



Analysis of Hannan consistent selection for Monte Carlo tree search in simultaneous move games

Vojtěch Kovařík¹ · Viliam Lisý¹

Received: 21 August 2017 / Revised: 7 July 2019 / Accepted: 17 July 2019 / Published online: 25 July 2019
© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2019

Abstract

Hannan consistency, or no external regret, is a key concept for learning in games. An action selection algorithm is Hannan consistent (HC) if its performance is eventually as good as selecting the best fixed action in hindsight. If both players in a zero-sum normal form game use a Hannan consistent algorithm, their average behavior converges to a Nash equilibrium of the game. A similar result is known about extensive form games, but the played strategies need to be Hannan consistent with respect to the counterfactual values, which are often difficult to obtain. We study zero-sum extensive form games with simultaneous moves, but otherwise perfect information. These games generalize normal form games and they are a special case of extensive form games. We study whether applying HC algorithms in each decision point of these games directly to the observed payoffs leads to convergence to a Nash equilibrium. This learning process corresponds to a class of Monte Carlo Tree Search algorithms, which are popular for playing simultaneous-move games but do not have any known performance guarantees. We show that using HC algorithms directly on the observed payoffs is not sufficient to guarantee the convergence. With an additional averaging over joint actions, the convergence is guaranteed, but empirically slower. We further define an additional property of HC algorithms, which is sufficient to guarantee the convergence without the averaging and we empirically show that commonly used HC algorithms have this property.

Keywords Nash equilibrium · Extensive form games · Simultaneous moves · Zero sum · Hannan consistency

Editor: Prasad Tadepalli.

✉ Vojtěch Kovařík
vojta.kovarik@gmail.com

Viliam Lisý
viliam.lisy@agents.fel.cvut.cz

¹ Artificial Intelligence Center, Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague, Zikova 1903/4, 166 36 Prague 6, Czech Republic

1 Introduction

Research on learning in games led to multiple algorithmic advancements, such as the variants of the counterfactual regret minimization algorithm (Zinkevich et al. 2007), which allowed for achieving human performance in poker (Moravčík et al. 2017; Brown and Sandholm 2018). Learning in games has been extensively studied in the context of normal form games, extensive form games, as well as Markov games (Fudenberg et al. 1998; Littman 1994).

One of the key concepts in learning in games is Hannan consistency (Hannan 1957; Hart and Mas-Colell 2000; Cesa-Bianchi and Lugosi 2006), also known as no external regret. An algorithm for repetitively selecting actions from a fixed set of options is Hannan consistent (HC), if its performance approaches the performance of selecting the best fixed option all the time. This property was first studied in normal form games known to the players, but it has later been shown to be achievable also if the algorithm knows only its own choices and the resulting payoffs (Auer et al. 2003).

If both players use a Hannan consistent algorithm in self-play in a zero-sum normal form game, the empirical frequencies of their action choices converge to a Nash equilibrium (NE) of the game¹ (Blum and Mansour 2007). A similar convergence guarantee can be provided for zero-sum imperfect information extensive form games (Zinkevich et al. 2007). However, it requires the algorithms selecting actions in each decision point to be Hannan consistent with respect to counterfactual values. Computing these values requires either an expensive traversal of all states in the game or a sampling scheme with reciprocal weighting (Lanctot et al. 2009) leading to high variance and slower learning theoretically (Gibson et al. 2012) and in larger games even practically (Bošanský et al. 2016). Therefore, using them should be avoided in simpler classes of games, where they are not necessary.

In this paper, we study zero-sum extensive form games with simultaneous moves, but otherwise perfect information (SMGs). These games generalize normal form games but are a special case of extensive form games. Players in SMGs simultaneously choose an action to play, after which the game either ends and the players receive a payoff, or the game continues to a next stage, which is again an SMG. These games are extensively studied in game playing literature (see Bošanský et al. 2016 for a survey). They are also closely related to Markov games (Littman 1994), but they do not allow cycles in the state space and the players always receive a reward only at the end of the game. We study whether applying HC algorithms in each decision point of an SMG directly to the observed payoffs leads to convergence to a Nash equilibrium as in normal form games, or whether it is necessary to use additional assumptions about the learning process.

Convergence properties of Hannan consistent algorithms in simultaneous-move games are interesting also because using a separate HC algorithm in each decision point directly corresponds to popular variants of Monte Carlo Tree Search (MCTS) (Coulom 2007; Kocsis and Szepesvári 2006) commonly used in SMGs. In SMGs, these algorithms have been used for playing card games (Teytaud and Flory 2011; Lanctot et al. 2014), variants of simple computer games (Perick et al. 2012), and in the most successful agents for General Game Playing (Finnsson and Björnsson 2008). However, there are no known theoretical guarantees for their performance in this class of games. The main goal of the paper is to understand better the conditions under which MCTS-like algorithms perform well when applied to SMGs.

¹ All of the games discussed in the paper might have multiple Nash equilibria, and regret minimization causes the empirical strategy to approach the set of these equilibria. For brevity, we will write that a strategy “converges to a NE” to express that there is eventually always some (but not necessarily always the same) NE close to the strategy.

1.1 Contributions

We focus on two player zero-sum extensive form games with simultaneous moves but otherwise perfect information. We denote a class of algorithms that learn a strategy from a sequence of simulations of self-play in these games as SM-MCTS. These algorithms use simple *selection policies* to choose next action to play in each decision point based on statistics of rewards achieved after individual actions. We study whether Hannan consistency of the selection policies is sufficient to guarantee convergence of SM-MCTS to a Nash equilibrium of the game. We also investigate whether a similar relation holds between the approximate versions of Hannan consistency and Nash equilibria.

