

The Complexity of Pure Maxmin Strategies in Two-Player Extensive-Form Games

Junkang Li

NukkAI, Paris, France

*Université Caen Normandie, ENSICAEN, CNRS, Normandie Univ,
GREYC UMR6072, F-14000 Caen, France*

JUNKANG.LI@NUKK.AI

Bruno Zanuttini

*Université Caen Normandie, ENSICAEN, CNRS, Normandie Univ,
GREYC UMR6072, F-14000 Caen, France*

BRUNO.ZANUTTINI@UNICAEN.FR

Véronique Ventos

NukkAI, Paris, France

VVENTOS@NUKK.AI

Abstract

Extensive-form games model strategic interaction between players, with an emphasis on the sequential aspect of decision-making: players take turns to move until an ending is reached, and receive a reward according to which ending is reached. We study the complexity of computing the pure maxmin value for such games, i.e. the maximum reward that a player can guarantee by playing a pure strategy, whatever their opponents play. We focus on two-player and two-team games and perform a systematic study depending on the degree of imperfect information of each player or team: perfect information, perfect recall, or perfect recall for each agent in a team (which we call multi-agent perfect recall). For each combination, we settle the complexity of deciding whether the maxmin value is at least as high as a given threshold. We give a complete complexity picture for three orthogonal settings: games represented explicitly by their game tree; games represented compactly by game rules, for which we propose two new formalisms; games in which the set of strategies of the opponents is restricted to a known set of opponent models.

1. Introduction

Game theory is concerned with strategic interaction between players, typically in a competitive setting (Bonanno, 2018; Maschler et al., 2020). In this paper, we are interested in algorithmic and complexity-theoretic aspects of game theory, for games represented in extensive form, that is, as trees of possible sequences of moves by the players and of resulting payoffs. Extensive-form games (EFGs) provide a natural representation of games. Additionally, they put emphasis on the sequence of moves, and in particular on the information available to a player when they are to choose a move, which is the focus of this work.

We are interested in EFGs with imperfect information, where players do not necessarily observe or remember all the moves made by their opponents or themselves. This is captured by information sets, which contain game positions that a player cannot distinguish due to what they cannot observe or remember. For instance, at the beginning of a typical card game, a player cannot distinguish two positions in which they have the same hand while their opponents have different ones.

We focus on two-player zero-sum games, and we study the solution concept called *maxmin*, according to which each player attempts to maximise their minimal reward against any strategy of their opponent, which corresponds to a notion of *robustness*. Maxmin is especially relevant in

settings where nothing should be assumed about the opponent’s strategy: in such cases, playing any strategy but a maxmin one exposes the player to the risk of getting a lower payoff than the maxmin value. In particular, maxmin strategies are the counterpart for zero-sum games of Stackelberg strategies for general-sum games; Stackelberg strategies have been adequately applied to real-world situations, e.g. those modelled by security games, in which a defender must choose a strategy to protect resources against an attacker (Pita et al., 2008; Basilico et al., 2012).

In this setting, we study the complexity of computing robust *pure* strategies. A pure strategy of a player is a strategy that deterministically prescribes a move for each information set of this player. In general, such strategies guarantee a lower payoff than mixed or behaviour strategies, which can prescribe stochastic choices. However, pure strategies are important in many settings. For instance:

- Pure strategies are optimal against nondeterministic opponents, as in contingent planning (Rintanen, 2004).
- Pure strategies are particularly important for pedagogical purposes, for instance when teaching human players to play a game, as experiments have shown that human reason in terms of pure rather than mixed or behaviour strategies (Dhami, 2019, Chapter 1).
- Pure strategies are a realistic model of the defender’s strategies in some security games, where the defender’s policy is changed infrequently and hence appears to be pure to the attacker (Schlenker et al., 2018).
- Pure strategies are by nature more predictable, interpretable, and verifiable than mixed strategies; these properties are of prime importance in high-risk safety-critical systems or in systems with legal or regulatory constraints, where stochasticity and opacity of policies are undesirable.
- Deploying pure strategies does not require additional sources of random bits as required by mixed and behaviour strategies, especially in the case of correlated team maxmin equilibria (Basilico et al., 2017) where the coordination between multiple players of the same team may be necessary.
- When pure strategies correspond to resources (e.g. to be bought before implementing the strategy), then it may not be operational to randomise over pure strategies, as this may require to have all resources at one’s disposal instead of only one (Lipton et al., 2003).
- When pure strategies correspond to protocols to be executed by human agents in a repeated game, randomising over pure strategies requires them to be familiar with all the protocols involved; such tasks may be cognitively difficult and error-prone (McCarthy et al., 2018).

Noticeably, pure Nash Equilibria have already been studied (Gottlob et al., 2005; Yun et al., 2022).

We investigate the complexity of computing the pure maxmin value of a given two-player zero-sum EFG in various settings. We mostly study the following decision problem: given an EFG and a threshold, decide whether the pure maxmin value for a given player (or team) is at least the threshold (PURE MAXMIN). The complexity of PURE MAXMIN is established with respect to different types of games, characterised by the degree of imperfect information of the two players, and the presence or absence of chance factors. In our work, three degrees of imperfect information are considered: perfect information (a player observes and remembers every move made by their opponents or themselves), perfect recall (a player does not necessarily observe all the moves but never forgets

anything they did or knew in the past), and multi-agent perfect recall (a team of agents with perfect recall).¹ Apart from completeness results for various complexity classes, we also provide (or refer to) polynomial-time algorithms for decision problems belonging to the complexity class P.

The decision problem PURE MAXMIN is studied under three orthogonal settings:

- First, we consider the standard setting in which EFGs are explicitly represented by their game tree; the complexity is measured with respect to the size of the whole game tree. This setting has already been partially tackled in the literature, e.g. in early work by Koller and Megiddo (1992), Koller et al. (1996), von Stengel (1996) (for pure, mixed, and behaviour strategies) and recent work by von Stengel and Koller (1997), Basilico et al. (2017), Celli and Gatti (2018) (for behaviour and mixed strategies in two-team games).
- Then, we consider the setting in which EFGs are represented compactly, for which we introduce a very generic representation called *oracle games*. This setting captures more naturally tabletop games or video games played by humans: in such games, the game tree is usually described by succinct game rules and is generated online (when the game is played) rather than offline. Intuitively, the complexity in this setting is measured with respect to the size of the game rules.
- Finally, we consider the setting in which the opponent does not consider all strategies, but only a limited set of (behaviour) strategies, known as *opponent models* in the literature; we are looking for pure maxmin strategies only against these models. This allows modelling situations in which a player has prior information about the behaviour or reasoning process of their opponent.

Near the end, we also briefly discuss the complexity of deciding whether the pure maxmin value of a given EFG is at most (PURE \leq -MAXMIN), or equal to (PURE =-MAXMIN) a given threshold.

The main contribution of our work is the following: (i) we thoroughly study the complexity of finding pure maxmin strategies in two-player extensive-form games under the aforementioned settings; (ii) we assemble known results scattered in the literature into one place and fill the gaps among these results; (iii) we present all the results in a coherent manner with rigorous proofs and sufficient references; (iv) for existent results, we strengthen them (e.g. by restricting the number of different payoffs, players, or turns needed in the construction for hardness results) or simplify their proof when possible.

The paper is organised as follows. We first present related work (Section 2) and technical background about EFGs and the notion of information in such games (Section 3). Then we study PURE MAXMIN for EFGs defined explicitly by their game tree (Section 4), for compactly represented EFGs (Section 5), and for EFGs defined explicitly by their tree but with known opponent models (Section 6). PURE \leq -MAXMIN and PURE =-MAXMIN are then studied (Section 7) before we conclude (Section 8). The proofs of auxiliary lemmas are deferred from the main text to the appendix. Finally, the complexity results of each section are summarised in a separate table.

2. Related Work

Our work concerns non-cooperative games from the field of game theory, a rigorous treatment of which is given by Maschler et al. (2020). These games model non-cooperative multi-agent

1. The more general case in which an agent can forget some information, a setting called *absent-mindedness* in the literature (Piccione & Rubinstein, 1997), is *not* considered in our work.

interaction, and mainly come in two categories: normal-form games, for one-round interaction during which agents pick their action concurrently; extensive-form games (EFGs), for multi-round interaction during which agents take turns to make decisions. The reward at the end of a game is determined by the actions taken by every agent during the game; the goal of each agent is to maximise their own reward.

Solution Concepts for Games One of the central questions of game theory concerns what will happen in a given game. From a theoretical and analytical point, we are interested in predicting what players will do under the assumptions of certain notions of reasonable or rational behaviours from the players (Maschler et al., 2020, page 84). In single-player games, the focus is usually on identifying optimal strategies, i.e. those maximising the expected payoff of a player. However, in multi-player games, the notion of optimal strategies of a player is not uniquely defined, since the optimality of a strategy can depend on the choices of other players. In this case, game theorists resort to *solution concepts* to identify certain subsets of outcomes deemed interesting in one sense or another (Shoham & Leyton-Brown, 2009, page 60).

In general, a solution concept is a rule or a systematic way to specify, given any game, predictions about players' behaviour and outcomes of the game to be expected (Myerson, 1997, page 107). Since the foundational work by von Neumann and Morgenstern (1994), various solution concepts have been proposed and studied.

The most prominent solution concept, and the one that has received the most attention from the literature, is Nash equilibrium (NE). An NE is a strategy profile such that no one has a strict unilateral incentive to deviate. Hence, NE is a solution concept about stability. Many refinements of NE have also been proposed and studied in the literature; see the work by van Damme (1991) for a detailed exposition.

For other examples of solution concepts, one may also refer to the books by Maschler et al. (2020, Chapter 7) and Shoham and Leyton-Brown (2009, Sections 3.3-4) for a quick overview. Moreover, the field of epistemic game theory (Perea, 2012) studies solution concepts coming from different assumptions of rationality, e.g. rationalisability (Pearce, 1984). Behavioural game theory (Dhami, 2019) concerns solution concepts that take human behaviour into account.

Maxmin and Stackelberg Strategies We focus on the solution concept of *maxmin* for EFGs, especially *pure* maxmin, where strategies for MAX are restricted to be pure. Maxmin predicts/prescribes that every player aims to maximise the minimum reward they get against all possible strategies of their opponents, whence the name “maxmin”.

By definition, maxmin is a solution concept about safety and robustness, since playing a maxmin strategy guarantees a certain amount of reward for a player. Importantly, in a zero-sum game, the notions of Nash equilibrium and mixed maxmin coincide, in the sense that a mixed strategy profile is an NE if and only if every strategy in it is a mixed maxmin strategy. Hence, the notion of maxmin is particularly attractive for zero-sum games.

Another closely related solution concept is that of Stackelberg strategies proposed by von Stackelberg (1934). A Stackelberg strategy, also called an optimal strategy to commit to, is a strategy that maximises a player's expected payoff when the opponents always best-respond with respect to their own reward.² In particular, Stackelberg strategies reduce to maxmin strategies in zero-sum games. However, these two notions differ in general-sum games: Stackelberg strategies

2. A formal definition in our notation is provided near the end of Section 8.

assume that one’s opponents are self-interested and aim to maximise their own payoffs; maxmin assumes that one’s opponents strive to minimise one’s payoff.

Team Games We also study the complexity of maxmin in multi-agent settings, i.e. when MAX or MIN consists of multiple agents, who all have perfect recall but do not share observation or information. This setting, called (*two-team games*), is worth a special mention since until quite recently, most of the literature on complexity or algorithms for equilibrium or maxmin focused solely on games involving two agents.

Maxmin for team MAX is called *team maxmin equilibrium* (TME) in the literature, first proposed by von Stengel and Koller (1997). Later, Basilico et al. (2017) and Celli and Gatti (2018) propose another maxmin notion called TMECor (“Cor” stands for “correlation”), which allows agents in team MAX to have access to a correlation device in order to coordinate in their mixed strategies (in a similar vein to correlated equilibrium, a reference for which is given by Maschler et al., 2020, Chapter 8). These works also study the inefficiency gap between NE, TME, and TMECor.³ Seemingly, these works have stirred the community’s interest in team games; many algorithms for computing TMECor of EFGs have been designed in recent years. To cite a few: Farina et al. (2018), Zhang et al. (2021), Farina et al. (2021), Zhang and Sandholm (2022), Zhang et al. (2023).

Opponent Models Part of our work concerns the complexity of computing optimal strategies when it is assumed that the opponents’ strategies are taken from a known, restricted set. Such known strategies are referred to in the literature as opponent models. Behavioural game theory, for instance, studies models of human behaviour in games (Dhami, 2019).

Opponent models can come in diverse forms. Iida et al. (1993, 1994) propose opponent models for games with perfect information, where models are given by the evaluation function and the search depth of the opponent. Sturtevant et al. (2006) propose opponent models given by opponent’s preferences over the outcomes of a game. Rebstock et al. (2019) use opponent models learnt from human in card games with imperfect information. Albrecht and Stone (2018) provide a survey of approaches for opponent modelling. Our work is related to these in the sense that we assume opponent models to be given (called “type-based reasoning” by Albrecht & Stone, 2018, Section 4.2).

Complexity of Solving Games Most work in the literature on the computational complexity of games concerns finding Nash equilibria, especially in normal-form games (e.g. Gilboa & Zemel, 1989; Daskalakis et al., 2009, who show the PPAD-completeness of computing NE). For more references, one may consult the introduction by Conitzer and Sandholm (2008), who also show that it is NP-complete to decide whether NEs with certain natural properties exist.

Koller and Megiddo (1992), Koller et al. (1996), von Stengel (1996) made the first major steps towards understanding the complexity landscape of solving two-player zero-sum EFGs. Apart from studying the complexity of computing pure, behaviour, and mixed maxmin strategies for some classes of EFGs, they also give polynomial-time algorithms for computing behaviour maxmin strategies of EFGs with perfect recall, based on linear programming. Building on the work by Koller and Megiddo (1992), Gimbert et al. (2020) and Zhang et al. (2023) study the complexity of TME and TMECor, thereby yielding a relatively complete picture of the complexity of behaviour and mixed maxmin for two-team zero-sum EFGs. Finally, we also established the complexity of computing (pure/behaviour/mixed) maxmin strategies in a specific subclass of two-team EFGs, namely those

3. Allowing more communication between agents of team MAX increases their maxmin value.

with incomplete information but only public actions (hence all the uncertainty comes from the initial state) (Li et al., 2024).

As for Stackelberg strategies, their computational complexity has been studied for normal-form games and Bayesian games by Conitzer and Sandholm (2006) and for EFGs by Letchford and Conitzer (2010). Since Stackelberg strategies generalise maxmin strategies for general-sum multi-player games, the latter work is very similar to ours in flavour. However, their results are independent of ours; our study can be considered to be more refined results for the zero-sum two-team case.

Peterson et al. (2001) study the complexity of compactly represented EFGs. However, they define games as alternating Turing machines with different kinds of specifications (number of alternations, amount of space, private or public band, etc.). They also invent a formalism called dependency quantified Boolean formulae, generalising the well-known formalism of quantified Boolean formulae (which compactly represent two-player EFGs of no chance and with perfect information) to allow multiple agents of MAX with different information. Deciding the truth value of such a formula is NEXP-complete, in contrast with the PSPACE-completeness for QBF.

Complexity of Other Models for Decision-Making The complexity of other models for decision-making under uncertainty (Kochenderfer, 2015; Kochenderfer et al., 2022) has also been studied in parallel. In general, these results showcase a jump in computational complexity when we go from single-agent models to multi-agent ones, or from fully observable models to partially observable ones; we will also observe this phenomenon in the complexity of EFGs.

For the complexity (and computability) of automatic planning, see the work by Mundhenk et al. (2000) for the complexity of finite horizon MDP and POMDP (with different observability, stationary or time/history-dependent policies, short or long horizon); Madani et al. (2003) for the undecidability of infinite horizon POMDP; Bernstein et al. (2002) for the complexity of Dec-POMDP; Brafman et al. (2013) for the complexity of qualitative Dec-POMDP; Goldsmith and Mundhenk (2007) for the complexity of POSG; and Rintanen (2004) for the complexity of propositional planning with full/no/partial observability and deterministic/nondeterministic transitions.

For the complexity of graph games (Apt & Grädel, 2011) for different observability and objectives, see the surveys by Chatterjee and Henzinger (2012) and Chatterjee et al. (2013).

3. Background on Game Theory

We now introduce notions that we will use throughout this work. For more details on game theory, one can refer to the textbook by Maschler et al. (2020), on which most of this section is based.

3.1 Games and Strategies

We adopt the following notations. For a tree T , we write $L(T)$ for the set of its leaves. For an internal vertex v , we write $C(v)$ for the set of children/successors of v . For a finite set S , we write $|S|$ for its cardinality and $\Delta(S)$ for its simplex, i.e. the set of all probability distributions over S .

3.1.1 EXTENSIVE-FORM GAMES

Informally, an extensive-form game (EFG) is a game played on a tree starting from the root, during which players take turns to choose a child of the current node until a leaf is reached; the leaf determines the payoff for each player. To model the idea that a player cannot distinguish one situation from another, the concept of information set can be used:

Definition 3.1 (Information Set and Available Actions). *In a tree, an information set is a pair $\langle IS, A \rangle$, where IS is a set of internal vertices with the same number of children, and A (called the set of available actions at IS) is a partition of the children of all vertices in IS ,⁴ such that no pair of two different children of the same vertex is in the same set in A : $\forall a \in A, \forall v \in IS, |a \cap C(v)| = 1$.*

Intuitively, two nodes in the same information set are indistinguishable by a player; if a player picks $a \in A$ at IS , then the next node will be the unique node in $a \cap C(v)$, where $v \in IS$ is the current node. This forces a player to pick the same action for every vertex in the same information set.

