

6. CONCLUSION

We have analysed the complexity of deciding whether a stochastic multiplayer game with ω -regular objectives has a Nash equilibrium whose payoff falls into a certain interval. Our results demonstrate that this problem is more complicated for multiplayer games than for two-player zero-sum games. In particular, the problem of deciding the existence of a Nash equilibrium where player 0 wins almost surely is undecidable for simple stochastic multiplayer games, whereas the same problem is decidable in polynomial time for two-player zero-sum simple stochastic games. On the positive side, we have shown that the strictly qualitative fragment of NE has a complexity that is comparable to the complexity of the qualitative decision problem for two-player zero-sum games.

Several directions for future research come to mind: First, one can study other restrictions of NE that might be decidable. For example, it is plausible that the restriction of NE to games with two players is decidable. Second, it would be interesting to extend our results to other game models such as *concurrent games* [55, 27] or games with quantitative payoff functions.

REFERENCES

- [1] E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. On the complexity of numerical analysis. *SIAM Journal on Computing*, 38(5):1987–2006, 2009.
- [2] D. Andersson and P. B. Miltersen. The complexity of solving stochastic games on graphs. In *Proceedings of the 20th International Symposium on Algorithms and Computation, ISAAC 2009*, volume 5878 of LNCS, pages 112–121. Springer-Verlag, 2009.
- [3] R. J. Aumann. Survey of repeated games. In *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, pages 11–42. Bibliographisches Institut Mannheim/Wien/Zürich, 1981.
- [4] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [5] C. Baier, M. Größer, M. Leucker, B. Bollig, and F. Ciesinski. Probabilistic controller synthesis. In *Proceedings of the 3rd IFIP International Conference on Theoretical Computer Science, IFIP TCS 2004*, pages 493–506. Kluwer Academic Publishers, 2004.
- [6] R. E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [7] T. Bewley and E. Kohlberg. On stochastic games with stationary optimal strategies. *Mathematics of Operations Research*, 3(2):104–125, 1978.
- [8] H. Björklund and S. Vorobyov. Combinatorial structure and randomized subexponential algorithms for infinite games. *Theoretical Computer Science*, 349:347–360, 2005.
- [9] H. Björklund, S. Sandberg, and S. Vorobyov. A discrete subexponential algorithm for parity games. In *Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science, STACS 2003*, volume 2607 of LNCS, pages 663–674. Springer-Verlag, 2003.

- [10] T. Brázdil, V. Brožek, V. Forejt, and A. Kučera. Stochastic games with branching-time winning objectives. In *Proceedings of the 21st IEEE Symposium on Logic in Computer Science, LICS 2006*, pages 349–358. IEEE Computer Society Press, 2006.
- [11] S. R. Buss and L. Hay. On truth-table reducibility to SAT. *Information and Computation*, 91(1):86–102, 1991.
- [12] J. Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the 20th annual ACM Symposium on Theory of Computing, STOC '88*, pages 460–469. ACM Press, 1988.
- [13] K. Chatterjee. Stochastic Müller games are PSPACE-complete. In *Proceedings of the 27th International Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2007*, volume 4855 of LNCS, pages 436–448. Springer-Verlag, 2007.
- [14] K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Simple stochastic parity games. In *Proceedings of the 12th Annual Conference of the European Association for Computer Science Logic, CSL 2003*, volume 2803 of LNCS, pages 100–113. Springer-Verlag, 2003.
- [15] K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Quantitative stochastic parity games. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms, SODA 2004*, pages 121–130. ACM Press, 2004.
- [16] K. Chatterjee, R. Majumdar, and M. Jurdziński. On Nash equilibria in stochastic games. In *Proceedings of the 13th Annual Conference of the European Association for Computer Science Logic, CSL 2004*, volume 3210 of LNCS, pages 26–40. Springer-Verlag, 2004.
- [17] K. Chatterjee, L. de Alfaro, and T. A. Henzinger. The complexity of stochastic Rabin and Streett games. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming, ICALP 2005*, volume 3580 of LNCS, pages 878–890. Springer-Verlag, 2005.
- [18] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.
- [19] A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2): 203–224, 1992.
- [20] V. Conitzer and T. Sandholm. Complexity results about Nash equilibria. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI 2003*, pages 765–771. Morgan Kaufmann, 2003.
- [21] C. A. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.
- [22] C. A. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. *IEEE Transactions on Automatic Control*, 43(10):1399–1418, 1998.
- [23] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [24] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
- [25] L. de Alfaro. How to specify and verify the long-run average behavior of probabilistic systems. In *Proceedings of the 13th IEEE Symposium on Logic in Computer Science, LICS '98*, pages 454–465. IEEE Computer Society Press, 1998.
- [26] L. de Alfaro and T. A. Henzinger. Concurrent omega-regular games. In *Proceedings of the 15th IEEE Symposium on Logic in Computer Science, LICS 2000*, pages 141–154. IEEE Computer Society Press, 2000.