In Theorem 1, we prove a negative result, which is in our opinion the most surprising and interesting part of the paper, which is the fact that Hannan consistency alone does not guarantee a good performance of the standard SM-MCTS. We present a Hannan consistent selection policy that causes SM-MCTS to converge to a solution far from the equilibrium.

We show that this could be avoided by a slight modification to SM-MCTS (which we call SM-MCTS-A), which updates the selection policies by the average reward obtained after a joined action of both players in past simulations, instead of the result of the current simulation. In Theorem 2, we prove that SM-MCTS-A combined with any Hannan consistent (resp. approximately HC) selection policy with guaranteed exploration converges to a subgame-perfect Nash equilibrium (resp. an approximate subgame-perfect NE) in this class of games. For selection policies which are only approximately HC, we present bounds on the eventual distance from a Nash equilibrium.

In Theorem 3, we show that under additional assumptions on the selection policy, even the standard SM-MCTS can guarantee the convergence. We do this by defining the property of having unbiased payoff observations (UPO), and showing that it is a sufficient condition for the convergence. We then empirically confirm that the two commonly used Hannan consistent algorithms, Exp3 and regret matching, satisfy this property, thus justifying their use in practice. We further investigate the empirical speed of convergence and show that the eventual distance from the equilibrium is typically much better than the presented theoretical guarantees. We empirically show that SM-MCTS generally converges as close to the equilibrium as SM-MCTS-A, but does it faster.

Finally, we give theoretical grounds for some practical improvements, which may be used with SM-MCTS, but have not been formally justified. These include removal of exploration samples from the resulting strategy and the use of average strategy instead of empirical frequencies of action choices. All presented theoretical results trivially apply also to perfect information games with sequential moves. The counterexample in Theorem 1, on the other hand, applies to Markov and general extensive form games.

1.2 Article outline

In Sect. 2 we describe simultaneous-move games, the standard class of SM-MCTS algorithms, and our modification SM-MCTS-A. We then describe the multi-armed bandit problem, the definition of Hannan consistency, and explain two of the common Hannan consistent bandit algorithms (Exp3 and regret matching). We also explain the relation of the studied problem with counterfactual regret minimization, multi-agent reinforcement learning, and shortest path problems.

In Sect. 3, we explain the relation of SM-MCTS setting with the multi-armed bandit problem. In this section, we also define the technical notation which is used in the proofs of our main results.

In Sect. 4, we provide a counterexample showing that for general Hannan consistent algorithms, SM-MCTS does not necessarily converge. While interesting in itself, this result also hints at which modifications are sufficient to get convergence. The positive theoretical results are presented in Sect. 5. First, we consider the modified SM-MCTS-A and present the asymptotic bound on its convergence rate. We follow by defining the unbiased payoff observations property and proving the convergence of SM-MCTS based on HC selection policies with this property. We then present an example which gives a lower bound on the quality of a strategy to which SM-MCTS(-A) converges.

In Sect. 6, we make a few remarks about which strategy should be considered as the output of SM-MCTS(-A). In Sect. 7, we present an empirical investigation of convergence of SM-MCTS and SM-MCTS-A, as well as empirical evidence supporting that the commonly used HC-algorithms guarantee the UPO property. Finally, Sect. 8 summarizes the results and highlights open questions which might be interesting for future research. Table 1 then recapitulates the notation and abbreviations used throughout the paper.

2 Background

We now introduce the game theory fundamentals and notation used throughout the paper. We define simultaneous-move games, describe the class of SM-MCTS algorithms and its modification SM-MCTS-A, and afterward, we discuss existing selection policies and their properties.

2.1 Simultaneous move games

A finite two-player zero-sum game with perfect information and simultaneous moves can be described by a tuple $G = (\mathcal{H}, \mathcal{L}, \mathcal{A}, \mathcal{T}, u, h_0)$, where \mathcal{H} is a set of inner states and \mathcal{L} denotes the terminal states.² $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ is the set of joint actions of individual players and we denote $\mathcal{A}_1(h) = \{1, \dots, m^h\}$ and $\mathcal{A}_2(h) = \{1, \dots, n^h\}$ the actions available to individual players in state $h \in \mathcal{H}$. The game begins in an initial state h_0 . The transition function $\mathcal{T} : \mathcal{H} \times \mathcal{A}_1 \times \mathcal{A}_2 \mapsto \mathcal{H} \cup \mathcal{L}$ defines the successor state given a current state and actions for both players. For brevity, we denote $\mathcal{T}(h, i, j) \equiv h_{ij}$. We assume that all the sets are finite and that the whole setting is modeled as a tree.³ The payoffs of players 1 and 2 are defined by the utility function $u = (u_1, u_2) : \mathcal{L} \rightarrow [0, 1]^2$. We assume constant-sum games (which are equivalent to zero-sum games): $\forall z \in \mathcal{L}, u_2(z) = 1 - u_1(z)$.

A *matrix game* is a special case of the setting above, where G only has a single stage, h —we have $\mathcal{H} = \{h\}$. In this case we can simplify the notation and represent the game by the matrix $M = \left(v_{ij}^M \right)_{i,j} \in [0, 1]^{m \times n}$, where for $(i, j) \in \mathcal{A}_1 \times \mathcal{A}_2 = \{1, \dots, m\} \times \{1, \dots, n\}$, we have $u_1(h_{ij}) = v_{ij}^M$. In other words, v_{ij}^M corresponds to the payoff received by player 1

² In our analysis, we will assume that the game does not contain chance nodes, where the next node is randomly selected by “nature” according to some probability distribution. We do this purely to reduce the amount of technical notation required—as long as the state-space is finite, all of the results hold even if the chance nodes are present. This can be proven in a straightforward way.