Definition 3.2 (EFG of Chance). *A two-player zero-sum EFG of chance (with imperfect information) is a tuple $G = \left(T, (V_i)_{i \in \{0, +, -\}}, (p_v)_{v \in V_0}, u_+, (\langle IS_i^j, A_i^j \rangle)_{i \in \{+, -\}}^{j=1, \dots, k_i} \right)$ where $T = (V, E, r)$ is a finite game tree; $(V_i)_{i \in \{0, +, -\}}$ is a partition of $V \setminus L(T)$; for every $v \in V_0$, $p_v \in \Delta(C(v))$; $u_+ : L(T) \rightarrow \mathbb{R}$ is a utility function; and for each player $i \in \{+, -\}$, $(\langle IS_i^j, A_i^j \rangle)_{j=1, \dots, k_i}$ is a set of information sets of player i such that $(IS_i^j)_{j=1, \dots, k_i}$ forms a partition of V_i .*

The index 0 corresponds to *Nature*, who chooses the successor at every $v \in V_0$ (called *chance node*) according to the probability distribution p_v . If a game has no chance node (i.e. $V_0 = \emptyset$), it is called an *EFG of no chance*. Hence, by definition, games of no chance form a strict subclass of games of chance. In the following, unless explicitly specified, an EFG refers to an EFG of chance.

The players indexed by $+$ and $-$ are called *player MAX* and *player MIN*, respectively. u_+ is interpreted to be the utility function of MAX. The game is zero-sum in the sense that the two players have opposite payoffs at every leaf (hence for every outcome). In other words, zero-sum games are completely adversarial. Henceforth, we write $u_- = u_+$ for the utility function of MIN.

Remark. *In general, games can be general-sum, i.e. what a player gains is not necessarily equal to what another player loses. However, as we are only interested in the solution concept of maxmin, we may assume without loss of generality that all two-player games we consider are zero-sum (cf. Definition 3.6 and the discussion thereafter).*

For $i \in \{+, -\}$, V_i is the set of *decision nodes* of player i . We will refer to a set IS_i^j as an information set of player i , leaving implicit the set A_i^j of actions available at IS_i^j . We denote the set of all information sets of player i in an EFG by $\mathcal{IS}_i := \{IS_i^1, \dots, IS_i^{k_i}\}$. By definition, the union of the sets in \mathcal{IS}_i yields V_i .

Definition 3.3 (Boolean EFG). *A two-player zero-sum EFG is said to be Boolean if $u_+(l) \in \mathbb{B} = \{0, 1\}$ for all leaves l .*

A Boolean game is a zero-sum game with only Boolean payoffs for player MAX. 1 as payoff usually signifies a win for MAX (and a loss for MIN), and 0 a loss for MAX (and a win for MIN).

Remark. *The values 0 and 1 are inessential; one can choose any two different constants: what captures the essence of Boolean games is that there are only two possible payoff values for MAX.*

Definition 3.4 (Timeability). *An EFG is said to be timeable if every two nodes in the same information set have the same depth in the game tree.*

Not every EFG is timeable; see the work by Jakobsen et al. (2016), who introduce this notion and argue that timeability is a necessary condition for EFGs to be practically implementable in the real world. However, all hardness results in this paper will be proven by constructing only timeable EFGs to strengthen their practical implication.

4. A partition of a set S is a set of non-empty subsets of S such that the subsets are pairwise disjoint and their union is S .

3.1.2 STRATEGIES AND MAXMIN

A pure strategy of a player i in a game maps each information set of player i to one of their available actions at that information set. Formally:

Definition 3.5 (Pure Strategy). *A pure strategy of player $i \in \{+, -\}$ in a game G is a mapping $s_i : \mathcal{IS}_i \rightarrow \bigcup_j A_i^j$ such that $s_i(\mathcal{IS}_i^j) \in A_i^j$ for every information set $\mathcal{IS}_i^j \in \mathcal{IS}_i$. The set of all pure strategies of player i in the game G will be denoted by \mathcal{S}_i^P .*

Every pure strategy profile $(s_+, s_-) \in \mathcal{S}_+^P \times \mathcal{S}_-^P$ induces a probability distribution (due to Nature's drawings at chance nodes) over the leaves to be reached. In particular, every profile reaches a unique leaf in a game of no chance; the path from the root to this leaf is called the *payout* under the profile.

Let us denote by $p_{(s_+, s_-)}$ the distribution induced by the profile (s_+, s_-) . Then, the *expected utility/payoff/reward* for MAX when MAX plays s_+ and MIN plays s_- reads $\mathcal{U}_+(s_+, s_-) := \sum_{l \in L(T)} p_{(s_+, s_-)}(l) u_+(l)$. The solution concept of maxmin is then defined as follows.

Definition 3.6 (Pure Maxmin Value). *The pure maxmin value for MAX in a game G is defined to be*

$$\underline{v}_+ := \max_{s_+ \in \mathcal{S}_+^P} \min_{s_- \in \mathcal{S}_-^P} \mathcal{U}_+(s_+, s_-). \quad (1)$$

Maxmin is a security/robustness concept: \underline{v}_+ is the largest payoff player MAX can guarantee by playing a pure strategy. Notice that the maxmin value for MAX depends solely on \mathcal{U}_+ , hence only on u_+ , not on u_- . As a result, when studying the maxmin value of a two-player game, we can always assume without loss of generality that the game is zero-sum.

In a Boolean game, if the maxmin value for MAX is 1, it means MAX can force a win, no matter how Nature draws at chance nodes and how MIN plays; if the maxmin value for MAX is 0, then MIN can force a loss for MAX.

Definition 3.7 (Best response and maxmin strategy). *Let $s_+^* \in \mathcal{S}_+^P$ and $s_-^* \in \mathcal{S}_-^P$.*

- s_-^* is called a (pure) best response to s_+^* , if $\mathcal{U}_+(s_+^*, s_-^*) = \min_{s_- \in \mathcal{S}_-^P} \mathcal{U}_+(s_+^*, s_-)$.
- s_+^* is called an optimal pure strategy or a pure maxmin strategy against \mathcal{S}_-^P , if s_+^* achieves the pure maxmin value, i.e. $\min_{s_- \in \mathcal{S}_-^P} \mathcal{U}_+(s_+^*, s_-) = \underline{v}_+$.

A *mixed strategy* of player $i \in \{+, -\}$ is an element of $\Delta(\mathcal{S}_i^P)$, i.e. a probability distribution over the pure strategies of player i . Instead of mixing pure strategies of a player, we may also consider mixing their choice of actions at each of their information sets; this yields a new notion of strategy called *behaviour strategy*. Concretely, a behaviour strategy of player i maps each information set of player i to a probability distribution over their available actions at that information set. The set of mixed and behaviour strategies of player i will be denoted by \mathcal{S}_i^M and \mathcal{S}_i^B , respectively.

Although we will not detail this here, the notion of expected utility (hence also the notions of maxmin value, best response, maxmin strategy) is well-defined with respect to mixed or behaviour strategies. Formal definitions can be found in the book by Maschler et al. (2020, Chapter 6).

Remark. *In all the settings we will consider, we have (with abuse of the subset relation) $\mathcal{S}_i^P \subseteq \mathcal{S}_i^B \subseteq \mathcal{S}_i^M$ for every player i (see also the remark after Definition 3.10). In addition, by linearity of the expected utility with respect to MIN's mixed strategies, best responses can be taken to be pure*

ones instead of mixed ones. Therefore, it is without loss of generality that we consider only MIN's pure strategies in Definition 3.6. To put it in another way, the pure maxmin value will not change if we replace \mathcal{S}_-^P in (1) by \mathcal{S}_-^B or \mathcal{S}_-^M . This fact has also been shown by Isbell (1957), Piccione and Rubinstein (1997); the argument before is provided for our work to be self-contained.

Remark. When MAX and MIN can use mixed strategies, the minimax theorem holds:

$$\max_{\sigma_+ \in \mathcal{S}_+^M} \min_{\sigma_- \in \mathcal{S}_-^M} \mathcal{U}_+(\sigma_+, \sigma_-) = \min_{\sigma_- \in \mathcal{S}_-^M} \max_{\sigma_+ \in \mathcal{S}_+^M} \mathcal{U}_+(\sigma_+, \sigma_-).$$

Hence the role of MAX and MIN is symmetric. This symmetry is broken when MAX can only use pure strategies since, in general, we only have

$$\max_{s_+ \in \mathcal{S}_+^P} \min_{\sigma_- \in \mathcal{S}_-^M} \mathcal{U}_+(s_+, \sigma_-) \leq \min_{\sigma_- \in \mathcal{S}_-^M} \max_{s_+ \in \mathcal{S}_+^P} \mathcal{U}_+(s_+, \sigma_-),$$

and the inequality is strict, for instance, in the game Matching Pennies. More evidence of this broken symmetry is provided by the lack of symmetry in the complexity results we will see, e.g. Table 1.

3.2 Information in Games

We now present notions about information⁵ that are implicitly encoded by the structure of the information sets of the players in an EFG. In particular, we present the notions of perfect information, perfect recall, and multi-agent perfect recall (team games).

3.2.1 DEGREES OF IMPERFECT INFORMATION

Definition 3.8 (Perfect Information). *A player is said to have perfect information (PI) if all their information sets are singletons. An EFG is said to be with perfect information if all players have perfect information.*

A player $i \in \{+, -\}$ is said to have *perfect recall* (PR) if whenever two paths from the root arrive at the same information set of i , they pass through the same information sets of i , in the same order, and the same action is chosen by i at each such information set in these two paths. Formally:

Definition 3.9 (Perfect Recall). *A player $i \in \{+, -\}$ is said to have perfect recall (PR) if the following condition holds: Let v and v' be two nodes in the same information set of i , let $v_1, v_2, \dots, v_n = v$ (respectively $v'_1, v'_2, \dots, v'_{n'} = v'$) be the decision nodes of i passed through by the unique path from the root to v (respectively the path from the root to v'). Then $n = n'$, and for $k = 1, 2, \dots, n - 1$: (i) v_k and v'_k are in the same information set of i , denoted by IS_i^j for some j ; (ii) the child of v_k in the path to v and the child of v'_k in the path to v' are in the same available action at IS_i^j .*

An EFG is said to be with perfect recall if all players have perfect recall.

Intuitively, perfect recall implies that a player never forgets anything they knew in the past: two paths passing through different sequences of information sets cannot end in the same information set, hence if the player can distinguish two of their decision nodes v and v' , then they will not confound a descendant of a v with one of v' . In addition, the player remembers the actions they have chosen in the past, since taking different actions at the same information set will lead to different information sets. Hence, a player with perfect recall always remembers what they saw and did in the past.

5. By “information”, we informally refer to what actions in the past a player can observe or retain in memory.

Remark. *Perfect recall is the necessary and sufficient condition for all mixed strategies of a player to have an equivalent behaviour strategy (Kuhn, 1953).*

A player i is said to have *multi-agent perfect recall* (MA-PR) if none of their information sets is intersected twice by a path starting from the root. Formally:

Definition 3.10 (Multi-Agent Perfect Recall). *A player $i \in \{+, -\}$ is said to have multi-agent perfect recall (MA-PR) if for every two different nodes in the same information set of i , neither node is the ancestor of the other node. An EFG is said to be with multi-agent perfect recall or to be a team game if all players have multi-agent perfect recall.*

Notice that by definition, PI implies PR, which itself implies MA-PR.

Remark. *Multi-agent perfect recall is the necessary and sufficient condition for all behaviour strategies of a player to have an equivalent mixed strategy (Maschler et al., 2020, Theorem 6.11).*

Remark. *In the literature, the lack of multi-agent perfect recall is called absent-mindedness, or, less accurately, imperfect recall. We will not study this setting; one may refer to the work by Piccione and Rubinstein (1997) and other papers in the same collection for decision-making under this setting.*

The term *multi-agent* is motivated by the fact that if none of the information sets of a player is intersected twice by a path, then this player can be regarded as a team of multiple agents with perfect recall and a shared reward, each agent controlling one of the information sets of the player. Conversely, agents with perfect recall and identical payoffs can be regarded as being controlled by a meta-player (i.e. their team) who has multi-agent perfect recall.

This multi-agent interpretation dates back to the work by Isbell (1957), which leads to the notion of team games (von Stengel & Koller, 1997) and is widely adopted in recent research on these games (Basilico et al., 2017; Celli & Gatti, 2018, for instance). In the literature, a *team* is defined to be an inclusion-wise maximal set of players with perfect recall and the same utility function; and recent research focuses on the computational complexity or algorithms for solution concepts such as TME and TMECor in two-team games with perfect recall;⁶ see the references in Section 2 on this subject.

As the notion of two-player games with multi-agent perfect recall and the one of two-team games with perfect recall are equivalent, we use these two notions interchangeably throughout this work; we prefer using the first term so as to say that we focus exclusively on two-player games. From now on, to avoid potential confusion, a *player* means a team with zero, one, or more *agents* with perfect recall and a common utility function who make decisions in a completely decentralised manner. Since all agents in the same team have the same payoffs by definition, we only show the rewards for player/team MAX in all our examples of two-player zero-sum EFGs.

3.2.2 CONCURRENT ACTIONS IN TEAM GAMES

The notion of multi-agent perfect recall (or team) allows constructing EFGs that model concurrent actions, which are useful for proving complexity results. By *concurrent actions*, we mean situations in which each agent in the same team has to make decisions concurrently and independently, without knowing which action the others have taken. We use the word *concurrent* instead of *simultaneous* to emphasise that the chronological order of the actions of the agents is not relevant to the game; when modelling such situations with EFG, we can allow agents to take turns in any order.

6. In our terminology, TME (respectively, TMECor) corresponds to the maxmin value for player MAX with multi-agent perfect recall when MAX is allowed to use all behaviour strategies (respectively, all mixed strategies).




Figure 1: Cooperative Matching Pennies. Dotted lines represent (non-singleton) information sets.

Example. Consider the games called cooperative Matching pennies in Figure 1. In these games, we have a team MAX of 2 agents and a team MIN of 0 agent. Notice that on the left, we let agent 1 move first, but this is inessential; we may as well let agent 2 move first as on the right. These two games model the same situation, in which agent 1 and 2 have to pick between heads and tails concurrently and independently. Notice that in these two games, the only pure strategies of team MAX to guarantee a win (i.e. a payoff of 1), are to let both agents 1 and 2 choose heads, or to let them both choose tails.⁷

Remark. Notice how concurrent actions in the example above allow imposing non-adaptivity: both agents of MAX must stick to the same answer, otherwise they lose. This is essentially why multi-agent perfect recall allows encoding difficult (e.g. NP-hard) decision problems. This technique of forcing non-adaptivity is also used to prove that multi-prover interactive proofs are more powerful than single-prover ones (Babai et al., 1991), or that multi-agent planning is computationally more difficult than single-agent planning (Bernstein et al., 2002; Brafman et al., 2013). In our work, we will show similar results concerning team games with the same technique.

4. Complexity of EFGs

We first study the complexity of deciding whether the pure maxmin value of a given zero-sum EFG is above a given threshold, when the game tree of the EFG is *explicitly* given as input.

4.1 Problem Setting

Throughout this section, we consider the following decision problem:

Definition 4.1 (PURE MAXMIN). *Let \mathcal{G} be a class of zero-sum EFGs. Then PURE MAXMIN(\mathcal{G}) is the following decision problem.*

Input: An EFG $G \in \mathcal{G}$ and a rational number m .
 Output: Does $\underline{v}_+ := \max_{s_+ \in \mathcal{S}_+^P} \min_{s_- \in \mathcal{S}_-^P} \mathcal{U}_+(s_+, s_-) \geq m$ hold in G ?

For complexity analyses, to account for the explicit representation of the game, we define the size of an instance $\langle G, m \rangle$ of PURE MAXMIN to be $\|G\| + \|m\|$, where for a rational number m , we define $\|m\|$ to be the number of bits in the representation of m , and for an EFG of chance $G :=$

7. Also, notice that if the two agents have access to a correlation device so that they can coordinate in their mixed strategies, then they can implement a mixture of the two pure strategies above to guarantee a win; this is the problem setting for the solution concept of TMECor (Basilico et al., 2017; Celli & Gatti, 2018).

$(T, (V_i)_{i \in \{0,+, -\}}, (p_v)_{v \in V_0}, u_+, (\langle IS_i^j, A_i^j \rangle)_{i \in \{+, -\}}^{j=1, \dots, k_i})$ we define $\|G\| := |T| + \max_{l \in L(T)} (\|u_+(l)\|) + \max_{v \in V_0, v' \in C(v)} (\|p_v(v')\|)$.

We will study the complexity of PURE MAXMIN for the classes \mathcal{G} defined by three parameters:

- MAX's degree of imperfect information: perfect information (PI), perfect recall (PR), or multi-agent perfect recall (MA-PR);⁸
- MIN's degree of imperfect information: PI, PR, or MA-PR;
- the existence of chance nodes: games of no chance, or games of chance.

Notice that we are interested in maxmin values with respect to all MIN's pure strategies \mathcal{S}_-^P , which is equivalent to considering all behaviour or mixed strategies of MIN for all subclasses of games we consider: since MIN has at least multi-agent perfect recall, every behaviour strategy of MIN has an equivalent mixed strategy (Maschler et al., 2020, Theorem 6.11); in addition, best responses of MIN in mixed strategies can always be taken to be pure strategies, for the expected utility is linear with respect to MIN's mixed strategies.

Moreover, we can further simplify the analysis for games of no chance. Given an EFG G of no chance, let $\text{PI}_{\text{MIN}}(G)$ be the EFG of no chance obtained from G by replacing the set of information sets of MIN by the set of all singleton nodes. Then, in $\text{PI}_{\text{MIN}}(G)$, the game tree, the payoff functions, MAX's information sets, and the set of MAX's pure strategies are the same as in G , but MIN has perfect information.