- [27] L. de Alfaro, T. A. Henzinger, and O. Kupferman. Concurrent reachability games. *Theoretical Computer Science*, 386(3):188–217, 2007.
- [28] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science, FoCS '91*, pages 368–377. IEEE Computer Society Press, 1991.
- [29] E. A. Emerson and C. S. Jutla. The complexity of tree automata and logics of programs. *SIAM Journal on Computing*, 29(1):132–158, 1999.
- [30] K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.
- [31] K. Etessami, M. Z. Kwiatkowska, M. Y. Vardi, and M. Yannakakis. Multi-objective model checking of Markov decision processes. *Logical Methods in Computer Science*, 4(4), 2008.
- [32] J. Filar and K. Vrieze. *Competitive Markov decision processes*. Springer-Verlag, 1997.
- [33] O. Friedmann. An exponential lower bound for the parity game strategy improvement algorithm as we know it. In *Proceedings of the 24th IEEE Symposium on Logic in Computer Science, LICS 2009*, pages 145–156. IEEE Computer Society Press, 2009.
- [34] M. R. Garey, R. L. Graham, and D. S. Johnson. Some NP-complete geometric problems. In *Proceedings of the 8th Annual ACM Symposium on Theory of Computing, STOC '76*, pages 10–22. ACM Press, 1976.
- [35] H. Gimbert and F. Horn. Solving simple stochastic games with few random vertices. *Logical Methods in Computer Science*, 5(2), 2009.
- [36] H. Gimbert and F. Horn. Solving simple stochastic tail games. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms, SODA 2010*, pages 847–862. ACM Press, 2010.
- [37] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games*, volume 2500 of LNCS. Springer-Verlag, 2002.
- [38] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.
- [39] L. A. Hemachandra. The strong exponential hierarchy collapses. *Journal of Computer and System Sciences*, 39(3):299–322, 1989.
- [40] F. Horn. *Random Games*. PhD thesis, Université Paris 7, 2008.
- [41] F. Horn. Explicit Muller games are PTIME. In *Proceedings of the 28th International Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2008*, volume 2 of *Leibniz International Proceedings in Informatics*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2008.
- [42] P. Hunter and A. Dawar. Complexity bounds for regular games. In *Proceedings of the 30th International Symposium on Mathematical Foundations of Computer Science, MFCS 2005*, volume 3618 of LNCS, pages 495–506. Springer-Verlag, 2005.
- [43] N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22(3):384–406, 1981.
- [44] M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, 1998.
- [45] M. Jurdziński, M. Paterson, and U. Zwick. A deterministic subexponential algorithm for solving parity games. *SIAM Journal on Computing*, 38(5):1519–1532, 2008.
- [46] N. Klarlund. Progress measures, immediate determinacy, and a subset construction for tree automata. *Annals of Pure and Applied Logic*, 69(2–3):243–268, 1994.

- [47] A. P. Maitra and W. D. Sudderth. Finitely additive stochastic games with Borel measurable payoffs. *International Journal of Game Theory*, 27(2):257–267, 1998.
- [48] D. A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- [49] A. McIver and C. Morgan. Games, probability and the quantitative μ -calculus. In *Proceedings of the 9th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, LPAR 2002*, volume 2514 of LNCS, pages 292–310. Springer-Verlag, 2002.
- [50] R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65(2):149–184, 1993.
- [51] A. W. Mostowski. Games with forbidden positions. Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.
- [52] J. F. Nash, Jr. Equilibrium points in N -person games. *Proceedings of the National Academy of Sciences of the USA*, 36:48–49, 1950.
- [53] A. Neyman and S. Sorin, editors. *Stochastic Games and Applications*, volume 570 of *NATO Science Series C*. Springer-Verlag, 2003.
- [54] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [55] L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the USA*, 39:1095–1100, 1953.
- [56] W. Thomas. Automata on infinite objects. In *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics, pages 133–192. Elsevier, 1990.
- [57] M. Ummels. Rational behaviour and strategy construction in infinite multiplayer games. Diploma Thesis, RWTH Aachen University, 2005.
- [58] M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *Proceedings of the 11th International Conference on Foundations of Software Science and Computation Structures, FOSSACS 2008*, volume 4962 of LNCS, pages 20–34. Springer-Verlag, 2008.
- [59] M. Ummels. *Stochastic Multiplayer Games: Theory and Algorithms*. PhD thesis, RWTH Aachen University, 2010.
- [60] M. Ummels and D. Wojtczak. Decision problems for Nash equilibria in stochastic games. In *Proceedings of the 18th Annual Conference of the European Association for Computer Science Logic, CSL '09*, volume 5771 of LNCS, pages 515–529. Springer-Verlag, 2009.
- [61] J. Vöge and M. Jurdziński. A discrete strategy improvement algorithm for solving parity games. In *Proceedings of the 12th International Conference on Computer Aided Verification, CAV 2000*, volume 1855 of LNCS, pages 202–215. Springer-Verlag, 2000.
- [62] K. W. Wagner. Bounded query classes. *SIAM Journal on Computing*, 19(5):833–846, 1990.
- [63] W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1–2):135–183, 1998.
- [64] W. Zielonka. Perfect-information stochastic parity games. In *Proceedings of the 7th International Conference on Foundations of Software Science and Computation Structures, FOSSACS 2004*, volume 2987 of LNCS, pages 499–513. Springer-Verlag, 2004.