³ That is, if $(h, i, j) \neq (h', i', j')$ then $h_{ij} \neq h_{i'j'}$.

Table 1 The abbreviations and common notation for quick reference

<i>Abbreviations</i>	
SM-MCTS(-A)	(averaged) Simultaneous-move Monte Carlo tree search
NE	Nash equilibrium
MAB	Multi-armed bandit
HC	Hannan consistent
Exp3	Exponential-weight algorithm for exploration and exploitation
RM	Regret matching algorithm
UPO	Unbiased payoff observations
CFR	The counterfactual regret minimization algorithm
EFG	Extensive form game
<i>Notation related to a multi-armed bandit problem (P)</i>	
$x_i(t)$	The reward for action i at time t
$i(t)$	The action chosen at time t
$\sigma, \bar{\sigma}, \hat{\sigma}$	Strategy (at the given time, average, and empirical frequencies)
$G(t), G_{\max}(t)$	The actually obtained (resp. maximum achievable) cumulative reward
$R(t)$	The external regret at time t
g, g_{\max}, r	The averaged variants of G, G_{\max} , and R
γ, ε	Exploration rate, ε -Hannan consistency parameter
<i>Matrix games</i>	
$M = (v_{ij}^M)_{i,j}$	Matrix game
u^M, v^M	Utility in M , minimax value of M
br	Best response
(P^M)	MAB problem corresponding to M
$\tilde{M}(t) = (v_i^{\tilde{M}} j(t))_{ij}$	Repeated matrix game with bounded distortion
$(P^{\tilde{M}}), R^{\tilde{M}}, g^{\tilde{M}}, \dots$	The corresponding MAB problem and related variables
<i>Notation related to simultaneous-move games...</i>	
h, v^h, t^h	State in the game, its value, and the number of visits
$i^h(t), j^h(t)$	Actions selected at h during its t -th visit
t_i^h, t_j^h, t_{ij}^h	Variables related to the number of uses of each (joint) action
D, d_h	Game depth, depth of the sub-tree rooted at h
<i>...and the corresponding MAB problems</i>	
$(P^h), G^h, x^h, \dots$	Player 1 MAB problem corresponding to SM-MCTS at h , related variables
r_1^h, r_2^h	Average regret corresponding to (P^h) , resp. player 2 variant of (P^h)
$h_{ij}, v_{ij}^h, x_{ij}^h, \dots$	State resulting from taking (i, j) at h , corresponding variables v^{hij}, x^{hij}, \dots
(\bar{P}^h)	Player 1 MAB problem corresponding to SM-MCTS-A at h
$\bar{x}_{ij}^h(n), \tilde{x}_{ij}^h(n)$	The arithmetical (resp. weighted) average of values $x_{ij}^h(1), \dots, x_{ij}^h(n)$

if player 1 chooses the row i and player 2 chooses the column j . A strategy $\sigma_p \in \Delta(\mathcal{A}_p)$ is a distribution over the actions in \mathcal{A}_p . If σ_1 is represented as a row vector and σ_2 as a column vector, then the expected value to player 1 when both players play with these strategies is $u_1(\sigma_1, \sigma_2) = \sigma_1 M \sigma_2$. Given a profile $\sigma = (\sigma_1, \sigma_2)$, define the utilities against best response

strategies to be $u_1(br, \sigma_2) = \max_{\sigma'_1 \in \Delta(\mathcal{A}_1)} \sigma'_1 M \sigma_2$ and $u_1(\sigma_1, br) = \min_{\sigma'_2 \in \Delta(\mathcal{A}_2)} \sigma_1 M \sigma'_2$. A strategy profile (σ_1, σ_2) is an ε -Nash equilibrium of the matrix game M if and only if

$$u_1(br, \sigma_2) - u_1(\sigma_1, \sigma_2) \leq \varepsilon \quad \text{and} \quad u_1(\sigma_1, br) - u_1(\sigma_1, \sigma_2) \leq \varepsilon \tag{1}$$

A behavioral strategy for player p is a mapping from states $h \in \mathcal{H}$ to a probability distribution over the actions $\mathcal{A}_p(h)$, denoted $\sigma_p(h)$. Given a profile $\sigma = (\sigma_1, \sigma_2)$, define the probability of reaching a terminal state z under σ as $\pi^\sigma(z) = \pi_1^\sigma(z)\pi_2^\sigma(z)$, where each $\pi_p^\sigma(z)$ is a product of probabilities of the actions taken by player p along the path to z . Define Σ_p to be the set of behavioral strategies for player p . Then for any strategy profile $\sigma = (\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$ we define the expected utility of the strategy profile (for player 1) as

$$u(\sigma) = u(\sigma_1, \sigma_2) = \sum_z \pi^\sigma(z) u_1(z) \tag{2}$$

An ε -Nash equilibrium profile (σ_1, σ_2) in this case is defined analogously to (1). In other words, none of the players can improve their utility by more than ε by deviating unilaterally. If $\sigma = (\sigma_1, \sigma_2)$ is an exact Nash equilibrium (an ε -NE with $\varepsilon = 0$), then we denote the unique value of G as $v^G = v^{h_0} = u(\sigma_1, \sigma_2) = \min_{\sigma'_2} \max_{\sigma'_1} u(\sigma'_1, \sigma'_2)$ (because of the last identity, NE in this setting could also be called “minimax optima”).