Lemma 4.2 (Hansen et al., 2007, Lem. 1). *Let G be an EFG of no chance in which MIN has multi-agent perfect recall and $s_+ \in \mathcal{S}_+^P$ be a pure strategy of MAX in G . Then s_+ has the same payoff against MIN's best responses in G as against MIN's best responses in $\text{PI}_{\text{MIN}}(G)$.*

Proof. As stated before, we may restrict our attention to MIN's pure best responses. It suffices to show that against a fixed pure strategy s_+ of MAX, every pure best response of MIN in $\text{PI}_{\text{MIN}}(G)$ corresponds to at least one pure strategy of MIN in G with the same payoff.

Let s'_- be a pure best response of MIN to s_+ in $\text{PI}_{\text{MIN}}(G)$. The pure strategy profile (s_+, s'_-) in $\text{PI}_{\text{MIN}}(G)$ uniquely determines a playout of the game. Since MIN has MA-PR in G , this path intersects every information set of MIN in G at most once. Hence, one can define a pure strategy s_- of MIN in G such that at every decision node of MIN in the path, MIN takes the same action as in s'_- .⁹ By construction, the playout is the same under the pure strategy profile (s_+, s_-) in G as under (s_+, s'_-) in $\text{PI}_{\text{MIN}}(G)$, hence the payoff for MAX is the same. \square

Remark. *Hansen et al. (2007) only consider EFGs in which MIN has perfect recall, but their argument, which is reproduced here in our terminology, also carries over to the case in which MIN has only MA-PR.*

8. The case in which a player has *absent-mindedness* (i.e. no multi-agent perfect recall) is not in the scope of our work. For the computational complexity of EFGs with absent-mindedness, see the recent work by Gimbert et al. (2020), Tewolde et al. (2023, 2024).

9. Concretely, at an information set IS_- of MIN in G , if IS_- intersects exactly once the playout in $\text{PI}_{\text{MIN}}(G)$ under (s_+, s'_-) , then $s_-(IS_-)$ is defined to be the action picked by MIN under s'_- in $\text{PI}_{\text{MIN}}(G)$ at the unique intersection of IS_- and the playout; otherwise, if IS_- does not intersect the playout, then $s_-(IS_-)$ can be defined to be any available action at IS_- .

		No chance	Chance		
		PI/PR/MA-PR	PI	PR	MA-PR
MAX \ MIN	PI	P	P [m: 4.7]	NP-c [h: 4.8]	Σ_2^P -c [h: 4.10]
	PR	P [m: 4.5]	NP-c [h: 4.9]	NP-c	Σ_2^P -c
	MA-PR	NP-c [h: 4.6]	NP-c	NP-c [m: 4.4]	Σ_2^P -c [m: 4.12]

Table 1: Complexity of PURE MAXMIN. All hardness results hold even under the restriction to Boolean timeable EFGs with at most 2 agents for both MAX and MIN. PI, PR, MA-PR stand for perfect information, perfect recall, multi-agent perfect recall, respectively. Only key membership (“m”) and hardness (“h”) results are referred to; the others can be deduced by monotonicity. Results in bold are new from this paper; other results are direct consequences of known results (see the citations in the referred statements).

The intuition behind this result is that when MAX plays a pure strategy in a game of no chance, MIN essentially faces a deterministic decision problem, hence MIN’s degree of imperfect information is not relevant (as long as it is no worse than MA-PR) when computing MIN’s best response. This result no longer holds if MAX can use mixed or behaviour strategies. For example, in the game Matching Pennies, MAX’s uniform strategy has a lower payoff if MIN has perfect information.

Corollary 4.3. *The pure maxmin values for MAX in G and in $PI_{\text{MIN}}(G)$ are the same.*

Since $PI_{\text{MIN}}(G)$ can clearly be built in polynomial time from G , it follows that for games of no chance, for a fixed degree of imperfect information of MAX, the degree of imperfect information of MIN (PI, PR, or MA-PR) does not influence the complexity of PURE MAXMIN.

4.2 Summary of Results

The complexity of PURE MAXMIN(\mathcal{G}) is summarised in Table 1. By definition, the complexity of each case is increasingly monotone in all three parameters: MAX’s degree of imperfect information (in this order: PI, PR, MA-PR), MIN’s degree of imperfect information, and the existence of chance nodes (in this order: no chance, chance). Hence, Table 1 only gives the references for the key hardness (“h”) and membership (“m”) results; the other results can be deduced using monotonicity. Note that results written in bold font are new from our work. The relevant citations for the other results, which can be directly deduced from the literature, will be given with their statements.

Notice that the table is asymmetric with respect to MAX and MIN: this is evidence of the fact that the minimax theorem no longer holds when MAX can only play pure strategies.¹⁰

4.3 EFGs of no Chance

We start with an NP upper bound for most cases in Table 1, which follows directly from the literature.

Proposition 4.4 (Koller & Megiddo, 1992). *PURE MAXMIN is in NP for EFGs of chance in which MAX has MA-PR and MIN has PR.*

10. It is interesting to compare this table to Table 5, which shows the complexity of MIXED MAXMIN and where the symmetry between MAX and MIN holds.

Proof. One can guess a pure strategy s_+ for MAX, then verify that it yields an expected payoff no less than the given threshold, by computing a best response to s_+ for MIN, which can be done in linear time when MIN has PR (Koller & Megiddo, 1992, Proposition 2.7). \square

When MAX has perfect recall, the problem is actually in P.

Proposition 4.5 (Hansen et al., 2007). *PURE MAXMIN is decidable in linear time, and a fortiori is in P, for EFGs of no chance in which MAX has PR and MIN has MA-PR.*

Proof. This follows directly from applying Corollary 4.3 to a result by Hansen et al. (2007, Theorem 1). The linear-time algorithm for computing the pure maxmin value is given below; it will also be used to prove Proposition 5.7.

Given a two-player game of no chance in which MAX has perfect recall and MIN has multi-agent perfect recall, we recursively define the function $ev: V \rightarrow \mathbb{R}$ on all vertices of the game tree:

- If v is a leaf, then the value of this node is given by the utility function of MAX: $ev(v) := u_+(v)$.
- If v is a decision node of MIN, then the value of this node is the minimum of the value of its children: $ev(v) := \min_{v' \in C(v)} ev(v')$.
- Otherwise, v is a decision node of MAX. Let IS_+ be the information set containing v . Then the value of this node is

$$ev(v) := \max_{a \in A_+} \min_{v' \in IS_+} ev(a(v')), \quad (2)$$

where $a(v') := a \cap C(v')$ denotes the vertex reached by taking action a at node v' .

The function ev is well-defined since MAX's information sets form a forest due to MAX's perfect recall (Koller & Megiddo, 1992, Proposition 3.1). In addition, ev has the same value for vertices in the same information set of MAX: the range of the operators max and min in (2) both depend solely on the information set containing the vertex given as input to ev .

It can be shown that $ev(r)$ is the pure maxmin value of the game (Hansen et al., 2007, Lemma 2; Theorem 1). Therefore, a linear-time algorithm for computing the pure maxmin value is given by an algorithm that computes $ev(r)$ using the recursive definition of ev . \square

If MAX only has MA-PR, then the problem becomes NP-hard.

Proposition 4.6 (von Stengel & Forges, 2008, Thm. 1.3). *PURE MAXMIN is NP-hard for EFGs of no chance in which MAX has MA-PR. The result holds even under the restriction to 2 agents for MAX and to Boolean timeable games.*

Proof. The proof is a straightforward adaptation of the proof of the same result with chance nodes by von Stengel and Forges (2008, Theorem 1.3). Given a 3-CNF formula with n clauses, their reduction builds a game in which the maxmin value is 1 if the formula is satisfiable, and at most $1 - 1/n$ otherwise: first, Nature chooses a clause (uniformly at random), then a first agent of MAX, who observes the clause, chooses a literal in it, and finally a second agent of MAX, who observes the variable but not the clause nor the polarity of the literal, chooses a value for the variable. The payoff for MAX is 1 if and only if the assignment to the variable satisfies the literal (and hence the clause). The intuition is that if (and only if) the formula is not satisfiable, then for every strategy of MAX, there is at least one clause not satisfied, which is picked by Nature with probability $1/n$.

The adaptation to our statement simply consists in replacing Nature by MIN: if (and only if) the formula is not satisfiable, then for every strategy of MAX, the best response of MIN consists in choosing an unsatisfied clause, yielding a payoff of 0 for MAX. \square

Remark. We also briefly present an alternative reduction from 3-COLOURING to highlight the idea that multi-agent coordination makes it hard to play a game optimally (see also Subsubsection 3.2.2). Given a graph, we construct an EFG of no chance: MIN picks a pair of vertices (not necessarily different); each agent of MAX is shown a different vertex from the pair; then concurrently and independently, each picks a colour for the vertex they observe. MAX wins if and only if (1) the two agents pick the same colour if they are shown the same vertex, and (2) they pick different colours if the pair of vertices are connected by an edge. It is straightforward to verify that MAX can guarantee a win in this game if and only if the given graph is 3-colourable.

Remark. The contrast between the NP-hardness in Proposition 4.6 (EFGs of no chance; MAX has MA-PR and MIN has PI) and the membership in P in Proposition 4.5 (EFGs of no chance; MAX has PI and MIN has MA-PR) demonstrates the asymmetry between MAX and MIN in the computation of maxmin values when MAX can only use pure strategies. This is in contrast to the symmetry between MAX and MIN implied by the minimax theorem when MAX can use mixed strategies.

4.4 EFGs of Chance

For EFGs of no chance and with perfect information for both players, the famous minimax algorithm and alpha-beta search (Knuth & Moore, 1975) compute the maxmin value in linear time; this result has also been generalized to games with chance nodes.

Proposition 4.7 (Ballard, 1983). PURE MAXMIN is decidable in linear time, and a fortiori is in P, for EFGs of chance in which both MAX and MIN have PI.

However, PURE MAXMIN becomes hard when at least one player has imperfect information.¹¹

Proposition 4.8. PURE MAXMIN is NP-hard for EFGs of chance in which MAX has PI and MIN has PR. The result holds even under the restriction to Boolean timeable games.

Proof. We give a reduction from the NP-complete problem SUBSET SUM, which is defined as follows:

Input: A multi-set of natural numbers $S = \{i_1, \dots, i_n\}$, a natural number k .

Output: Is there a subset $J \subseteq S$ that sums up to k (i.e. $k = \sum_{j \in J} j$).

Let $S = \{i_1, \dots, i_n\}$ and k form an instance of SUBSET SUM. We build a game in which, intuitively, a strategy of MAX is a subset J of S , and MIN chooses to verify either $k \geq \sum_{j \in J} j$ or $k \leq \sum_{j \in J} j$.

Concretely, consider the following game:

- *Players:* Nature; MAX with perfect information; MIN with perfect recall.
- *Game tree:* At the root, Nature chooses uniformly at random an element $j \in S$. MAX observes j and chooses between \checkmark (encoding the choice of some $J \ni j$) or \times (encoding $J \not\ni j$). Finally, without observing j nor the choice of MAX, MIN chooses \leq or \geq .

¹¹ In contrast, computing optimal behaviour/mixed strategies for MAX is still in P when both players have perfect recall (Koller & Megiddo, 1992, Section 3).

- *Payoffs:* For each Nature's choice $j \in S$, MAX's payoff is as follows:
 - If MIN has chosen \geq , MAX receives nj if they have chosen \checkmark , otherwise 0.
 - If MIN has chosen \leq , MAX receives $2k - nj$ if they have chosen \checkmark , otherwise $2k$.
- *Threshold:* The threshold of maxmin value is k .

The construction is polynomial-time in the input (S, k) . Indeed, the game tree is of size $O(|S|)$. In addition, the construction yields a timeable EFG of chance in which MAX has perfect information and MIN has perfect recall.

Observe that the pure strategies of MAX are in bijection with the subsets of S . For each subset $J \subseteq S$, if MAX plays the pure strategy corresponding to J via choosing \checkmark (respectively \times) for $j \in J$ (respectively $j \notin J$), then MAX gets an expected payoff of $\sum_{j \in J} (\frac{1}{n} \times nj) + \sum_{j \notin J} (\frac{1}{n} \times 0) = \sum_{j \in J} j$ if MIN chooses to play \geq , and $\sum_{j \in J} (\frac{1}{n} \times (2k - nj)) + \sum_{j \notin J} (\frac{1}{n} \times 2k) = 2k - \sum_{j \in J} j$ if MIN chooses to play \leq . Hence, the maxmin value is

$$\max_{J \subseteq S} \min \left(\sum_{j \in J} j, 2k - \sum_{j \in J} j \right) \leq \max_{J \subseteq S} k = k,$$

with equality if and only if $\sum_{j \in J} j = k$ for some $J \subseteq S$.

Therefore, the maxmin value is k if and only if there is a subset J of S that sums up to exactly k ; if no such subset exists, the maxmin value is at most $k - 1$. This concludes our proof of NP-hardness. Finally, NP-hardness also holds for Boolean games because one can use Lemma A.1 to compile the game above into a Boolean one. \square

Proposition 4.9. *PURE MAXMIN is NP-hard for EFGs of chance in which MAX has PR and MIN has PI. The result holds even under the restriction to Boolean timeable games.*

Proof. This follows directly from the reduction in the proof of Proposition 4.8 by changing the order of the players: MIN first; then Nature; finally (observing Nature's but not MIN's choice) MAX. Now the game obtained is a timeable EFG of chance in which MAX has PR and MIN has PI. \square

Remark. *This also follows from a similar result proven by Frank and Basin (2001, Section 6) for EFGs with one-sided incomplete information, by a reduction from the NP-complete problem CLIQUE.*

We now consider games of chance in which MIN only has multi-agent perfect recall. Koller and Megiddo (1992, Proposition 2.10) show that if MAX also only has multi-agent perfect recall, then PURE MAXMIN is Σ_2^P -hard. We use a different reduction to strengthen their hardness result and show that actually PURE MAXMIN remains Σ_2^P -hard even if MAX has perfect information.

Proposition 4.10. *PURE MAXMIN is Σ_2^P -hard for EFGs of chance in which MAX has PI and MIN has MA-PR. The result holds even under the restriction to 2 agents for MIN and to Boolean timeable games.*

For this proof, we will consider a reduction from a variant of the domino tiling problem. A *tile* is a square with sides coloured, formally an element $t = (c_t, c_r, c_b, c_l) \in C^4$ for some set of colours C . We refer to c_t (respectively c_r, c_b, c_l) by $t(t)$ (respectively $r(t), b(t), l(t)$), standing for “top” (respectively “right”, “bottom”, “left”) of t .

Definition 4.11 (Legal Tiling). *Let C be a set of colours, including a distinguished element $w \in C$, let $T \subseteq C^4$ be a set of tiles, and let $m \geq 2$ be an integer. A tiling of $S \subseteq S_m := \{1, \dots, m\} \times \{1, \dots, m\}$ is a mapping $\tau: S \rightarrow T$. Such a tiling is said to be legal if the colours of adjacent tiles match, and the colour on the sides of S_m is always w :¹²*

$$\begin{aligned} \forall (r, c) \in S: (r, c+1) \in S &\implies r(\tau(r, c)) = l(\tau(r, c+1)); \\ \forall (r, c) \in S: (r+1, c) \in S &\implies t(\tau(r, c)) = b(\tau(r+1, c)); \\ \forall r \in \{1, \dots, m\}: (r, 1) \in S &\implies l(\tau(r, 1)) = w \wedge (r, m) \in S \implies r(\tau(r, m)) = w; \\ \forall c \in \{1, \dots, m\}: (1, c) \in S &\implies b(\tau(1, c)) = w \wedge (m, c) \in S \implies t(\tau(m, c)) = w. \end{aligned}$$

The decision problem TILING consists in, given $(C, w, T, 1^m)$,¹³ deciding whether there is a legal tiling of S_m . TILING is known to be NP-complete (van Emde Boas, 1997). Variants of tiling problems provide complete problems for many complexity classes (Schwarzentruber, 2019). For example, the variant FINITE TILING EXTENSION is Σ_2^P -complete (Schaefer & Umans, 2002):

- Input:* A finite set C with a distinguished element $w \in C$, a set of tiles $T \subseteq C^4$, and a natural number m expressed in unary.
- Output:* Is there a *non-extendable* legal tiling of the first row of S_m , that is, a legal tiling of $S_{1,m} := \{1\} \times \{1, \dots, m\}$ that cannot be extended to a legal tiling of S_m ?

Proof of Proposition 4.10. Given an instance $(C, w, T, 1^m)$ of FINITE TILING EXTENSION, we build a game in which MAX wins if they can choose a legal tiling of $S_{1,m}$ such that whatever tiling τ of the whole board S_m MIN chooses, either τ is not legal, or it does not extend MAX's tiling.

Concretely, we build the following game.

- *Players:* Nature; MAX with perfect information; MIN with multi-agent perfect recall of 2 agents, labelled by 1 and 2.
- *Game tree:* The game begins with a chance node and proceeds as follows:
 1. At the root, Nature chooses uniformly at random a column $c \in \{1, \dots, m\}$, and only shows it to MAX.
 2. MAX chooses a tile $t \in T$.
 3. Nature then chooses uniformly at random two positions $(r_1, c_1), (r_2, c_2) \in S_m$, then concurrently and independently shows agent i of MIN the position (r_i, c_i) .
 4. Without observing c nor t , each agent of MIN concurrently and independently chooses a tile $(t_1, t_2 \in T)$, and the game ends.
- *Payoffs:* The payoff for MAX at the leaf induced by $c, t, r_1, c_1, t_1, r_2, c_2, t_2$ is defined to be:
 1. $2m^5$, if $(r_1, c_1) = (r_2, c_2)$ but $t_1 \neq t_2$ (MAX wins if agent 1 and agent 2 of MIN are inconsistent);
 2. otherwise, $2m^4$, if $(r_1, c_1) = (1, c)$ but $t_1 \neq t$ (MAX wins if MIN's tiling does not extend MAX's);¹⁴

12. We number rows from bottom to top, and columns from left to right.

13. 1^m is the unary expression of m .