A subgame rooted in state $h \in \mathcal{H}$ is the game $(\mathcal{N}, \mathcal{H}, \mathcal{L}, \mathcal{A}, \mathcal{I}, u, h)$. This is the same game as the original one, except that it starts at the node h instead of h_0 . We denote by v^h the value of this subgame. An ε -Nash equilibrium profile σ is called subgame-perfect if σ is ε -Nash equilibrium of the subgame rooted at h for every $h \in \mathcal{H}$ (even at the nodes which are unreachable under this strategy).

Two-player perfect information games with simultaneous moves are sometimes appropriately called stacked matrix games because at every state h there is a joint action set $\mathcal{A}_1(h) \times \mathcal{A}_2(h)$ where each joint action (i, j) either leads to a terminal state with utility $u_1(h_{ij})$ or to a subgame rooted in h_{ij} . This subgame is itself another stacked matrix game and its unique value $v^{h_{ij}}$ can be determined by backward induction (see Fig. 1). Thus finding the optimal strategy at h is the same as finding the optimal strategy in the matrix game $(v^{h_{ij}})_{i,j}$, where $v^{h_{ij}}$ is either $u_1(h_{ij})$ or $v^{h_{ij}}$.




Fig. 1 Example game tree of a game with perfect information and simultaneous moves. Only the leaves contain actual rewards—the values in the inner nodes are achieved by optimal play in the corresponding subtree, they are not part of the definition of the game

2.2 Simultaneous Move Monte Carlo tree search

Monte Carlo Tree Search (MCTS) is a simulation-based state space search algorithm often used in game trees. The main idea is to iteratively run simulations from the current state h_0 until the end of the game, incrementally growing a tree $\mathbf{T} \subset \mathcal{H} \cup \mathcal{L}$ rooted at the current state. In the basic form of the algorithm, \mathbf{T} only contains h_0 and a single leaf is added each iteration. Each iteration starts at h_0 . If MCTS encounters a node $h \in \mathbf{T}$ whose all children in $\mathcal{H} \cup \mathcal{L}$ are already in \mathbf{T} , it uses the statistics maintained at h to select one of its children that it then transitions to. If MCTS encounters a node $h \in \mathbf{T}$ that has children that aren't yet in \mathbf{T} , it adds one of them to \mathbf{T} and transitions to it. Then we apply a rollout policy (for example, a random action selection) from this new leaf of \mathbf{T} to some terminal state of the game. The outcome of the simulation is then returned as a reward to the new leaf and all its predecessors.

In Simultaneous Move MCTS (SM-MCTS), the main difference is that a joint action of both players is selected and used to transition to a following state. The algorithm has been previously applied, for example in the game of Tron (Perick et al. 2012), Urban Rivals (Teytaud and Flory 2011), and in general game-playing (Finnsson and Björnsson 2008). However, guarantees of convergence to a NE remain unknown, and Shafiei et al. (2009) show that the most popular selection policy (UCB) does not converge to a NE, even in a simple one-stage game. The convergence to a NE depends critically on the selection and update policies applied. We describe variants of two popular selection algorithms in Sect. 2.3.

2.2.1 Description of SM-MCTS algorithm

In Algorithm 1, we present a generic template of MCTS algorithms for simultaneous-move games (SM-MCTS). We then proceed to explain how specific algorithms are derived from this template.

Algorithm 1 describes a single iteration of SM-MCTS. \mathbf{T} represents the incrementally built MCTS tree, in which each state is represented by one node. Every node $h \in \mathbf{T}$ maintains algorithm-specific statistics about the iterations that previously used this node. In the terminal states, the algorithm returns the value of the state for the first player (line 2). If the current state has a node in the current MCTS tree \mathbf{T} , the statistics in the node are used to select an action for each player (line 4). These actions are executed (line 5) and the algorithm is called recursively on the resulting state (line 6). The result of this call is used to update the statistics maintained for state h (line 7). If the current state is not stored in tree \mathbf{T} , it is added to the tree (line 10) and

```

SM-MCTS( $h$ —current state of the game)
1: if  $h \in \mathcal{L}$  then // current state is terminal
2:   return  $u_1(h)$ 
3: if  $h \in \mathbf{T}$  then // current state is in memory
4:    $(i, j) \leftarrow \text{Select}(h)$ 
5:    $h' \leftarrow h_{ij}$  // get the next game state
6:    $x \leftarrow \text{SM-MCTS}(h')$ 
7:    $\text{Update}(h, i, j, x)$ 
8:   return  $x$ 
9: else // state is not in memory
10:   $\mathbf{T} \leftarrow \mathbf{T} \cup \{h\}$  // add state to memory
11:   $x \leftarrow \text{Rollout}(h)$  // finish the game randomly

12: return  $x$ 

```

Algorithm 1: Simultaneous Move Monte Carlo Tree Search

its value is estimated using the rollout policy (line 11). The rollout policy is usually uniform random action selection until the game reaches a terminal state, but it can also be based on domain-specific knowledge. Finally, the result of the Rollout is returned to higher levels of the tree.

The template can be instantiated by specific implementations of the updates of the statistics on line 7 and the selection based on these statistics on line 4. Selection policies can be based on many different algorithms, but the most successful ones use algorithms for solving the multi-armed bandit problem introduced in Sect. 2.3. Specifically, we firstly decide on an algorithm A and choose all of its parameters except for the number of available actions. Then we run a separate instance of this algorithm for each node $h \in \mathbf{T}$ and each of the players. In each node, the action for each player is selected based only on the history. The update procedure then uses the values $u_1(h)$ and $1 - u_1(h)$ to update the statistics of players 1 and 2 at h respectively (this way, both the values are in the interval $[0, 1]$).

This work assumes that, except for the SM-MCTS back-propagation, the selection algorithms do not communicate with each other in any way. We refrain from analyzing the (potentially more powerful) algorithms which have access to non-local variables (e.g., a global clock, abstraction-based information sharing) because such analysis would be significantly more complicated, but also because the primary goal of this paper is to understand the limits of the simpler selection policies.

2.2.2 SM-MCTS-A algorithm

We also propose an “averaged” variant of the algorithm, which we denote as SM-MCTS-A (see Algorithm 2). Its main advantage over SM-MCTS is that it guarantees convergence to a NE under much weaker conditions (Theorem 2). Focusing on the differences between these algorithms helps us better understand what is missing to guarantee convergence of SM-MCTS.

SM-MCTS-A works similarly to SM-MCTS, except that during the back-propagation phase, every visited node $h \in \mathbf{T}$ sends back to its parent the average reward obtained from h

```

SM-MCTS-A( $h$ —current state of the game)
1: if  $h \in \mathcal{L}$  then // current state is terminal
2:   return ( $u_1(h), u_1(h)$ )
3: if  $h \in \mathbf{T}$  then // current state is in memory
4:   ( $i, j$ )  $\leftarrow$  Select( $h$ ) // generate actions

5:    $t^h \leftarrow t^h + 1$  // increase visit counter
6:   ( $x, \bar{x}$ )  $\leftarrow$  SM-MCTS( $h_{ij}$ ) // get the latest and average reward

7:   Update( $h, i, j, \bar{x}$ ) // update by the average reward

8:    $G^h \leftarrow G^h + x$  // update cumulative reward
9:   return ( $x, G^h/t^h$ ) // return reward & current state average reward
10: else // state is not in memory
11:    $\mathbf{T} \leftarrow \mathbf{T} \cup \{h\}$  // add state to memory
12:    $t^h, G^h \leftarrow 0$  // initialize visit counter and cumulative reward
13:    $x \leftarrow$  Rollout( $h$ ) // finish the game randomly

14:   return ( $x, x$ )

```

Algorithm 2: SM-MCTS-A, a variant of SM-MCTS which is guaranteed to converge to an equilibrium (under the assumptions of Theorem 2).

so far (denoted G^h/t^h). The parent is then updated by this “averaged” reward G^h/t^h , rather than by the current reward x (line 7). In practice, this is achieved by doing a bit of extra book-keeping—storing the cumulative reward G^h and the number of visits t^h and updating them during each visit of h (lines 5 and 8) and back-propagating both the current and average rewards (line 9).

We note that in our previous work (Lisý et al. 2013) we prove a result similar to Theorem 2 here. However, the algorithm that we used earlier is different from SM-MCTS-A algorithm described here. In particular, SM-MCTS-A uses averaged values for decision making in each node but propagates backward the non-averaged values (unlike the previous version, which also updates the selection algorithm based on the averaged values, but then it propagates backward these averaged numbers—and on the next level, it takes averages of averages and so on). Consequently, this new version is much closer to the non-averaged SM-MCTS used in practice, and it has faster empirical convergence.

2.3 Multi-armed bandit problem

The multi-armed bandit (MAB) problem is one of the basic models in online learning. It often serves as the basic model for studying fundamental trade-offs between exploration and exploitation in an unknown environment (Auer et al. 1995, 2002). In practical applications, the algorithms developed for this model have recently been used in online advertising (Pandey et al. 2007), convex optimization (Flaxman et al. 2005), and, most importantly for this paper, in Monte Carlo tree search algorithms (Kocsis and Szepesvári 2006; Browne et al. 2012; Gelly and Silver 2011; Teytaud and Flory 2011; Coulom 2007). More details can be found in an extensive survey of the field by Bubeck and Cesa-Bianchi (2012).

Definition 1 (*Adversarial multi-armed bandit problem*) Multi-armed bandit problem is specified by a set of actions $\{1, \dots, K\}$ and a sequence of reward vectors $x(1), x(2), \dots$, where $x(t) = (x_i(t))_{i=1}^K \in [0, 1]^K$ for each $t \in \mathbb{N}$. In each time step, an agent selects an action $i(t)$ and receives the payoff $x_{i(t)}(t)$.

Note that the agent does not observe the values $x_i(t)$ for $i \neq i(t)$. There are many special cases and generalizations of this setting (such as the stochastic bandit problem), however, in this paper, we only need the following variant of this concept:

The (adaptive) *adversarial MAB problem* (P) is identical to the setting above, except that each $x_i(t+1)$ is a random variable that might depend on $x(1), \dots, x(t)$ and $i(1), \dots, i(t)$.

A bandit algorithm is any procedure which takes as an input the number $K \in \mathbb{N}$, the sequence of actions $i(1), \dots, i(t)$ played thus far and the rewards received as a result and returns the next action $i(t+1) \in \{1, \dots, K\}$ to be played.

Bandit algorithms usually attempt to optimize their behavior with respect to some notion of regret. Intuitively, their goal is the minimization of the difference between playing the strategy given by the algorithm and playing some baseline strategy, which can use information not available to the agent. For example, the most common notion of regret is the *external regret*, which is the difference between playing according to the prescribed strategy and playing the fixed optimal action all the time.