14. Due to the first condition, the two agents of MIN must choose the same tile t' for $(1, c)$ in any of their optimal strategies, hence this second condition checks whether $t' = t$ by looking only at the action of MIN 1.

3. otherwise, $-m^4$, if $r_1 = r_2 = 1$ and $c_2 = c_1 + 1$ but $r(t_1) \neq l(t_2)$ (MAX loses if their tiling is illegal with respect to adjacency, which is verified with MIN's choices);
4. otherwise, $-m$, if $c = 1$ and $l(t) \neq w$, or $c = m$ and $r(t) \neq w$, or $b(t) \neq w$ (MAX loses if their tiling is illegal with respect to colour w);
5. otherwise, 1, if $r_1 = r_2 \neq 1$ and $c_2 = c_1 + 1$ but $r(t_1) \neq l(t_2)$ (MAX wins if MIN's tiling is illegal with respect to horizontal adjacency on a row different from 1);
6. otherwise, 1, if $c_1 = c_2$ and $r_2 = r_1 + 1$ but $t(t_1) \neq b(t_2)$ (MAX wins if MIN's tiling is illegal with respect to vertical adjacency);
7. otherwise, 1, if $r_1 \neq 1$, $c_1 = 1$ and $l(t_1) \neq w$, or $r_1 \neq 1$, $c_1 = m$ and $r(t_1) \neq w$, or $r_1 = m$ and $t(t_1) \neq w$ (MAX wins if MIN's tiling is illegal on (r_1, c_1) with respect to colour w , on a row different from 1);
8. otherwise, 0.

- *Threshold:* The threshold of maxmin value is $1/m^4$.

The construction is polynomial-time in the input $(C, w, T, 1^m)$. Indeed, the game tree is of size $\mathcal{O}(m^5|T|^3)$, and the computation of the payoffs of the leaves is polynomial-time. In addition, the construction yields a timeable EFG of chance in which MAX has perfect information and MIN has multi-agent perfect recall of 2 agents.

We will show that the pure maxmin value of this game is at least $1/m^4$ if there is a non-extendable legal tiling of $S_{1,m}$ (and at most 0 otherwise). First, observe that the pure strategies of MAX are in bijection with the tilings of $S_{1,m}$. Similarly, the pure strategies of each agent of MIN are in bijection with the tilings of S_m .

\implies Suppose first that there is a non-extendable legal tiling τ_+ of $S_{1,m}$. Let (τ_1, τ_2) denote an arbitrary pure strategy of MIN, where τ_1 and τ_2 are tilings of S_m . We show in the following that τ_+ yields an expected payoff of at least $1/m^4$ against (τ_1, τ_2) .

First observe that, since τ_+ is a legal tiling of $S_{1,m}$, MAX gets a strictly negative payoff only under the 3rd condition in the definition of the payoff, which happens with probability at most $1/m^3$. Hence, MAX cannot get an expected payoff smaller than $(1/m^3) \times (-m^4) = -m$ from the corresponding leaves.

- If $\tau_1 \neq \tau_2$ holds, then with probability at least $(1/m^2)^2 = 1/m^4$, Nature chooses $(r_1, c_1) = (r_2, c_2)$ over which they differ, and MAX gets $2m^5$ (the 1st condition of payoff). Since the smallest expected payoff MAX can get from negative leaves is $-m$, the expected payoff of MAX against (τ_1, τ_2) is at least $1/m^4 \times 2m^5 - m > 1/m^4$.
- Now assume $\tau_1 = \tau_2$, but τ_1 is not an extension of τ_+ . Then, with probability at least $1/m^3$, Nature chooses $c, r_1 = 1$ and $c_1 = c$ such that τ_1 differs from τ_+ on $(1, c)$, and MAX wins $2m^4$ (the 2nd condition of payoff). As above, it follows that the expected payoff of MAX against (τ_1, τ_2) is at least $1/m^3 \times 2m^4 - m > 1/m^4$.
- Finally, assume that $\tau_1 = \tau_2$ holds and τ_1 is an extension of τ_+ . Then τ_1 is illegal by definition of τ_+ . Since τ_1 extends τ_+ , which is legal, the third condition of the payoff is never satisfied; MAX cannot get a negative payoff at any leaf. If τ_1 is illegal with respect to adjacency, with probability at least $1/m^4$, Nature chooses (r_1, c_1) and (r_2, c_2) to witness the illegality, and

MAX gets 1 (the 5th or 6th condition of payoff). It follows again that the expected payoff of MAX against (τ_1, τ_2) is at least $1/m^4$. Otherwise, τ_1 is illegal with respect to colour w , which is witnessed by Nature's choice (r_1, c_1) with probability at least $1/m^2$ and MAX gets 1 (the 7th condition of payoff); so again MAX has expected payoff of at least $1/m^2 > 1/m^4$.

Hence, playing τ_+ guarantees a payoff of at least $1/m^4$ for MAX, which means the maxmin value of the game is at least $1/m^4$.

\Leftarrow Conversely, suppose that the maxmin value of the game is at least $1/m^4$, and let τ_+ be a pure strategy of MAX achieving this value.

- We claim that τ_+ is legal. Indeed, otherwise MIN can play the strategy (τ, τ) , where $\tau: S_m \rightarrow T$ is an arbitrary extension of τ_+ . Under this profile, MAX cannot get a payoff of $2m^5$ nor $2m^4$, so they get at most 1 from every leave. Then if τ_+ is illegal with respect to adjacency, with probability at least $1/m^4$ Nature chooses $r_1 = r_2 = 1$ and $c_1, c_2 = c_1 + 1$ witnessing the illegality of τ on the first row, resulting in a negative payoff of $-m^4$ for MAX (the 3rd condition of payoff). It follows that MAX's expected payoff is at most $1/m^4 \times (-m^4) + (1 - 1/m^4) \times 1 \leq 0$, contradicting the assumption. Now if τ_+ is illegal with respect to colour w , with probability at least $1/m$ Nature chooses c witnessing this, yielding $-m$ to MAX (the 4th condition of payoff), and again MAX's expected payoff is at most $1/m \times (-m) + (1 - 1/m) \times 1 \leq 0$.
- We also claim that any extension $\tau: S_m \rightarrow T$ of τ_+ is illegal. Indeed, otherwise MIN can play the strategy (τ, τ) , where τ is a legal extension of τ_+ , which yields a payoff of 0 to MAX at all leaves, again contradicting the assumption.

Hence, τ_+ is a non-extendable legal tiling of the first row. In addition, by the argument above, we see that if no such τ_+ exists, then MAX gets at most 0 against MIN's best responses.

To show that this result still holds under the restriction to Boolean games, we can use the gadgets in Lemma A.1 to compile all integers payoffs in the construction above into Boolean ones. \square

We conclude this section by recalling the upper bound for the most general setting.

Proposition 4.12 (Koller & Megiddo, 1992, Prop. 2.10). *PURE MAXMIN is in Σ_2^P for EFGs of chance in which MAX and MIN have MA-PR.*

5. Complexity of Compactly Represented Games

In this section, we study the complexity of PURE MAXMIN for EFGs represented in compact form. As it turns out, the complexity of PURE MAXMIN is very robust to the exact compact representation chosen. Hence, we formulate hardness results for a very restricted class of representations, and membership results for a very general one. We first introduce these representations and show that they encompass natural representations, then we give the complexity results, which, as it happens, parallel the results for non-compact EFGs.

5.1 Compact Representations of Games

For most tabletop and video games, the game tree is rarely defined explicitly, but rather implicitly by the game rules; such rules allow computing the decision maker, the children, the payoffs, etc., at

a given node of the tree, which can therefore be generated online. In other words, we may say that the trees of these games are represented compactly by their game rules.

Typically, for such a compact representation of a game, the corresponding game tree is exponentially larger, or more. In the following, we introduce two formalisms to capture the intuition of what *compactly represented games* means.

5.1.1 COMPACT BOOLEAN GAMES

We first introduce a minimal formalism for compactly represented games, which will be used to prove all hardness results in this section. Our formalism is inspired by *quantified Boolean formulae* (QBF). A QBF is a formula of the form $Q_1 x_1 Q_2 x_2 \cdots Q_n x_n \varphi(x_1, x_2, \dots, x_n)$, where for all i , $Q_i \in \{\forall, \exists\}$ and x_i is a Boolean variable (with values in $\mathbb{B} = \{0, 1\}$), and φ is a plain (unquantified) Boolean formula. A QBF can be regarded as a game of no chance with perfect information in which MAX and MIN take turns to pick a value for each variable: MAX is responsible for variables with the existential quantifier \exists and MIN for the universal quantifier \forall , and MAX's goal is to render the formula φ true.

Peterson et al. (2001) propose a generalisation of QBF called *dependency quantified Boolean formula* (DQBF), which allows dependencies for existential variables. For example, in the DQBF $\forall x_1 \forall x_2 \exists y_1(x_1) \exists y_2(x_2) \varphi(x_1, y_1, x_2, y_2)$, the value of y_1 can only depend on the value of x_1 , and the value of y_2 only on x_2 . In general, a DQBF represents a game of no chance in which MIN has perfect information but MAX only has multi-agent perfect recall. In the example above, y_1 and y_2 can be regarded as being controlled by two agents of MAX. Notice that DQBF is indeed a generalisation of QBF, in the sense that every QBF of size n can be transformed into an equivalent DQBF of size $O(n^2)$ by allowing each existential variable to depend on all previous variables.

As we intend to model games of chance with multi-agent perfect recall for both MAX and MIN, we introduce a generalisation of DQBF that we call *compact Boolean game* (CBG), which can be written into the following form to mimic DQBF:

$$P_1 x_1(D_1) P_2 x_2(D_2) \cdots P_n x_n(D_n) \varphi_+(x_1, x_2, \dots, x_n),$$

where P_j is the owner of x_j , D_j is the set of variables on which x_j depends, and φ_+ is an unquantified Boolean circuit. More formally:

Definition 5.1 (Compact Boolean Game). *A compact Boolean game (CBG) is a tuple of the form $\gamma = \langle X, P, D, \varphi_+ \rangle$, with X an ordered list of variables, $P: X \rightarrow \{0, +, -\}$ an owner function, $D: X \rightarrow \mathcal{P}(X)$ a dependency function such that for all $x \in X$, $D(x)$ only contains variables that precede x in the ordering and $D(x) = \emptyset$ if $P(x) = 0$, and φ_+ a Boolean circuit with inputs in X and one Boolean output.¹⁵ A CBG is said to be of no chance, if it does not have chance variables: $\forall x \in X, P(x) \neq 0$.*

Intuitively, a CBG represents a game in which for each x_j in turn, player $P(x_j)$ chooses a Boolean value after observing only the values chosen for the variables in $D(x_j)$ (and Nature chooses uniformly at random). After all variables have been played, MAX receives the payoff given by the Boolean output of the circuit $\varphi_+(x_1, x_2, \dots, x_n)$ (and MIN receives the opposite).

Therefore, every CBG induces a Boolean EFG of chance. It is straightforward to define the EFG of chance induced by a CBG, with each node in the game tree given by a variable index j and an

¹⁵. $\mathcal{P}(X)$ is the powerset of X , i.e. the set of all subsets of X .

assignment to x_1, \dots, x_{j-1} , and with two nodes in the same information set if and only if they have the same index j and the same value for all variables in $D(x_j)$. Moreover, the EFG thus defined is timeable: only nodes at the same depth can be in the same information set. In particular, it implies that not all two-player EFGs (even those without absent-mindedness) can be represented by CBGs.¹⁶

It is clear that the players in the EFG defined by a CBG have multi-agent perfect recall and:

- a player $i \in \{+, -\}$ has perfect information if and only if $D(x_j) = \{x_1, \dots, x_{j-1}\}$ for all $x_j \in X$ such that $P(x_j) = i$ (which intuitively means player i observes all values chosen in the past);
- a player $i \in \{+, -\}$ has perfect recall if and only if for all $x_j, x_{j'} \in X$, $j < j'$ and $P(x_j) = P(x_{j'}) = i$ implies $D(x_j) \subseteq D(x_{j'})$ (i.e. player i never forgets any value they have observed before) and $x_j \in D(x_{j'})$ (i.e. player i never forgets any value they have chosen before).

We will prove hardness results for CBGs with an arbitrary circuit φ_+ . However, these hardness results even hold for CBGs the circuit φ_+ of which is restricted to certain languages (CNF, DNF, ROBDD),¹⁷ since we can compile φ_+ into these languages in polynomial time using the Tseitin transformation (Tseitin, 1983) without changing the characteristics of the game; see Lemma A.2 for the details.

Note that the restriction to these languages is in some sense minimal: it is indeed easy to show that the computation of the maxmin value for a CBG is essentially trivial if the circuit φ_+ is restricted to be a single term (conjunction of literals) or a single clause (disjunction of literals).

5.1.2 ORACLE GAMES

CBGs are intended to be a very constrained representation of games. At the other extreme, we now define oracle games, a minimally constrained representation for which we will show membership complexity results. Intuitively, an oracle game (respectively a valid oracle game) is a game with an exponential number of vertices $0, 1, \dots, 2^n - 1$, represented in binary over n digits, for which all components are given by oracles (respectively by polynomial-space oracles with a polynomial horizon).

Definition 5.2 (Oracle Game). *A two-player oracle game (OG) is a tuple of the form $G := \langle n, C, P, p, u, IS \rangle$, where n is a positive integer, and C, P, p, u, IS are algorithms such that for all $v, v' \in V(G) := \{0, 1, \dots, 2^n - 1\}$:*¹⁸

- on input v , C returns an ordered list of elements of $V(G)$ (children of v);
- on input v , P returns one of $0, +, -$ (decision-maker at v);
- on inputs v, v' with $P(v) = 0$ and $v' \in C(v)$, p returns a rational number in $[0, 1]$ (probability of v' being the child of v drawn by Nature);
- on input v with $C(v) = \emptyset$, u returns a rational number (utility of leaf v for MAX);
- on inputs v, v' and $i \in \{+, -\}$, IS returns a Boolean value (whether v, v' are in the same information set of i).

16. This fact actually strengthens our hardness results, which will all be proven using the restrictive formalism of CBG.

17. ROBDD stands for “reduced ordered binary decision diagram”; see the work by Darwiche and Marquis (2002).

18. If the number of nodes of a game is not a power of 2, extra nodes are simply disconnected from the root 0.

Definition 5.3 (Valid OG). An OG $G := \langle n, C, P, p, u, IS \rangle$ is said to be valid if the following conditions hold:

1. algorithms C, P, p, u, IS are all deterministic algorithms that run in space at most n and terminate in time at most 2^n ;
2. the output of algorithms p and u is of size at most n (probabilities and utilities have a representation of linear size);
3. the binary relation $\{(v, v') \in V(G)^2 \mid v' \in C(v)\}$ is a tree with node 0 as the root;
4. for all sequences $v_1 \in V(G), v_2 \in C(v_1), \dots, v_k \in C(v_{k-1})$, k is at most n (that is, the game horizon is linear);
5. for all $v \in V(G)$ with $P(v) = 0$, $\sum_{v' \in C(v)} p(v, v') = 1$ holds (p returns a probability distribution at chance nodes);
6. for $i \in \{+, -\}$, $IS(v, v', i) = 1$ only if $P(v) = P(v') = i$;
7. for $i \in \{+, -\}$, the binary relation $\{(v, v') \in V(G)^2 \mid IS(v, v', i) = 1\}$ is an equivalence relation such that for every (v, v') in this relation, $|C(v)| = |C(v')|$ holds (vertices in the same information set have the same number of children).

Observe that we require the oracles to run in linear rather than polynomial space, and similarly we require a linear horizon. However, this assumption is without loss of generality up to polynomial-time reductions. Indeed, if an oracle runs in space n^c rather than n for some constant c , then one can define the OG game with n^c instead of n as its first component (akin to the idea of *padding* used in complexity theory). Similarly, up to a replacement of n by dn for some constant d , this definition is independent of the computational model on which the oracles are supposed to run. For complexity analyses, we define the *size* of a valid oracle game to be n ; in particular, we do not count the representations of the oracles, for which we make no specific assumption.¹⁹

The definition allows one to naturally capture families of games defined by the same rules (oracles) but different sizes (n), e.g. the family of Checkers games played on an $n \times n$ board. If such a family of OGs is valid, this means that it has “reasonable” game rules (essentially, computable in polynomial space). Observe that “more reasonable” game rules, e.g. computable in polynomial time, are also encompassed; hence the membership results given in this section also apply to such games; our definition of OG games with linear space and exponential time makes them more general.

Let us elaborate on the last point. Let OG be the class of all valid OGs, and OG_{poly} be the class of all valid OGs with a polynomial time bound on the oracles (instead of a polynomial space bound). Then by definition, $OG_{\text{poly}} \subseteq OG$, which means for all decision problems that take a valid OG as input, the complexity of such problems for OG is an upper bound on their complexity for OG_{poly} . Contrastingly, lower complexity bounds may fail to go from OG over to OG_{poly} : some problem may be, say, NP-complete for OG but polynomial for OG_{poly} . However, as we will show in the following, the complexity of all decision problems we consider is the same for OG as for the class of all CBGs, which is a subclass of OG_{poly} (up to a polynomial-time translation, see Subsubsection 5.1.3). Hence, as it turns out, all our complexity results (membership and hardness) about OG also apply to OG_{poly} .

19. A reasonable encoding would be given by the input to a fixed Universal Turing Machine.

In particular, our results imply that allowing game rules to be implementable in exponential time (but still in polynomial space) instead of polynomial time does not strictly increase the worst-case complexity of solving them. However, it may be interesting for further work on compact games to consider a definition of a valid OG parameterized by a complexity class for the oracles.

The interpretation of a valid OG as an EFG is straightforward. The actions at an information set can be denoted by integers in such a way that the k -th action maps every vertex v in the information set to the k -th child of v (which is well-defined due to Items 6 and 7 of Definition 5.3). Moreover, given the requirements on the oracles, the following result is straightforward.

Lemma 5.4. *The EFG of chance defined by a given valid OG G can be computed in deterministic exponential time. In particular, this EFG has at most exponential size in the size n of G .*

Let us emphasise that it can be decided in polynomial space whether a given OG is valid, by verifying that no property is violated; for instance, it can be checked that p runs in space at most n by enumerating all pairs of vertices and for each one, running the algorithm until more than n space, or more than 2^n time, is used (if ever); all of this can be done in polynomial space.

Similarly, it can also be decided in polynomial space whether a player in a given OG has perfect information, perfect recall, or multi-agent perfect recall, by verifying that there is no counterexample: for PI, two different nodes in the same information set; for PR, an invalid pair of paths to the same information set; for MA-PR, a path intersecting the same information set twice.

5.1.3 RELATIONSHIPS BETWEEN COMPACT REPRESENTATIONS

It is easy to see that given a CBG $\gamma := \langle X, P, D, \varphi_+ \rangle$ with n_γ variables, an OG $G := \langle n_G, C, P, p, u, IS \rangle$ that defines the same EFG as γ can be computed efficiently.

Indeed, the vertices in the EFG can be encoded efficiently over $n_\gamma + \log n_\gamma$ bits, by encoding the level of a vertex over $\log n_\gamma$ bits, and the assignment of the previous variables over n_γ bits. Given this propositional representation, the oracles in G can be defined to be the straightforward polynomial-space algorithms that, given input in this representation, retrieve the corresponding information from γ ; for instance, u can be implemented by evaluating φ_+ on the values of all variables stored in the representation of the node. Taking $cn_\gamma^{c'}$ to be the maximum space used by these algorithms on the computational model at hand and defining $n_G := \max(n_\gamma + \log n_\gamma, cn_\gamma^{c'})$ indeed yields a valid OG (observe that the horizon of the underlying EFG is n_γ).

To summarise, a family of CBGs can be transformed in polynomial time into a family of valid OGs that define the same family of EFGs. We will in particular use this result to derive that if computing maxmin is hard for CBGs of some type, then it is hard as well for valid OGs of this type.

In addition, since CBG is a relatively minimal formalism to represent compact games, while OG is a relatively powerful one, if we prove that computing the maxmin value for CBGs is as hard as for OGs, then this implies that for any other conceivable intermediate formalism, the complexity of computing the maxmin value is as hard as for both CBG and OG. More concretely, in the following, we prove membership results for OG and hardness results for CBG.

We leave as future work to clarify the precise relationship between oracle games (as we define them) and standard compact languages for defining games, such as the Game Description Language GDL (Genesereth & Thielscher, 2014), the language of Toss (Kaiser & Stafiniak, 2011), POGDDL (Richards & Amir, 2012), or MA-PDDL (Kovacs, 2012). In particular, we conjecture that propositional GDL (restricted to games with a polynomial horizon) can be reduced to OG, while full GDL

cannot (even with the restriction to a polynomial horizon) due to the computational complexity of generating states (nodes in the tree) using logic programming. However, the precise connection remains open; a good starting point may be the complexity results by Bonnet and Saffidine (2014).

5.2 Summary of Results

In the remainder of this section, we consider the following variant of PURE MAXMIN called PURE C-MAXMIN, where “C” stands for “compact”.

Definition 5.5 (PURE C-MAXMIN). *Let \mathcal{G} be a class of zero-sum EFGs. Then PURE C-MAXMIN(\mathcal{G}) is the following decision problem.*

Input: *A valid oracle game G that defines an EFG $G \in \mathcal{G}$, and a rational number m .*

Output: *Does $\underline{v}_+ := \max_{s_+ \in \mathcal{S}_+^P} \min_{s_- \in \mathcal{S}_-^P} \mathcal{U}_+(s_+, s_-) \geq m$ hold in G ?*

Recall that the size of a valid OG $G := \langle n, C, P, p, u, IS \rangle$ is defined to be n .

The complexity class NEXP contains all languages that can be decided in nondeterministic exponential-time. The class NEXP^{NP} (Hemachandra, 1989) contains all languages that can be decided by an exponential-time nondeterministic Turing machine with the help of an NP oracle, which decides the membership of a language in NP in one step.²⁰ Equivalently, NEXP^{NP} contains all languages L such that

$$x \in L \iff \exists y_1 \in \{0, 1\}^{2^{p(|x|)}}, \forall y_2 \in \{0, 1\}^{2^{p(|x|)}}, (x, y_1, y_2) \in L',$$

where p is a polynomial and L' is a language in P. This class is likely to be different from NEXP (Hemachandra, 1989, Note 3, page 312).

The complexity of PURE C-MAXMIN(\mathcal{G}) is summarised in Table 2; see Subsection 4.2 for hints about how to read it. Observe that this table is parallel to Table 1, in the sense that all polynomial-time (respectively NP-complete, Σ_2^P -complete) problems become PSPACE-complete (respectively NEXP-complete, NEXP^{NP} -complete) under compact representations. However, except for membership results, this is not per definition of succinct representations, and we indeed have to prove all hardness results (even if the proofs are simple).

5.3 Membership Results

Proposition 5.6. *PURE C-MAXMIN is in PSPACE for valid OGs of chance in which both MAX and MIN have PI.*

Proof. In this case, we can use backward induction to compute the maxmin value (Proposition 4.7). Since the game has a polynomial horizon by assumption, and the children of a node can be iterated over in PSPACE, this can indeed be done in PSPACE. \square

Proposition 5.7. *PURE C-MAXMIN is in PSPACE for valid OGs of no chance in which MAX has PR and MIN has MA-PR.*

20. However, the Turing machine making use of such an oracle still needs time to write down the queries to the oracle, which, in this case, takes exponential time (Arora & Barak, 2009, Chapter 3).

		No chance	Chance		
		PI/PR/MA-PR	PI	PR	MA-PR
MAX	MIN				
	PI	PSPACE-c [h: 5.10]	PSPACE-c [m: 5.6]	NEXP-c [h: 5.12]	NEXP^{NP}-c [h: 5.14]
	PR	PSPACE-c [m: 5.7]	NEXP-c [h: 5.13]	NEXP-c	NEXP^{NP}-c
	MA-PR	NEXP-c [h: 5.11]	NEXP-c	NEXP-c [m: 5.8]	NEXP^{NP}-c [m: 5.9]

Table 2: Complexity of PURE C-MAXMIN. All hardness results hold even under the restriction to CBGs with CNF (respectively DNF, ROBDD) circuits. PI, PR, MA-PR stand for perfect information, perfect recall, multi-agent perfect recall, respectively. Only key membership (“m”) and hardness (“h”) results are referred to; the others can be deduced by monotonicity. Results in bold are new from this paper; other results are direct consequences of known results (see the citations in the referred statements).

Proof. We consider the algorithm of Proposition 4.5, which computes the maxmin value by a bottom-up induction on the game tree. Since the game tree of an oracle game has polynomial depth and the children of each information set can be iterated over in PSPACE, this computation can be done in PSPACE from the oracle representation. \square

The following two results are obtained directly by using Lemma 5.4 to first compute, in deterministic exponential time, the EFG defined by the given valid OG, then running a nondeterministic polynomial-time algorithm (Proposition 4.4), respectively a nondeterministic polynomial-time algorithm with an NP-oracle (Proposition 4.12) on this EFG of exponential size.

Proposition 5.8. PURE C-MAXMIN is in NEXP for valid OGs of chance in which MAX has MA-PR and MIN has PR.

Proposition 5.9. PURE C-MAXMIN is in NEXP^{NP} for valid OGs of chance in which both MAX and MIN have MA-PR.

5.4 Hardness Results

As discussed above, we give all hardness results for compact Boolean games. Since all proofs use a polynomial-time reduction from a decision problem to PURE C-MAXMIN for a family of CBGs, all these results also imply that PURE C-MAXMIN is at least as hard for the corresponding family of valid OGs, as discussed in Subsubsection 5.1.3.

The two first hardness results are essentially given by the hardness of QBF and DQBF for PSPACE and NEXP, respectively.

Proposition 5.10 (Stockmeyer & Meyer, 1973). PURE C-MAXMIN is PSPACE-hard for CBGs of no chance in which both MAX and MIN have PI.

Proof. It is easy to see that for a given QBF $Q_1 x_1 Q_2 x_2 \cdots Q_n x_n \varphi(x_1, x_2, \dots, x_n)$, the CBG $\langle X, P, D, \varphi \rangle$, with $X = (x_1, \dots, x_n)$, $P(i) = +$ for $Q_i = \exists$ and $P_i = -$ for $Q_i = \forall$, and $D(x_i) = \{x_1, \dots, x_{i-1}\}$ for all i , is a game of no chance and with perfect information. Moreover, it has a maxmin value of 1 if and only if the QBF is valid. We conclude using the PSPACE-hardness of deciding the validity of a QBF (Stockmeyer & Meyer, 1973). \square

Proposition 5.11 (Peterson et al., 2001, Thm. 5.2.1). *PURE C-MAXMIN is NEXP-hard for CBGs of no chance in which MAX has MA-PR and MIN has PR.*

Proof. Peterson et al. (2001, Theorem 5.2.1) show that deciding the validity of a DQBF of the form

$$\begin{aligned} & \forall x_1^1 \cdots \forall x_1^{k_1} \quad \forall x_2^1 \cdots \forall x_2^{k_2} \\ & \exists y_1^1(x_1^1, \dots, x_1^{k_1}) \cdots \exists y_1^{\ell_1}(x_1^1, \dots, x_1^{k_1}) \quad \exists y_2^1(x_2^1, \dots, x_2^{k_2}) \cdots \exists y_2^{\ell_2}(x_2^1, \dots, x_2^{k_2}) \\ & \varphi(x_1^1, \dots, x_1^{k_1}, x_2^1, \dots, x_2^{k_2}, y_1^1, \dots, y_1^{\ell_1}, y_2^1, \dots, y_2^{\ell_2}) \end{aligned}$$

is already NEXP-hard. We define the CBG of no chance $\langle X, P, D, \varphi \rangle$, where

$$X := (x_1^1, \dots, x_1^{k_1}, x_2^1, \dots, x_2^{k_2}, y_1^1, \dots, y_1^{\ell_1}, y_2^1, \dots, y_2^{\ell_2}),$$

and for all j , $P(x_1^j) = P(x_2^j) := -$, $P(y_1^j) = P(y_2^j) := +$, and

$$\begin{aligned} D(x_1^j) &:= \{x_1^1, \dots, x_1^{j-1}\}, \\ D(x_2^j) &:= \{x_1^1, \dots, x_1^{k_1}, x_2^1, \dots, x_2^{j-1}\}, \\ D(y_1^j) &:= \{x_1^1, \dots, x_1^{k_1}, y_1^1, \dots, y_1^{j-1}\}, \\ D(y_2^j) &:= \{x_2^1, \dots, x_2^{k_2}, y_2^1, \dots, y_2^{j-1}\}. \end{aligned}$$

It is easy to see that MAX has MA-PR and MIN has PI in this CBG. In addition, its pure maxmin value is 1 if and only if the DQBF is valid, which concludes. \square

Proposition 5.12. *PURE C-MAXMIN is NEXP-hard for CBGs of chance in which MAX has PI and MIN has PR.*

Proof. We give a reduction from $\text{TILING}(2^n, 2^n)$, which is NEXP-complete (Schwarzentruber, 2019, Theorem 3) and is defined as follows:

Input: A finite set C with a distinguished element $w \in C$, a set of tiles $T \subseteq C^4$, a tile $t_{1,1} \in T$, and a natural number n expressed in unary.

Output: Is there a legal tiling τ of $S_{2^n} = \{1, \dots, 2^n\} \times \{1, \dots, 2^n\}$ with $\tau(1, 1) = t_{1,1}$?

Given an instance $(C, w, T, t_{1,1}, n)$, we build a game as follows (the encoding as a CBG is given afterwards):

- *Players:* Nature, MAX with perfect information, MIN with perfect recall.
- *Game tree:* At the root, Nature picks a cell $(r, c) \in S_{2^n}$, uniformly at random. Then MAX observes (r, c) and chooses $t \in T$. Finally, MIN observes neither (r, c) nor t , and chooses $(r_1, c_1), (r_2, c_2) \in S_{2^n}$ and $t_1, t_2 \in T$.
- *Payoffs:* MAX's payoff is:
 1. 0, if t is not legal at (r, c) with respect to colour w (for instance, $c = 1$ and $l(t) \neq w$);
 2. otherwise, 0, if $(r, c) = (1, 1)$ and $t \neq t_{1,1}$;
 3. otherwise, $1/2$, if (i) (r_1, c_1) and (r_2, c_2) are adjacent cells, (ii) t_1, t_2 have different colours on their common side (for instance, $(r_1, c_1 + 1) = (r_2, c_2)$ and $r(t_1) \neq l(t_2)$), and (iii) $(r, c, t) = (r_1, c_1, t_1)$ or $(r, c, t) = (r_2, c_2, t_2)$;

4. otherwise, 1.

- *Threshold:* The threshold of maxmin value is $1 - 1/2^{2n+1}$.

We now show that if $(C, w, T, t_{1,1}, n)$ is a positive instance of $\text{TILING}(2^n, 2^n)$, then the game has a maxmin value of at least $1 - 1/2^{2n+1}$, and otherwise of at most $1 - 1/2^{2n}$.

Assume first that there is a legal tiling τ consistent with $t_{1,1}$, and MAX plays according to it (that is, always chooses t to be $\tau(r, c)$). Since MIN observes nothing, their strategy is to always choose the same $r_1, c_1, r_2, c_2, t_1, t_2$. Since τ is legal and consistent with $t_{1,1}$, MAX never receives a payoff of 0, and receives $1/2$ only if the third condition of the payoff function is satisfied. If (r_1, c_1) and (r_2, c_2) are adjacent, since τ is legal and t_1, t_2 have different colours on the common side, it cannot be the case that both $t_1 = \tau(r_1, c_1)$ and $t_2 = \tau(r_2, c_2)$ hold. Assume by symmetry $t_2 \neq \tau(r_2, c_2)$. Then MAX receives 1 in all cases, except possibly when Nature chooses (r, c) to be (r_1, c_1) (and $t_1 = \tau(r_1, c_1)$ holds). The latter happens with probability $1/(2^n)^2 = 1/2^{2n}$, so that MAX has an expected payoff of at least $1/2 \times 1/2^{2n} + 1 \times (1 - 1/2^{2n}) = 1 - 1/2^{2n+1}$.

Conversely, assume that there is no legal tiling consistent with $t_{1,1}$, and let τ be the tiling defined by MAX's strategy. Then if τ is not legal with respect to colour w , let (r_0, c_0) be a cell witnessing this; MAX receives 0 when Nature chooses (r, c) to be (r_0, c_0) , which happens with probability $1/2^{2n}$. Similarly, if τ is not consistent with $t_{1,1}$, then MAX receives 0 with probability $1/2^{2n}$. Since at all other leaves MAX can receive at most 1, this means MAX's expected payoff is at most $1 - 1/2^{2n}$. Otherwise, τ is not legal with respect to adjacency; let $(r_1, c_1), (r_2, c_2)$ be two cells witnessing this, and consider MIN's strategy to play $r_1, c_1, \tau(r_1, c_1), r_2, c_2, \tau(r_2, c_2)$. Then MAX receives $1/2$ whenever Nature picks (r, c) to be either (r_1, c_1) or (r_2, c_2) , which happens with probability $2/2^{2n}$. Since at all other leaves MAX can receive at most 1, this yields an expected payoff of at most $1/2 \times 2/2^{2n} + 1 \times (1 - 2/2^{2n}) = 1 - 1/2^{2n}$.

We now show that a CBG equivalent to this game can be computed in polynomial time from the instance of $\text{TILING}(2^n, 2^n)$.

First, observe that we can assume the number $|T|$ of tiles to be a power of 2, since adding dummy colours and the corresponding monochromatic tiles does not change the existence of a legal tiling (these could not be part of one). Hence, the choice of each of $r, c, t, r_1, c_1, t_1, r_2, c_2, t_2$ can be replaced by the choice of n or $\log|T|$ bits, encoding the corresponding number; since n is encoded in unary and T is given explicitly in the input, this only requires a polynomial number of Boolean variables. The owner and dependency functions for a bit are defined as for the corresponding index or tile in the game above. So there remains to define the Boolean circuit for the CBG.

As a first step, we define a Boolean circuit for each condition in the definition of MAX's payoff, in terms of the Boolean variables introduced above. This yields four circuits: φ_+^1 outputs 1 if the input Boolean variables are such that r, c, t , as encoded by these variables, satisfy that t is not legal at (r, c) with respect to colour w , and similarly for $\varphi_+^2, \varphi_+^3, \varphi_+^4$; details are left to the reader, but clearly, such circuits can be constructed in polynomial time given that n is encoded in unary and T is given explicitly in the input. To cope with the payoff $1/2$ in the third condition of the payoff function, we introduce an extra variable y owned by Nature and ordered last, and finally define φ_+ to be $\neg\varphi_+^1 \wedge \neg\varphi_+^2 \wedge ((y \wedge \varphi_+^3) \vee (\neg\varphi_+^3 \wedge \varphi_+^4))$.²¹ Then it is easy to verify that the expected utility of any pair of strategies as induced by φ_+ is exactly the same as in the game above. \square

21. The idea of using a chance variable to model the payoff $1/2$ is similar in spirit to the one in Lemma A.1.

By using the same proof as for Proposition 5.12, but having MIN’s variables ordered first and letting MAX observe only Nature’s choice of (r, c) , we obtain the same result for the dual setting.