Definition 2 (*External Regret*) The *external regret* for playing a sequence of actions $i(1), i(2), \dots, i(t)$ in (P) is defined as

$$R(t) = \max_{i=1, \dots, K} \sum_{s=1}^t x_i(s) - \sum_{s=1}^t x_{i(s)}(s).$$

By $r(t)$ we denote the average external regret $r(t) := \frac{1}{t} R(t)$.

2.4 Hannan consistent algorithms

A desirable goal for any bandit algorithm is the classical notion of Hannan consistency. Having this property means that for high enough t , the algorithm performs nearly as well as it would if it played the optimal constant action since the beginning.

Definition 3 (*Hannan consistency*) An algorithm is ε -Hannan consistent for some $\varepsilon \geq 0$ if $\limsup_{t \rightarrow \infty} r(t) \leq \varepsilon$ holds with probability 1, where the “probability” is understood with respect to the randomization of the algorithm. An algorithm is Hannan consistent if it is 0-Hannan consistent.

We now present regret matching and Exp3, two of the ε -Hannan consistent algorithms previously used in MCTS. The proofs of Hannan consistency of variants of these two algorithms, as well as more related results, can be found in a survey by Cesa-Bianchi and Lugosi (2006, Sect. 6). The fact that the variants presented here are ε -HC is not explicitly stated there, but it immediately follows from the last inequality in the proof of Theorem 6.6 in the survey.

2.4.1 Exponential-weight algorithm for exploration and exploitation

Input: K - number of actions; γ - exploration parameter
 1: $\forall i : G_i \leftarrow 0$ // initialize cumulative sum estimates
 2: **for** $t \leftarrow 1, 2, \dots$ **do** // in each iteration

 3: $\forall i : \mu_i \leftarrow \frac{\exp(\frac{\gamma}{K} G_i)}{\sum_{j=1}^K \exp(\frac{\gamma}{K} G_j)}$ // compute the new strategy
 4: $\forall i : \mu'_i \leftarrow (1-\gamma)\mu_i + \frac{\gamma}{K}$ // add uniform exploration

 5: Sample action $i(t)$ from distribution μ' and receive reward r
 6: $G_{i(t)} \leftarrow G_{i(t)} + r / \mu'_{i(t)}$ // update the cumulative sum estimates

Algorithm 3: Exponential-weight algorithm for Exploration and Exploitation (Exp3) algorithm for regret minimization in adversarial bandit setting

The most popular algorithm for minimizing regret in adversarial bandit setting is the Exponential-weight algorithm for Exploration and Exploitation (Exp3) proposed by Auer et al. (2003), further improved by Stoltz (2005) and then yet further by Bubeck and Cesa-Bianchi (2012, Sect. 3). The algorithm has many different variants for various modifications of the setting and desired properties. We present a formulation of the algorithm based on the original version in Algorithm 3.

Exp3 stores the estimates of the cumulative reward of each action over all iterations, even those in which the action was not selected. In the pseudo-code in Algorithm 3, we denote this value for action i by G_i . It is initially set to 0 on line 1. In each iteration, a probability distribution μ is created proportionally to the exponential of these estimates. The distribution

is combined with a uniform distribution with probability γ to ensure sufficient exploration of all actions (line 4). After an action is selected and the reward is received, the estimate for the performed action is updated using *importance sampling* (line 6): the reward is weighted by one over the probability of using the action. As a result, the expected value of the cumulative reward estimated only from the time steps where the agent selected the action is the same as the actual cumulative reward over all the time steps.

In practice, the optimal choice of the exploration parameter γ strongly depends on the computation time available for each decision and the specific domain (Tak et al. 2014). However, the amount of exploration γ directly translates to γ -Hannan consistency of the algorithm. We later show that, asymptotically, smaller γ yields smaller worst-case error when the algorithm is used in SM-MCTS.

2.4.2 Regret matching

An alternative learning algorithm that allows minimizing regret in adversarial bandit setting is regret matching (Hart and Mas-Colell 2001), later generalized as *polynomially weighted average forecaster* (Cesa-Bianchi and Lugosi 2006). Regret matching (RM) corresponds to selection of the parameter $p = 2$ in the more general formulation. It is a general procedure originally developed for playing known general-sum matrix games in Hart and Mas-Colell (2000). The algorithm computes, for each action in each step, the regret for not playing another fixed action every time the action has been played in the past. The action to be played in the next round is selected randomly with probability proportional to the positive portion of the regret for not playing the action.

The average strategy⁴ resulting from this procedure has been shown to converge to the set of coarse correlated equilibria in general-sum games. As a result, it converges to a Nash equilibrium in a zero-sum game. The regret matching procedure in Hart and Mas-Colell (2000) requires the exact information about all utility values in the game, as well as the action selected by the opponent in each step. In Hart and Mas-Colell (2001), the authors modify the regret matching procedure and relax these requirements. Instead of computing the exact values for the regrets, the regrets are estimated in a similar way as the cumulative rewards in Exp3. As a result, the modified regret matching procedure is applicable to the MAB problem.

Input: K - number of actions; γ - the amount of exploration

```

1:  $\forall_i : R_i \leftarrow 0$  // initialize regret estimates
2: for  $t \leftarrow 1, 2, \dots$  do // in each iteration
3:  $\forall_i : R_i^+ \leftarrow \max\{0, R_i\}$ 
4: if  $\sum_{j=1}^K R_j^+ = 0$  then // all regrets are non-positive
5:  $\forall_i : \mu_i \leftarrow 1/K$  // use uniform strategy
6: else // otherwise, update the strategy based on regrets
7:  $\forall_i : \mu_i \leftarrow (1 - \gamma) \frac{R_i^+}{\sum_{j=1}^K R_j^+} + \frac{\gamma}{K}$ 
8: Sample action  $i(t)$  from distribution  $\mu$  and receive reward  $r$ 
9:  $\forall_i : R_i \leftarrow R_i - r$  // update regrets
10:  $R_{i(t)} \leftarrow R_{i(t)} + r/\mu_{i(t)}$ 

```

Algorithm 4: Regret matching variant for regret minimization in adversarial bandit setting.