Proposition 5.13. *PURE C-MAXMIN is NEXP-hard for CBGs of chance in which MAX has PR and MIN has PI.*

Finally, we turn to the case in which MIN only has multi-agent perfect recall.

Proposition 5.14. *PURE C-MAXMIN is NEXP^{NP}-hard for CBGs of chance in which MAX has PI and MIN has MA-PR.*

Proof. This can be shown by a polynomial-time reduction from the exponential version of FINITE TILING EXTENSION introduced in Subsection 4.4; this exponential version is defined as follows:

Input: A finite set C with a distinguished element $w \in C$, a set of tiles $T \subseteq C^4$, and a natural number n expressed in unary.

Output: Is there a *non-extendable* legal tiling of the first row of S_{2^n} , that is, a legal tiling of $S_{1,2^n} := \{1\} \times \{1, \dots, 2^n\}$ that cannot be extended to a legal tiling of S_{2^n} ?

This problem is NEXP^{NP}-hard: indeed, as observed by Goldsmith and Mundhenk (2007, Section 2.3), the “master reduction” from computations of Turing machines to tilings (van Emde Boas, 1997) can be used to lift the proof of Σ_2^P -hardness to the exponential case.

With this in hand, we build exactly the same game as we did for the non-compact case in the proof of Proposition 4.10, and the proof that an equivalent CBG can be computed in polynomial time goes exactly as in the proof of Proposition 5.12. \square

6. Complexity against Opponent Models

We now consider the situation in which MIN is only allowed to choose from a finite set of behaviour strategies. This setting corresponds to playing a game against an opponent whose behaviour or reasoning is captured by a model we know. Such a model is called an *opponent model* (cf. Section 2). The same setting also captures the problem of planning in uncertain environments or with adversarial cost functions (McMahan et al., 2003). Concretely, we study the following problem.

Definition 6.1 (PURE OM-MAXMIN). *Let \mathcal{G} be a class of zero-sum EFGs. Then PURE OM-MAXMIN(\mathcal{G}) is the following decision problem.*

Input: An EFG $G \in \mathcal{G}$, a rational number m , and a finite set $\mathcal{S}_-^O \subseteq \mathcal{S}_-^B$ of MIN’s behaviour strategies in G .

Output: Does $\underline{v}_+ := \max_{s_+ \in \mathcal{S}_+^P} \min_{\pi_- \in \mathcal{S}_-^O} \mathcal{U}_+(s_+, \pi_-) \geq m$ hold in G ?

As in Section 4, we consider the complexity for EFGs; we define the size of the input of this problem to be the sum of $\|G\|$, $\|m\|$, and the sizes of the strategies in \mathcal{S}_-^O ; the size of a behaviour strategy $\pi_- \in \mathcal{S}_-^O$ is defined to be the sum of the sizes of the rational numbers (probabilities) that define π_- , over all MIN’s nodes in the game tree.²²

In the following, strategies in $\mathcal{S}_-^O \subseteq \mathcal{S}_-^B$ will be referred to as opponent models (shortened as OMs). A game with OMs is said to be of no chance, if the original game is of no chance, and all OM

22. We do not consider the complexity for compactly represented games, since this would require defining compact representations of opponent models, which is out of the scope of this paper.

MAX \ $ S_-^O $	1			constant (≥ 2)			unbounded		
	No chance								
	PI	P		P		P		P	
	PR	P		P		P [m: 6.5]			
	MA-PR	P [m: 6.3]		P [m: 6.6]		NP-c [m: 6.2, h: 6.7]			
Chance									
	PI	P		NP-c [h: 6.8]		NP-c			
	PR	P [m: 6.3]		NP-c		NP-c			
	MA-PR	NP-c [h: 6.4]		NP-c		NP-c [m: 6.2]			

Table 3: Complexity of PURE OM-MAXMIN. All hardness results hold even under the restriction to Boolean timeable games with at most 2 agents for MAX. PI, PR, MA-PR stand for perfect information, perfect recall, multi-agent perfect recall, respectively. Only key membership (“m”) and hardness (“h”) results are referred to; the others can be deduced by monotonicity. Results in bold are new from this paper; other results are direct consequences of known results (see the citations in the referred statements).

strategies in S_-^O are pure strategies. This terminology is consistent, since a player with a non-pure OM strategy is indistinguishable from Nature, a player whose behaviour strategy is known.

Observe that an instance of PURE MAXMIN can be considered as an instance of PURE OM-MAXMIN, by defining the set of opponent models to be the set of all MIN’s pure strategies. However, this set is, in general, exponential in the size of the game tree. As a consequence, this does not give a polynomial-time reduction from the former to the latter problem.

6.1 Summary of Results

The complexity of PURE OM-MAXMIN(\mathcal{G}) is summarised in Table 3; see Subsection 4.2 for hints about how to read it. Note that it does not make sense to distinguish between different degrees of imperfect information of MIN, but it makes sense to study the complexity with respect to the number of OMs (1, a constant but at least 2, or unbounded).

Since we only consider pure strategies for MAX, this problem is trivially in NP, which provides an upper bound for all cases in Table 3.

Proposition 6.2. PURE OM-MAXMIN is in NP.

Proof. One can guess a pure strategy of MAX, then verify that it yields an expected payoff no less than the threshold against all the OMs given in the input, all in time linear in the size of the input. \square

6.2 Complexity of Best Responses in EFGs

We begin by considering the case $|S_-^O| = 1$, i.e. MIN’s behaviour strategy is fixed. Under this circumstance, a two-player game is transformed into a one-player game, which has no chance if the original game has no chance and MIN’s known strategy is a pure strategy.

Let us write π_- for the only strategy in \mathcal{S}_-^0 . The maxmin value reads $v_+ = \max_{s_+ \in \mathcal{S}_+^P} \mathcal{U}_+(s_+, \pi_-)$, which is the value of MAX's pure best responses to π_- . Hence, studying the complexity of the case $|\mathcal{S}_-^0| = 1$ is equivalent to studying the complexity of finding the best responses of a player.

Proposition 6.3 (Koller & Megiddo, 1992, Sec. 3.3). *PURE OM-MAXMIN with only one OM is decidable in linear time, and is a fortiori in P, for EFGs of no chance in which MAX has MA-PR, and for EFGs of chance in which MAX has PR.*

However, if MAX only has MA-PR and there is chance (due to chance nodes and/or due to MIN's behaviour strategy π_-), then the problem becomes intractable. The following is essentially the result by von Stengel and Forges (2008, Theorem 1.3), but viewing Nature as a MIN player with a single OM corresponding to the uniformly random behaviour strategy.

Proposition 6.4 (von Stengel & Forges, 2008, Thm. 1.3). *PURE OM-MAXMIN with only one OM is NP-hard for EFGs of chance in which MAX has MA-PR. The result holds even under the restriction to 2 agents for MAX and to Boolean timeable games.*

6.3 Complexity of EFGs with Multiple OMs

We now consider the case $|\mathcal{S}_-^0| \geq 2$, starting with games of no chance. In this problem setting, the OMs are pure strategies of MIN. Hence, we can write $\mathcal{S}_-^0 = \{s_{-,1}, \dots, s_{-,|\mathcal{S}_-^0|}\}$, where $s_{-,i} \in \mathcal{S}_-^P$ for $1 \leq i \leq |\mathcal{S}_-^0|$.

Proposition 6.5. *PURE OM-MAXMIN with multiple OMs is in P for EFGs of no chance in which MAX has PR.*

Proof. We reduce this problem to the equivalent one without OM.²³

Let (G, m, \mathcal{S}_-^0) be an instance, and write T for its game tree. For each OM $s_{-,i} \in \mathcal{S}_-^0$, write T_i for the tree obtained from T by removing all edges which do not correspond to the move by $s_{-,i}$ at MIN's nodes (thereby also removing the subtrees reached by such edges). Moreover, for each information set IS_+^j of MAX in G , write $IS_{+,i}^j$ for the corresponding information set in T_i (restricted to the nodes in T_i).

Now define the game G' (without OM) as follows:

- the tree of G' has a root for player MIN, with one child per OM $s_{-,i} \in \mathcal{S}_-^0$, leading to the subtree T_i (that is, the game starts with MIN choosing their future behaviour);
- all information sets of MIN in G' are singletons;
- the information sets of MAX are the sets $\bigcup_{s_{-,i} \in \mathcal{S}_-^0} IS_{+,i}^j$, for all information sets IS_+^j of MAX in G (that is, MAX has the same information as in G , and does not observe the behaviour/OM chosen by MIN);
- the utility function is as in G .

Now, if MIN chooses $s_{-,i}$ at the root and MAX plays some strategy s_+ , the node reached in G' is the node of T_i corresponding to the node reached in G if MAX plays s_+ against the OM $s_{-,i}$. It follows that the maxmin value of G' is the same as the one of G . Since the construction is clearly polynomial, and G' is an EFG of no chance in which MAX has perfect recall and MIN has perfect information, we conclude with Proposition 4.5. \square

23. We are truly grateful to an anonymous reviewer for suggesting this reduction.

If MAX only has multi-agent perfect recall, Proposition 4.5 no longer applies. However, if the number of OMs is bounded by a constant, then the decision problem is still in P.

Proposition 6.6. *For every $k \geq 1$, PURE OM-MAXMIN with k OMs is in P for EFGs of no chance in which MAX has MA-PR.*

Proof. Let us write $\mathcal{S}_-^O = \{s_{-,1}, \dots, s_{-,k}\}$, where $s_{-,i} \in \mathcal{S}_-^P$ for $1 \leq i \leq k$. We first observe that for a fixed pure strategy of MAX and a fixed OM, a unique leaf of the game tree T is reached. Hence, the outcome of a given strategy s_+ of MAX against all OMs can be represented by a k -tuple of leaves (l_1, \dots, l_k) (possibly with repetitions) such that l_i is the leaf reached by the ployout under the profile $(s_+, s_{-,i})$. For every constant k , there are only $\mathcal{O}(|L(T)|^k)$ such tuples. Hence, we can enumerate them to decide whether there is one k -tuple for which (i) the value $\min_{1 \leq i \leq k} u_+(l_i)$ is at least m and (ii) the tuple indeed corresponds to the outcomes of some pure strategy of MAX against the k OMs.

Algorithm 1: Polynomial-time algorithm for verifying that a k -tuple (l_1, \dots, l_k) is reachable under some pure strategy of MAX against the k OMs.

```

1  $s_+ \leftarrow$  empty mapping from  $IS_+$  to actions
2 for  $i = 1, \dots, k$ :
3    $pl \leftarrow$  the unique path from  $r$  to  $l_i$ 
4   for  $v \in pl$ :
5      $a \leftarrow$  the action taken at  $v$  in  $pl$ 
6     if  $P(v) = -$ :
7       if  $a \neq s_{-,i}(v)$ :
8         return False           /*  $l_i$  is not consistent with  $s_{-,i}$  */
9       continue               /*  $v$  is good, look at the next one */
10     $IS_+ \leftarrow$  the information set of MAX containing  $v$ 
11    if  $s_+(IS_+)$  is not set yet:
12       $s_+(IS_+) \leftarrow a$ 
13    elif  $s_+(IS_+) \neq pl(v)$ :
14      return False           /* no  $s_+$  can reach  $(l_1, \dots, l_k)$  */
15 return True
    
```

It remains to show that (ii) can be decided in polynomial time. The whole algorithm is given as Algorithm 1. We first check that on the unique path from the root to l_i , all decisions taken at MIN's nodes are indeed as prescribed by the i -th OM $s_{-,i}$. We then check that these k paths indeed prescribe the same action at every information set of MAX. Clearly, Algorithm 1 runs in polynomial time (more precisely, in $\mathcal{O}(k|T|)$ time), which completes the proof. \square

It turns out that a constant number of OMs is essentially the best we can do when MAX only has multi-agent perfect recall.

Proposition 6.7. *PURE OM-MAXMIN with multiple OMs is NP-hard for EFGs of no chance in which MAX has MA-PR. The result holds even under the restriction to 2 agents for MAX, to Boolean timeable games, and to polynomially (in the size of the game tree) many OMs.*

Proof. Consider the reduction in Proposition 4.6. Notice that in the game obtained by the reduction, MIN has n pure strategies, one for each clause. If we take all these strategies as OMs, then a 3-CNF formula is satisfiable if and only if MAX has a pure strategy with payoff 1 against all these OMs. \square

Remark. Another way to interpret this result is by comparing it to Proposition 6.4: the proof by von Stengel and Forges (2008) is essentially the same as the proof of Proposition 6.7, but involves a single behaviour OM (or Nature) corresponding to the uniform mixture of the n pure OMs above.

We finally turn to the case of multiple OMs, for games of chance. Surprisingly, in the presence of chance, computing the pure maxmin value against only 2 OMs is already NP-hard, even if MAX has perfect information.

Proposition 6.8. PURE OM-MAXMIN with 2 OMs is NP-hard for EFGs of chance in which MAX has PI. The result holds even under the restriction to Boolean timeable games.

Proof. The proof follows from the proof of Proposition 4.8. Indeed, it is shown there that PURE MAXMIN is NP-hard even if MIN has only 2 strategies; seeing these as OMs gives a polynomial-time reduction from SUBSET SUM to PURE OM-MAXMIN with 2 OMs. \square

7. Other Variants of PURE MAXMIN

We finally briefly discuss two natural variants of PURE MAXMIN for EFGs.

Definition 7.1 (PURE \leq -MAXMIN). Let \mathcal{G} be a class of zero-sum EFGs. Then PURE \leq -MAXMIN(\mathcal{G}) is the following decision problem.

Input: An EFG $G \in \mathcal{G}$ and a rational number m .
 Output: Does $\underline{v}_+ := \max_{s_+ \in \mathcal{S}_+^P} \min_{s_- \in \mathcal{S}_-^P} \mathcal{U}_+(s_+, s_-) \leq m$ hold in G ?

Notice that PURE \leq -MAXMIN(\mathcal{G}) is *not* the complement of PURE MAXMIN(\mathcal{G}), which would consist in deciding whether the maxmin value of a game is *strictly* smaller than a given threshold. Still, the complexity class of PURE \leq -MAXMIN(\mathcal{G}) turns out to be the complement of that of PURE MAXMIN(\mathcal{G}).

Definition 7.2 (PURE =-MAXMIN). Let \mathcal{G} be a class of zero-sum EFGs. Then PURE =-MAXMIN(\mathcal{G}) is the following decision problem.

Input: An EFG $G \in \mathcal{G}$ and a rational number m .
 Output: Does $\underline{v}_+ = m$ hold in G ?

The results are summarized in Table 4. As it turns out, all results are parallel to those for PURE MAXMIN in Table 1.

Proposition 7.3. The P-membership results in Table 4 hold.

Proof. This follows from the fact that there is a polynomial-time algorithm for *computing* the maxmin value in these cases (cf. Proposition 4.5 and Proposition 4.7). \square

Proposition 7.4. The coNP-completeness and Π_2^P -completeness results in Table 4 hold. They hold even under the restrictions to 2 agents for MAX (in case MAX has MA-PR) and to Boolean timeable games.

		No chance	Chance		
		PI/PR/MA-PR	PI	PR	MA-PR
MAX \ MIN	PI	P	P	coNP-c/DP-c	Π_2^P -c/ D_2^P -c
	PR	P	coNP-c/DP-c	coNP-c/DP-c	Π_2^P -c/ D_2^P -c
	MA-PR	coNP-c/DP-c*	coNP-c/DP-c	coNP-c/DP-c	Π_2^P -c/ D_2^P -c

Table 4: Complexity of PURE \leq -MAXMIN and PURE =-MAXMIN, on the left and right of each cell, respectively. All hardness results hold even under the restriction to timeable games, and, except for the DP-completeness result marked with a star, under the further restriction to Boolean games. PI, PR, MA-PR stand for perfect information, perfect recall, multi-agent perfect recall, respectively.

Proof. For membership in coNP, we can check that for all pure strategies of MAX, the best response of MIN, computable in linear time in these cases (Koller & Megiddo, 1992, Proposition 2.7), yields at most m for MAX.

For membership in Π_2^P , we can check that for all pure strategies of MAX, there exists a strategy of MIN which yields at most m for MAX.


For hardness, the argument is similar for all cases. Consider for example EFGs of chance in which MAX has PI and MIN has PR. In the proof of Proposition 4.8, we give a reduction from SUBSET SUM to PURE MAXMIN, such that if the instance is positive, the maxmin value of the EFG obtained is at least k , and otherwise it is at most $k - 1$. Hence, this also gives us a reduction from the coNP-complete complement of SUBSET SUM to PURE \leq -MAXMIN with $k - 1$ as threshold.

The other cases are similar:

- for EFGs of chance in which MAX has PR and MIN has PI (cf. Proposition 4.9), EFGs from “yes” instances of SUBSET SUM have a maxmin value of at least k while those from “no” instances have a maxmin value of at most $k - 1$;
- for EFGs of no chance in which MAX has MA-PR and MIN has PI (cf. Proposition 4.6), EFGs from “yes” instances of the satisfiability problem have a maxmin value of 1, while those from “no” instances have a maxmin value of 0;
- finally, for EFGs of chance in which MAX has PI and MIN has MA-PR (cf. Proposition 4.10), EFGs from “yes” instances of TILING EXTENSION have a maxmin value of at least $1/m^4$ while those from “no” instances have a maxmin value of at most 0.

□

To describe the complexity of PURE =-MAXMIN, we make use of the classes D_k^P . For each $k \geq 1$, D_k^P is the set of languages L that can be written as $L = L_1 \cap L_2$ with $L_1 \in \Sigma_k^P$ and $L_2 \in \Pi_k^P$, or equivalently as the difference of two languages in Σ_k^P (or two languages in Π_k^P). The most well-known one is D_1^P , which is also called DP in the literature. Intuitively, DP corresponds to problems that concern the optimal value of an optimisation problem in NP (Papadimitriou, 1994).