⁴ The average strategy is defined as $\bar{\mu}(t) := \frac{1}{t} \sum_{s=1}^t \mu(s)$, where $\mu(s)$ is the strategy used at iteration s .

We present the algorithm in Algorithm 4. The algorithm stores the estimates of the regrets for not playing action i in all time steps in the past in variables R_i . On lines 3–7, it computes the strategy for the current time step. If there is no positive regret for any action, a uniform strategy is used (line 5). Otherwise, the strategy is chosen proportionally to the positive part of the regrets (line 7). The uniform exploration with probability γ is added to the strategy as in the case of Exp3. It also ensures that the addition on line 10 is bounded.

Cesa-Bianchi and Lugosi (2006) prove that regret matching eventually achieves zero regret in the adversarial MAB problem, but they provide the exact finite time bound only for the perfect-information case, where the agent learns rewards of all arms.

2.5 An alternative to SM-MCTS: counterfactual regret minimization

Counterfactual regret minimization (CFR) is an iterative algorithm for computing approximate Nash equilibria in zero-sum extensive-form games with imperfect information (EFGs). Since we use EFGs only at a few places, we overload the defined notation with corresponding concepts from EFGs. In the EFG setting, the elements of \mathcal{H} are typically called *histories* rather than states. Unlike in SMGs, each non-terminal history $h \in \mathcal{H}$ only has a single acting player. The second difference is that in EFGs, histories are partitioned into *information sets*. Instead of observing the current history h directly, the acting player only sees the information set I that h belongs to. The EFG framework is more general than the SMG one since each simultaneous decision in an SMG can be modeled by two consecutive decisions in an EFG (where the player who acts second does not know the action chosen by the first player).

Let σ be the strategy profile of the players and denote $\pi^\sigma(h)$ the probability of reaching $h \in \mathcal{H}$ from the root of the game under the strategy profile σ , $\pi^\sigma(h, z)$ the probability of reaching history z given the game has already reached history h . We use the lower index at π to denote the players who contribute to the probability, i.e., $\pi_p^\sigma(h)$ is player p 's contribution⁵ to the probability of reaching h and $\pi_{-p}^\sigma(h)$ is the contribution of the opponent of p and chance if it is present, i.e., $\pi^\sigma(h) = \pi_p^\sigma(h) \cdot \pi_{-p}^\sigma(h)$. We further denote by ha the history reached after playing action a in history h . The *counterfactual value* of player p playing action a in an information set I under a strategy σ is the expected reward obtained when player p first chooses the actions to reach I , plays action a , and then plays based on σ_p , while the opponent and chance play according to σ_{-p} all the time:

$$v_p^\sigma(I, a) = \sum_{(h, z) \in Z_I} \pi_{-p}^\sigma(h) \pi^\sigma(ha, z) u_p(z),$$

where $Z_I = \{(h, z) | z \in \mathcal{Z}, h \in I, h \text{ is prefix of } z\}$ are the terminal histories that visit information set I by a prefix h .

Counterfactual regret in an information set is the external regret (see Definition 2) with respect to the counterfactual values. Counterfactual regret minimization algorithms minimize counterfactual regret in each information set, which provably leads to convergence of average strategies to a Nash equilibrium of the EFG (Zinkevich et al. 2007). The variant of counterfactual regret minimization most relevant for this paper is *Monte Carlo Counterfactual Regret Minimization* (MCCFR) and more specifically outcome sampling. MCCFR minimizes the counterfactual regrets by minimizing their unbiased estimates obtained by sampling. In the case of outcome sampling, these estimates are computed based on sampling

⁵ It is the product of the probabilities of the actions executed by player p in history h .

a single terminal history, as in MCTS. Let $q(z)$ be the probability of sampling a history z . The *sampled counterfactual value* is:

$$\tilde{v}_p^\sigma(I, a) = \begin{cases} \frac{1}{q(z)} \pi_{-p}^\sigma(h) \pi^\sigma(ha, z) u_p(z) & \text{if } (h, z) \in Z_I \\ 0 & \text{otherwise.} \end{cases}$$

The sampling probability typically decreases exponentially with the depth of the tree. Therefore, $\frac{1}{q(z)}$ will often be large, which causes high variance in sampled counterfactual value. This has been shown to cause slower convergence both theoretically (Gibson et al. 2012) and practically (Bošanský et al. 2016).

2.6 Relation to multi-agent reinforcement learning

Our work is also related to multi-agent reinforcement learning (MARL) in Markov games (Littman 1994). The goal of reinforcement learning is to converge to the optimal policy based on rewards obtained in simulations. Markov games are more general than SMGs in allowing immediate rewards and cycles in the state space. However, any SMG can be viewed as a Markov game. Therefore, the negative results presented in this paper apply to Markov games as well.

To the best of our knowledge, the existing algorithms in MARL literature do not help with answering the question of convergence of separate Hannan consistent strategies in individual decision points. They either explicitly approximate and solve the matrix games for individual stages of the Markov games (e.g., Littman 1994) or do not have convergence guarantees beyond repeated matrix games (Bowling and Veloso 2002).