 Figure 2: An EFG with subgames G_1 and G_2 .

By generalising the argument by Papadimitriou (1994, Theorem 17.1), one can prove the D_k^P -hardness of a language L by showing that there are a Σ_k^P -hard language L_1 , a Π_k^P -hard language L_2 , and a polynomial-time reduction f such that $\langle x, y \rangle \in L_1 \times L_2 \iff f(\langle x, y \rangle) \in L$.

When $\text{PURE MAXMIN}(\mathcal{G})$ is NP-complete, we have observed that $\text{PURE } \leq\text{-MAXMIN}(\mathcal{G})$ is coNP-complete, hence $\text{PURE } =\text{-MAXMIN}(\mathcal{G})$ is in DP (since a game has a maxmin value k if and only if its maxmin value is both at least k and at most k). We show below that it is also DP-hard.²⁴

Proposition 7.5. *The DP-completeness results in Table 4 hold for EFGs of chance in which MAX has PI and MIN has PR, and for EFGs of chance in which MAX has PR and MIN has PI. They hold even under the restriction to Boolean timeable games.*

Proof. The membership results are shown as discussed above. For hardness, we will give a reduction from two instances of SUBSET SUM such that the EFG of chance thus obtained has a certain maxmin value if and only if the first instance of SUBSET SUM is a “yes” instance and the second one is not.

Let (S_1^0, k_1^0) and (S_2^0, k_2^0) be two arbitrary instances of SUBSET SUM, and define the two instances (S_1, k) and (S_2, k) by $S_1 = \{2ik_2^0 \mid i \in S_1^0\}$, $S_2 = \{2k_1^0j \mid j \in S_2^0\}$, and $k = 2k_1^0k_2^0$. Clearly, (S_1, k) (respectively (S_2, k)) is a positive instance if and only if so is (S_1^0, k_1^0) (respectively (S_2^0, k_2^0)).

Consider first the case in which MAX has PI and MIN has PR. Let G_1 (respectively G_2) be the timeable EFG of chance built from (S_1, k) (respectively (S_2, k)) in the proof of Proposition 4.8. Then G_1 (respectively G_2) has a maxmin value of at least k if (S_1, k) (respectively (S_2, k)) is a positive instance of SUBSET SUM, and at most $k - 2$ instead (since all elements in S_1 and S_2 are even).

Now consider the EFG G depicted in Figure 2: at the root, MIN chooses whether to play G_1 ; if they do, then the game proceeds as in G_1 . Otherwise, MAX chooses whether to play the game G_2 ; if they do, then the game proceeds as in G_2 . Otherwise, MAX receives a payoff of $k - 1$.

Since both G_1 and G_2 are timeable EFGs of chance in which MAX has PI and MIN has PR, the game G is also a timeable EFG of chance in which MAX has PI and MIN has PR. In addition, the construction is polynomial-time in the two instances of SUBSET SUM in input.

It remains to show that G has a maxmin value of exactly $k - 1$ if and only if (S_1, k) is a positive instance of SUBSET SUM and (S_2, k) is not. This is straightforward to verify:

- if (S_1, k) is a negative instance, then MIN can choose to play G_1 so that the maxmin value of G is at most $k - 2 < k - 1$;
- if both (S_1, k) and (S_2, k) are positive instances, the maxmin value of G is at least $k > k - 1$ since MAX can get at least k from both G_1 and G_2 ;

24. The fact that a decision problem is both NP-hard and coNP-hard, is, in general, *not* enough for it to be DP-hard.

- if (S_1, k) is a positive instance but (S_2, k) is not, then it is optimal for MIN to choose not to play G_1 , and for MAX to choose not to play G_2 . Hence, the maxmin value of G is exactly $k - 1$.

This concludes the DP-hardness of =-MAXMIN for timeable EFGs of chance in which MAX has PI and MIN has PR. The same construction as above with Proposition 4.9 also shows the DP-hardness of =-MAXMIN for timeable EFGs of chance in which MAX has PR and MIN has PI.

Finally, these two problems remain DP-hard for Boolean games since we can compile away the constant $k - 1$ and the non-Boolean payoffs in G_1 and G_2 using Lemma A.1. \square

Proposition 7.6. *The DP-completeness results in Table 4 hold for EFGs in which MAX has MA-PR. They hold even under the restriction to 2 agents for MAX and to timeable games. For EFGs of chance, they hold even under the further restriction to Boolean games.*

Proof. For EFG of no chance in which MAX has MA-PR and MIN has PI, we use the reduction in the proof of Proposition 4.6, which yields a timeable EFG of no chance with 2 agents for MAX; its maxmin value is 1 if the formula is satisfiable, and 0 otherwise.

The construction then is the same as in Figure 2, but the payoff of the leftmost leaf is replaced by $1/2$. By construction, the reduction is to an EFG with three values (0, $1/2$ and 1). For games of chance, the constant $1/2$ can be compiled into Boolean payoffs using Lemma A.1. \square

Remark. *When chance nodes are not allowed, we cannot compile the constant $1/2$ into Boolean payoffs. As a result, we could not prove that =-MAXMIN for Boolean EFGs of no chance in which MAX has MA-PR and MIN has PI is DP-complete. Actually, we conjecture that it is $\text{P}_{1\text{-tt}}^{\text{NP}}$ -complete (Beigel, 1991), that is, complete for the class of problems which are solvable in polynomial time with 1 call to an NP-oracle (in particular, it is both NP- and coNP-hard). However, the details are out of scope of this paper.*

Similarly, when $\text{PURE MAXMIN}(\mathcal{G})$ is Σ_2^{P} -complete, we have observed that $\text{PURE } \leq\text{-MAXMIN}(\mathcal{G})$ is Π_2^{P} -complete, hence $\text{PURE } \text{=-MAXMIN}(\mathcal{G})$ is in D_2^{P} ; we show below that it is also D_2^{P} -hard.

Proposition 7.7. *The D_2^{P} -completeness results in Table 4 hold. They hold even under the restriction to 2 agents for MIN and to Boolean timeable games.*

Proof. The membership results are shown as discussed above. For hardness, we use a construction similar to the one in Proposition 7.5, but with a reduction from a pair of instances of TILING EXTENSION (cf. Proposition 4.10).

Given the timeable EFGs of chance G_1 reduced from $(C_1, w_1, T_1, 1^{m_1})$ and G_2 reduced from $(C_2, w_2, T_2, 1^{m_2})$, let $m := \max(m_1, m_2)$ and let G be the game built as in Figure 2, but with the payoff of the leftmost leaf replaced by $1/2m^4$. Then G is timeable and has a pure maxmin value of exactly $1/2m^4$ if and only if the first instance of TILING EXTENSION is a “yes” instance and the second one is not. To show that the problem remains hard for Boolean games, we can use Lemma A.1 to compile the payoffs in G into Boolean ones. \square

8. Conclusion

We have thoroughly investigated the computational complexity of finding a lower bound on the pure maxmin value in two-player zero-sum EFGs. For each degree of information (perfect information,

	MIN	PI	PR	MA-PR
MAX				
	PI	P	P	coNP-c
	PR	P	P	coNP-c
	MA-PR	NP-c	NP-c	$\Sigma_2^P\text{-c}/\Delta_2^P\text{-c}$

Table 5: Complexity of BEHAVIOUR MAXMIN and MIXED MAXMIN for EFGs of chance. In the last cell, BEHAVIOUR MAXMIN and MIXED MAXMIN are Σ_2^P -complete and Δ_2^P -complete, respectively. PI, PR, MA-PR stand for perfect information, perfect recall, multi-agent perfect recall, respectively.

perfect recall, multi-agent perfect recall) for MAX and MIN, and for games of no chance or of chance, we have either given a polynomial-time algorithm, or shown completeness for a certain complexity class. This allows us to have a complete landscape of this problem (Table 1). In addition, we have studied the complexity landscape of the same decision problem, but under two other settings: when the EFGs are defined by some compact representations (for which we have proposed a very generic definition) (Table 2); when MIN is known to pick strategies from a finite set known to MAX (Table 3). We have also studied the complexity of verifying an upper bound or the exact value of pure maxmin (Table 4).

Some hardness results presented in this work are already known in the literature (Koller & Megiddo, 1992, for instance). However, we have strengthened many of these results in different ways: by giving a simpler reduction; by restricting the degree of imperfect information or the number of strategies of a player; by restricting to Boolean payoffs; etc. We emphasise that all the hardness results in our work hold under strong restrictions: timeable games; at most 2 agents for each player; only Boolean payoffs;²⁵ only chance nodes with uniform distribution; only one turn per agent of MAX or MIN.

Related work on the complexity of maxmin We have focused on the complexity of pure maxmin, but not on behaviour maxmin or mixed maxmin, since the latter ones are well-studied in the literature. For reference (and comparison with pure maxmin), we give the complexity results concerning behaviour and mixed maxmin for EFGs of chance in Table 5; all come from the work by Koller and Megiddo (1992) and Zhang et al. (2023, Appendix C).²⁶ A few clarifications follow:

- For BEHAVIOUR MAXMIN and MIXED MAXMIN, the existence of chance nodes does not change the complexity, unlike the case for PURE MAXMIN:
 - Koller and Megiddo (1992, Proposition 2.6) show that the NP- and coNP-hardness in Table 5 also hold for EFGs of no chance; our proof for Proposition 4.6 also implies this.
 - Carminati et al. (2024) show that the Σ_2^P - and Δ_2^P -hardness in Table 5 also hold for EFGs of no chance.

25. Except for one case: the DP-hardness of =MAXMIN is not known to hold for Boolean EFGs of no chance in which MAX has MA-PR and MIN has PI; see the remark after Proposition 7.6.

26. In the literature on team games (Zhang et al., 2023, for instance), behaviour maxmin and mixed maxmin are called *TME* (team maxmin equilibrium) and *TMECor* (where “Cor” stands for “correlation”), respectively.

- The non-polynomial upper bounds in Table 5 are proven by Zhang et al. (2023, Appendix C) for the *promise problems* that concern the ε -approximation of the maxmin value, where ε is a rational number given as input.²⁷; this is to circumvent the difficulty caused by the possibility that the maxmin value may not be a rational number. Gimbert et al. (2020) show that the exact version is more difficult: $\exists\mathbb{R}$ -complete for EFGs with multi-agent perfect recall for MAX and no MIN, SQUARE-ROOT-SUM-hard when both MAX and MIN are multi-agent.

Perspectives The main perspective for this work is to study these decision problems related to the maxmin value with finer-grained notions of complexity. In particular, we would like to investigate the complexity of these problems in the setting of parametrised complexity. Natural parameters for a game are, for instance, the maximal number of alternations between MAX and MIN along any branch; the branching factor; the maximal size of an information set; the number of distinct utility values. For some of these parameters, we already know that the problems will be hard under the fixed-parameter setting as well, since they are already hard for constant values of the parameters (number of alternations, number of utility values).

It is also of interest to continue and extend our study to more general settings. More concretely, we may add two more dimensions to the complexity tables, allowing for the possibility of general-sum and more than two teams, by adopting Stackelberg strategies as the universal solution concept for all the cases. For the purpose of reference, in our notation, the set of *Stackelberg strategies* or *optimal strategies to commit to* for MAX can be defined to be the set

$$\arg \max_{\zeta_+ \in \mathcal{S}_+} \max_{s_- \in \text{BR}(\zeta_+)} \mathcal{U}_+(\zeta_+, s_-),$$

where \mathcal{S}_+ is the set of pure or mixed strategies of MAX depending on the setting, and $\text{BR}(\zeta_+)$ is the set of MIN's best responses to ζ_+ in \mathcal{S}_-^P according to MIN's utility function u_- .²⁸

Finally, an important direction for future work is to investigate the link between our framework of oracle games and various classes of the game description language GDL. Drawing such connections would allow giving a complexity picture for games compactly represented in GDL (possibly with additional information, for instance a known horizon).

Acknowledgments We are deeply grateful to anonymous reviewers for many suggestions that have helped us to improve the paper, including, in particular, additional references on pure strategies and description languages for games, and the simpler proof of Proposition 6.5. We also thank Brian Hu Zhang for clarifying a few points about the notions and proofs in their paper; our discussion has been conducive to establishing Table 5.

Appendix A. Proofs

A.1 Compiling away Non-Boolean Payoffs

Some proofs in this work involve EFGs of chance with non-Boolean but integer payoffs. We will show that these payoffs can always be compiled into Boolean ones with the help of chance nodes.

27. Promise problem is a generalisation of the notion of decision problem; see the survey by Goldreich (2005) for an overview.

28. The idea behind this definition is that MAX commits to a certain strategy that they communicate to MIN, then MIN plays to the best of MIN's own interest; hence, MIN picks a best response to MAX's committed strategy; the max operator over $\text{BR}(\zeta_+)$ is due to the common assumption that MIN breaks ties in MAX's favour (Conitzer & Sandholm, 2006). It follows from the definition that Stackelberg strategies coincide with maxmin strategies in zero-sum games.




Figure 3: A Boolean chance tree that represents the value $5/8$.

Lemma A.1. *Let G be a two-player timeable EFG of chance with only integer payoffs. Then there is a Boolean timeable EFG G' and two constants $C > 0$ and D , constructible in time polynomial in the size of G , such that*

- *Every player has the same degree of imperfect information and set of strategies in G' as in G .*
- *For all strategies ζ_+ and ζ_- (whether pure, mixed, or behaviour), the expected utilities $\mathcal{U}_+(\zeta_+, \zeta_-)$ in G and $\mathcal{U}'_+(\zeta_+, \zeta_-)$ in G' satisfy $\mathcal{U}'_+(\zeta_+, \zeta_-) = C \times \mathcal{U}_+(\zeta_+, \zeta_-) + D$.*

Proof. We first transform the payoffs into rational numbers in the interval $[0, 1]$ with bounded binary expansion. Notice that every affine transformation with positive scaling for all payoffs induces the same affine transformation on the maxmin value. Let $n \geq 0$ be the largest payoff in G in terms of absolute value, which means all payoffs in G are integers in $[-n, n]$. Let $d \in \mathbb{N}$ be the smallest integer such that $2^d \geq 2n$. Notice that the value of d is linearly bounded by the size of G . By adding n to all payoffs in G , we shift all payoffs into integers in $[0, 2n]$; then by dividing all payoffs by 2^d , we transform the payoffs into rational numbers in the interval $[0, 1]$ with a binary expansion over at most d digits. If MAX's expected payoff in the initial game G under a strategy profile (ζ_+, ζ_-) is v , then MAX's expected payoff under the same profile in the new game G' is $v' = (v + n)/2^d = Cv + D$ with $C = 1/2^d$ and $D = n/2^d$.

Now, we show how to get rid of the fractional payoffs in G' with the help of chance nodes. Let $x = i/2^d$ with $0 \leq i \leq 2^d$. Then a leaf with a payoff of x can be replaced by a chance tree T_x of depth d that contains only chance nodes with uniform Bernoulli distribution and Boolean payoffs in the following way: if the d -digit binary representation of i reads $i_1 i_2 \dots i_d$, then the right node of the chance tree at depth $1 \leq j \leq d$ has value i_j , and the left node at depth d has value 0. See Figure 3 for an example with $x = 5/8$ (i.e. $d = 3$ and $i = 5$ with binary representation 101); all internal nodes are binary chance nodes that lead to either child with probability $1/2$.

By construction, T_x for $x = i/2^d$ is of size $O(d)$, and the root of T_x has an expected value of x . Since there is no decision node of MAX nor of MIN in T_x , one can replace a leaf with a payoff of x in G' by T_x without changing the expected payoff \mathcal{U} under any strategy profile of the players nor their degree of information. Finally, timeability is clearly preserved from G to G' . \square

A.2 Tseitin Transformation for Compact Boolean Games

Lemma A.2. *Let $\gamma = \langle X, P, D, \varphi_+ \rangle$ be a CBG. Then there exists a CBG $\gamma' := \langle X', P', D', \varphi'_+ \rangle$ with φ'_+ in CNF (respectively DNF, ROBDD), which has the same maxmin value as γ , and such that (i) γ' is of no chance if and only if γ is of no chance; (ii) MAX and MIN have the same degree of imperfect information in γ' as in γ . Furthermore, γ' can be constructed from γ in polynomial time.*

Proof. Let $\exists y_1 \dots \exists y_p \psi_+(X \cup \{y_1, \dots, y_p\})$ be a formula which is satisfied by an assignment to X if and only if φ_+ outputs 1 on this assignment, and in which ψ_+ is in CNF, y_1, \dots, y_p are variables

not in X and p is polynomial in the size of φ_+ . Such a formula can be obtained from φ_+ in polynomial time using the Tseitin transformation (Tseitin, 1983).

To build a CBG γ' with a CNF goal from γ , we define $X' := X \cup \{y_1, \dots, y_n\}$ (with the ordering from X then y_1, \dots, y_n), $P'(x) := P(x)$ for $x \in X$ and $P'(y_i) := +$ for $i = 1, \dots, p$, $D'(x) := D(x)$ for $x \in X$ and $D'(y_i) := X \cup \{y_1, \dots, y_{i-1}\}$ for $i = 1, \dots, p$, and finally $\varphi'_+ := \psi_+$. Observe that the construction is polynomial time, does not introduce any chance variable, and preserves PI, PR, and MA-PR for both players (since all information is revealed at the extra nodes). We now show that γ' has the same maxmin value as γ .

Given a strategy ζ_+ for MAX in γ , we simply define ζ'_+ for MAX in γ' to extend ζ_+ by choosing y_1, \dots, y_p such that ψ_+ is true at all leaves of γ at which φ_+ outputs 1, and arbitrary values at the other leaves. This can be done because in γ' , all information is revealed to MAX at the leaves of γ . Moreover, ζ'_+ is pure if and only if ζ_+ is pure. Finally, by construction, this ensures that the utility of ζ'_+ in γ' is the same as the utility of ζ_+ in γ , against any strategy of MIN (observe that the set of strategies for MIN is the same in γ and γ'), and hence that the maxmin value is preserved.

For DNF, we use the dual construction. Given a CBG with an arbitrary goal φ_+ , we first build a formula $\exists y_1 \dots \exists y_p \psi_+(X \cup \{y_1, \dots, y_p\})$ as above, where ψ_+ is in CNF, but which is satisfied by an assignment to X if and only if φ_+ outputs 0 on this assignment; such a formula can be obtained from $\neg\varphi_+$ using the Tseitin transformation. Then we build a DNF χ_+ equivalent to $\neg\psi_+$ using De Morgan's laws. Finally, we define γ' by $X' := X \cup \{y_1, \dots, y_p\}$, $P'(x) := P(x)$ for $x \in X$ and $P'(y_i) := -$ for $i = 1, \dots, p$, D' as in the construction for CNF, and finally $\varphi'_+ := \chi_+$. Then MAX has the same set of strategies in γ' as in γ . Now given a strategy ζ_- for MIN in γ , we define ζ'_- for MIN in γ' to extend ζ_- by choosing y_1, \dots, y_p such that χ_+ is false at all leaves of γ at which φ_+ outputs 0, and arbitrary values at other leaves; indeed, if and only if φ_+ outputs 0, then by construction there exists y_1, \dots, y_p such that ψ_+ is true, that is, such that χ_+ is false. It follows that MIN can force a loss for MAX in γ' exactly at those leaves in which MAX loses in γ , as desired.

Finally, for ROBDD, we use the construction given by Darwiche and Marquis (2002, middle of page 258): for a given DNF φ_+ , the construction introduces auxiliary variables y_1, \dots, y_p not in X , and builds in polynomial time an ROBDD ψ_+ over $X \cup \{y_1, \dots, y_p\}$ such that $\exists y_1 \dots \exists y_p \psi_+(X \cup \{y_1, \dots, y_p\})$ is equivalent to φ_+ . Hence, with the same reasoning as for CNF above, we conclude that a CBG with a DNF goal can be reformulated in polynomial time into a CBG with an ROBDD goal. By the previous paragraph, it follows that any CBG can be reformulated into a CBG with an ROBDD goal. \square

References

- Albrecht, S. V., & Stone, P. (2018). Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258, 66–95.
- Apt, K. R., & Grädel, E. (Eds.). (2011). *Lectures in Game Theory for Computer Scientists*. Cambridge University Press.
- Arora, S., & Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge University Press.
- Babai, L., Fortnow, L., & Lund, C. (1991). Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1, 3–40.

- Ballard, B. W. (1983). The ϵ -minimax search procedure for trees containing chance nodes. *Artificial Intelligence*, 21(3), 327–350.
- Basilico, N., Celli, A., Nittis, G. D., & Gatti, N. (2017). Team-maxmin equilibrium: Efficiency bounds and algorithms. In Singh, S., & Markovitch, S. (Eds.), *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI 2017)*, pp. 356–362. AAAI Press.
- Basilico, N., Gatti, N., & Amigoni, F. (2012). Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial Intelligence*, 184-185, 78–123.
- Beigel, R. (1991). Bounded queries to sat and the boolean hierarchy. *Theoretical Computer Science*, 84(2), 199–223.
- Bernstein, D. S., Givan, R., Immerman, N., & Zilberstein, S. (2002). The complexity of decentralized control of markov decision processes. *Mathematics of Operations Research*, 27(4), 819–840.
- Bonanno, G. (2018). *Game Theory*. Kindle Direct Publishing.
- Bonnet, É., & Saffidine, A. (2014). On the complexity of general game playing. In Cazenave, T., Winands, M. H. M., & Björnsson, Y. (Eds.), *Proc. 3rd Workshop on Computer Games (CGW@ECAI 2014)*, Vol. 504 of *Communications in Computer and Information Science*, pp. 90–104. Springer.
- Brafman, R. I., Shani, G., & Zilberstein, S. (2013). Qualitative planning under partial observability in multi-agent domains. In desJardins, M., & Littman, M. L. (Eds.), *Proc. 27th AAAI Conference on Artificial Intelligence (AAAI 2013)*. AAAI Press.
- Carminati, L., Zhang, B. H., Cacciamani, F., Li, J., Farina, G., Gatti, N., & Sandholm, T. (2024). Efficient representations for team and imperfect-recall equilibrium computation. Submission under review.
- Celli, A., & Gatti, N. (2018). Computational results for extensive-form adversarial team games. In McIlraith, S. A., & Weinberger, K. Q. (Eds.), *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI 2018)*, pp. 965–972. AAAI Press.
- Chatterjee, K., Doyen, L., & Henzinger, T. A. (2013). A survey of partial-observation stochastic parity games. *Formal Methods in System Design*, 43(2), 268–284.
- Chatterjee, K., & Henzinger, T. A. (2012). A survey of stochastic ω -regular games. *Journal of Computer and System Sciences*, 78(2), 394–413.
- Conitzer, V., & Sandholm, T. (2006). Computing the optimal strategy to commit to. In Feigenbaum, J., Chuang, J. C., & Pennock, D. M. (Eds.), *Proc. 7th ACM Conference on Electronic Commerce (EC 2006)*, pp. 82–90. ACM.
- Conitzer, V., & Sandholm, T. (2008). New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2), 621–641.
- Darwiche, A., & Marquis, P. (2002). A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17, 229–264.
- Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1), 195–259.

- Dhami, S. (2019). *The Foundations of Behavioral Economic Analysis Volume IV: Behavioral Game Theory*. Oxford University Press.
- Farina, G., Celli, A., Gatti, N., & Sandholm, T. (2018). Ex ante coordination and collusion in zero-sum multi-player extensive-form games. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., & Garnett, R. (Eds.), *Proc. Advances in Neural Information Processing Systems (NeurIPS 2018)*, pp. 9661–9671.
- Farina, G., Celli, A., Gatti, N., & Sandholm, T. (2021). Connecting optimal ex-ante collusion in teams to extensive-form correlation: Faster algorithms and positive complexity results. In Meila, M., & Zhang, T. (Eds.), *Proc. 38th International Conference on Machine Learning (ICML 2021)*, Vol. 139 of *Proceedings of Machine Learning Research*, pp. 3164–3173. PMLR.
- Frank, I., & Basin, D. A. (2001). A theoretical and empirical investigation of search in imperfect information games. *Theoretical Computer Science*, 252(1-2), 217–256.
- Genesereth, M. R., & Thielscher, M. (2014). *General Game Playing*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Gilboa, I., & Zemel, E. (1989). Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1), 80–93.
- Gimbert, H., Paul, S., & Srivathsan, B. (2020). A bridge between polynomial optimization and games with imperfect recall. In Seghrouchni, A. E. F., Sukthankar, G., An, B., & Yorke-Smith, N. (Eds.), *Proc. 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, pp. 456–464. IFAAMAS.
- Goldreich, O. (2005). On promise problems (a survey in memory of Shimon Even [1935-2004]). *Electronic Colloquium on Computational Complexity*, TR05-018.
- Goldsmith, J., & Mundhenk, M. (2007). Competition adds complexity. In Platt, J. C., Koller, D., Singer, Y., & Roweis, S. T. (Eds.), *Proc. Advances in Neural Information Processing Systems (NeurIPS 2007)*, pp. 561–568. Curran Associates, Inc.
- Gottlob, G., Greco, G., & Scarcello, F. (2005). Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research*, 24, 357–406.
- Hansen, K. A., Miltersen, P. B., & Sørensen, T. B. (2007). Finding equilibria in games of no chance. In Lin, G. (Ed.), *Proc. 13th Annual International Conference on Computing and Combinatorics (COCOON 2007)*, Vol. 4598 of *Lecture Notes in Computer Science*, pp. 274–284. Springer.
- Hemachandra, L. A. (1989). The strong exponential hierarchy collapses. *Journal of Computer and System Sciences*, 39(3), 299–322.
- Iida, H., Uiterwijk, J. W. H. M., van den Herik, H. J., & Herschberg, I. S. (1993). Potential applications of opponent-model search, part 1: The domain of applicability. *Journal of the International Computer Games Association*, 16(4), 201–208.
- Iida, H., Uiterwijk, J. W. H. M., van den Herik, H. J., & Herschberg, I. S. (1994). Potential applications of opponent-model search, part 2: Risks and strategies. *Journal of the International Computer Games Association*, 17(1), 10–14.
- Isbell, J. (1957). Finitary games. *Contributions to the Theory of Games*, 3, 79–96.

- Jakobsen, S. K., Sørensen, T. B., & Conitzer, V. (2016). Timeability of extensive-form games. In Sudan, M. (Ed.), *Proc. ACM Conference on Innovations in Theoretical Computer Science (ITCS 2016)*, pp. 191–199. ACM.
- Kaiser, L., & Stafiniak, L. (2011). First-order logic with counting for general game playing. In Burgard, W., & Roth, D. (Eds.), *Proc. 25th AAAI Conference on Artificial Intelligence (AAAI 2011)*, pp. 791–796. AAAI Press.
- Knuth, D. E., & Moore, R. W. (1975). An analysis of alpha-beta pruning. *Artificial Intelligence*, 6(4), 293–326.
- Kochenderfer, M. J. (2015). *Decision Making Under Uncertainty: Theory and Application*. MIT Press.
- Kochenderfer, M. J., Wheeler, T. A., & Wray, K. H. (2022). *Algorithms for Decision Making*. MIT Press.
- Koller, D., & Megiddo, N. (1992). The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4(4), 528–552.
- Koller, D., Megiddo, N., & von Stengel, B. (1996). Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14(2), 247–259.
- Kovacs, D. L. (2012). A multi-agent extension of pddl3.1. In *Proc. 3rd Workshop on the International Planning Competition (IPC), 22nd International Conference on Automated Planning and Scheduling (ICAPS-2012)*, pp. 19–27. ICAPS.
- Kuhn, H. W. (1953). Extensive games and the problem of information. In *Contributions to the Theory of Games (AM-28), Volume II*, pp. 193–216. Princeton University Press, Princeton.
- Letchford, J., & Conitzer, V. (2010). Computing optimal strategies to commit to in extensive-form games. In Parkes, D. C., Dellarocas, C., & Tennenholtz, M. (Eds.), *Proc. 11th ACM Conference on Electronic Commerce (EC 2010)*, pp. 83–92. ACM.
- Li, J., Zanuttini, B., & Ventos, V. (2024). Combinatorial games with incomplete information. In Larson, K. (Ed.), *Proc. 33rd International Joint Conference on Artificial Intelligence (IJCAI 2024)*, pp. 2895–2903. ijcai.org.
- Lipton, R. J., Markakis, E., & Mehta, A. (2003). Playing large games using simple strategies. In Menascé, D. A., & Nisan, N. (Eds.), *Proc. 4th ACM Conference on Electronic Commerce (EC-2003)*, pp. 36–41. ACM.
- Madani, O., Hanks, S., & Condon, A. (2003). On the undecidability of probabilistic planning and related stochastic optimization problems. *Artificial Intelligence*, 147(1-2), 5–34.
- Maschler, M., Solan, E., & Zamir, S. (2020). *Game Theory* (2nd edition). Cambridge University Press.
- McCarthy, S. M., Laan, C. M., Wang, K., Vayanos, P., Sinha, A., & Tambe, M. (2018). The price of usability: Designing operationalizable strategies for security games. In Lang, J. (Ed.), *Proc. 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pp. 454–460. ijcai.org.
- McMahan, H. B., Gordon, G. J., & Blum, A. (2003). Planning in the presence of cost functions controlled by an adversary. In Fawcett, T., & Mishra, N. (Eds.), *Proc. 20th International Conference on Machine Learning (ICML 2003)*, pp. 536–543. AAAI Press.

- Mundhenk, M., Goldsmith, J., Lusena, C., & Allender, E. (2000). Complexity of finite-horizon Markov decision process problems. *Journal of the ACM*, 47(4), 681–720.
- Myerson, R. B. (1997). *Game theory: Analysis of Conflict*. Harvard University Press.
- Papadimitriou, C. H. (1994). *Computational Complexity*. Addison-Wesley.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4), 1029–1050.
- Perea, A. (2012). *Epistemic Game Theory: Reasoning and Choice*. Cambridge University Press.
- Peterson, G., Reif, J., & Azhar, S. (2001). Lower bounds for multiplayer noncooperative games of incomplete information. *Computers & Mathematics with Applications*, 41(7), 957–992.
- Piccione, M., & Rubinstein, A. (1997). On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior*, 20(1), 3–24.
- Pita, J., Jain, M., Marecki, J., Ordóñez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., & Kraus, S. (2008). Deployed ARMOR protection: the application of a game theoretic model for security at the los angeles international airport. In Berger, M., Burg, B., & Nishiyama, S. (Eds.), *Proc. 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008), Industry and Applications Track*, pp. 125–132. IFAAMAS.
- Rebstock, D., Solinas, C., Buro, M., & Sturtevant, N. R. (2019). Policy based inference in trick-taking card games. In *Proc. IEEE Conference on Games (CoG 2019)*, pp. 1–8. IEEE.
- Richards, M., & Amir, E. (2012). Information set generation in partially observable games. In Hoffmann, J., & Selman, B. (Eds.), *Proc. 26th AAAI Conference on Artificial Intelligence (AAAI 2012)*, pp. 549–555. AAAI Press.
- Rintanen, J. (2004). Complexity of planning with partial observability. In Zilberstein, S., Koehler, J., & Koenig, S. (Eds.), *Proc. 14th International Conference on Automated Planning and Scheduling (ICAPS 2004)*, pp. 345–354. AAAI Press.
- Schaefer, M., & Umans, C. (2002). Completeness in the polynomial-time hierarchy: A compendium. *SIGACT news*, 33(3), 32–49.
- Schlenker, A., Thakoor, O., Xu, H., Fang, F., Tambe, M., Tran-Thanh, L., Vayanos, P., & Vorobeychik, Y. (2018). Deceiving cyber adversaries: A game theoretic approach. In André, E., Koenig, S., Dastani, M., & Sukthankar, G. (Eds.), *Proc. 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2018)*, pp. 892–900. IFAAMAS.
- Schwarzentruber, F. (2019). The complexity of tiling problems. *CoRR*, abs/1907.00102.
- Shoham, Y., & Leyton-Brown, K. (2009). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- Stockmeyer, L. J., & Meyer, A. R. (1973). Word problems requiring exponential time: Preliminary report. In Aho, A. V., Borodin, A., Constable, R. L., Floyd, R. W., Harrison, M. A., Karp, R. M., & Strong, H. R. (Eds.), *Proc. 5th Annual ACM Symposium on Theory of Computing (STOC 1973)*, pp. 1–9. ACM.
- Sturtevant, N. R., Zinkevich, M., & Bowling, M. H. (2006). Prob-maxn: Playing n-player games with opponent models. In *Proc. 21st National Conference on Artificial Intelligence (AAAI 2006)*, pp. 1057–1063. AAAI Press.

- Tewolde, E., Oesterheld, C., Conitzer, V., & Goldberg, P. W. (2023). The computational complexity of single-player imperfect-recall games. In *Proc. 32nd International Joint Conference on Artificial Intelligence (IJCAI 2023)*, pp. 2878–2887. ijcai.org.
- Tewolde, E., Zhang, B. H., Oesterheld, C., Zampetakis, M., Sandholm, T., Goldberg, P., & Conitzer, V. (2024). Imperfect-recall games: Equilibrium concepts and their complexity. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, pp. 2994–3004. ijcai.org.
- Tseitin, G. S. (1983). On the complexity of derivation in propositional calculus. In Siekmann, J., & Wrightson, G. (Eds.), *Automation of reasoning: 2: Classical papers on computational logic 1967–1970*, pp. 466–483. Springer.
- van Damme, E. (1991). *Stability and Perfection of Nash Equilibria*. Springer Berlin Heidelberg.
- van Emde Boas, P. (1997). The convenience of tilings. In *Complexity, Logic, and Recursion Theory*, pp. 331–363. CRC Press.
- von Neumann, J., & Morgenstern, O. (1994). *Theory of Games and Economic Behavior*. Princeton University Press.
- von Stackelberg, H. (1934). *Marktform und Gleichgewicht*. Springer.
- von Stengel, B. (1996). Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2), 220–246.
- von Stengel, B., & Forges, F. (2008). Extensive-form correlated equilibrium: Definition and computational complexity. *Mathematics of Operations Research*, 33(4), 1002–1022.
- von Stengel, B., & Koller, D. (1997). Team-maxmin equilibria. *Games and Economic Behavior*, 21(1), 309–321.
- Yun, B., Vesic, S., & Oren, N. (2022). Representing pure Nash equilibria in argumentation. *Argument & Computation*, 13(2), 195–208.
- Zhang, B. H., Farina, G., & Sandholm, T. (2023). Team belief DAG: generalizing the sequence form to team games for fast computation of correlated team max-min equilibria via regret minimization. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., & Scarlett, J. (Eds.), *Proc. International Conference on Machine Learning (ICML 2023)*, Vol. 202 of *Proceedings of Machine Learning Research*, pp. 40996–41018. PMLR.
- Zhang, B. H., & Sandholm, T. (2022). Team correlated equilibria in zero-sum extensive-form games via tree decompositions. In *Proc. 36th AAAI Conference on Artificial Intelligence (AAAI 2022)*, pp. 5252–5259. AAAI Press.
- Zhang, Y., An, B., & Cerný, J. (2021). Computing ex ante coordinated team-maxmin equilibria in zero-sum multiplayer extensive-form games. In *Proc. 35th AAAI Conference on Artificial Intelligence (AAAI 2021)*, pp. 5813–5821. AAAI Press.