2.7 Relation to stochastic shortest path problem

Recently, some authors considered variants of Markov decision processes (MDP, see for example Puterman 2014), where the rewards may change over time (stochastic shortest path problem, discussed for example in Neu et al. 2010) or even more generally, where the rewards *and* the transition probabilities may change over time (Yu and Mannor 2009; Abbasi et al. 2013).

SM-MCTS can be viewed as a special case of this scenario, where the transition probabilities change over time, but the rewards remain the same. Indeed, assuming the role of one of the players, we can view each node of the game tree as a state in MDP. From a state h , we can visit its child nodes with a probability which depends on the strategy of the other player. This strategy is unknown to the first player and will change over time.

Our setting is more specific than the general version of MDP—the state space contains no loops, as it is, in fact, a tree. Moreover, the rewards are only received at the terminal states and correspond to the value of these states. Consequently, an algorithm which would perform well in this special case of MDPs with variable transition probabilities could also be successfully used for solving simultaneous-move games. However, to the best of our knowledge, so far all such algorithms require additional assumptions, which do not hold in our case.

3 Application of the MAB problem to SM-MCTS(-A)

To analyze SM-MCTS(-A) we need to know how it is affected by the selection policies it uses (line 4 in Algorithms 1 and 2) and by the properties of the game it is applied to. In

this section, we first introduce some additional notation related to the MAB problem and SM-MCTS(-A). We then frame the events at h as a separate MAB problem (P^h) (resp. (\bar{P}^h) for SM-MCTS-A) in such a way that applying the bandit algorithm from h to (P^h) yields exactly the output observed at line 4.

Throughout the paper we will use the following notation for quantities related to MAB problems: In any MAB problem (P), the reward assignment $x_i(t)$ is such that

$$x_i(t) = \text{the reward the agent would receive in } (P) \text{ if they chose the action } i \text{ at time } t. \quad (3)$$

We define the notions of *cumulative payoff* $G(\cdot)$ and *maximum cumulative payoff* $G_{\max}(\cdot)$ and relate these quantities to the external regret⁶:

$$\begin{aligned} G(t) &:= \sum_{s=1}^t x_{i(s)}(s) \\ G_{\max}(t) &:= \max_{i \in \mathcal{I}_1} \sum_{s=1}^t x_i(s), \\ R(t) &= G_{\max}(t) - G(t). \end{aligned} \quad (4)$$

We also define the corresponding average notions and relate them to the average regret:

$$\begin{aligned} g(t) &:= G(t)/t, \\ g_{\max}(t) &:= G_{\max}(t)/t, \\ r(t) &= g_{\max}(t) - g(t). \end{aligned} \quad (5)$$

Next, we introduce the additional notation related to SM-MCTS(-A). By $i^h(t)$ and $j^h(t)$ we denote the action chosen by player 1 (resp. 2) during the t -th visit of h . To track the number of uses of each action, we set⁷

$$\begin{aligned} t_i^h &:= 1 + \left| \left\{ 1 \leq s \leq t - 1 \mid i^h(s) = i \right\} \right|, \\ t_{ij}^h &:= 1 + \left| \left\{ 1 \leq s \leq t - 1 \mid i^h(s) = i \ \& \ j^h(s) = j \right\} \right| \end{aligned} \quad (6)$$

and define t_j^h analogously to t_i^h . Note that when $(i, j) = (i^h(t), j^h(t))$, t_{ij}^h is actually equal to the number of times this joint action has been used up to (and including) the t -th iteration. For $(i, j) \neq (i^h(t), j^h(t))$, this is equal to the same number increased by 1.

We now define (P^h). When referring to the quantities from (4) and (5) which correspond to (P^h), we will add the superscript h (for example G_{\max}^h, g^h, r^h). To keep the different levels of indices manageable, we will sometimes write e.g., g_{ij}^h and x_{ij}^h instead of $g^{h_{ij}}$ and $x^{h_{ij}}$. To indicate whether an average regret is related to players 1 or 2, we denote the corresponding quantities as r_1^h and r_2^h .

Since we want the reward assignment corresponding to (P^h) to coincide with what is happening at h during SM-MCTS(-A), (3) requires us to define $x_i^h(t)$ as

⁶ For definition of the external regret, see Definition 2.

⁷ Note that we define t_i^h as the number of uses of i up to the $(t - 1)$ -th visit of h , increased by 1, even though the more natural candidate would be simply the number of uses of i up to the t -th visit. However, this version will simplify the notation later and for the purposes of computing the empirical frequencies, the difference between the two definitions becomes negligible with increasing t .

$$\begin{aligned}
 x_i^h(t) &:= \text{the reward } x \text{ from line 6 in Figure 1 (resp. 2) that we would get} \\
 &\quad \text{during the } t\text{-th visit of } h \text{ if we switched the first action at line 4} \\
 &\quad \text{to } i \text{ (while keeping the choice of player 2 as } j = j^h(t)\text{)}. \tag{7}
 \end{aligned}$$

In particular, the reward for the selected action $i = i^h(t)$ is

$$x^h(t) = x_{i^h(t)}^h(t) = \text{the reward obtained during the } t\text{-th visit of } h. \tag{8}$$

For $i \neq i^h(t)$, the unobserved reward $x_i^h(t)$ corresponds the value we would receive if we ran SM-MCTS(-A)($h_{ij^h(t)}$) during the t -th visit of h . Before the t -th visit of h , its child $h_{ij^h(t)}$ has been visited $(t_{ij^h(t)}^h - 1)$ -times. The hypothetical visit from (7) would therefore be the $t_{ij^h(t)}^h$ -th one, implying that

$$x_i^h(t) = x_{i^h(t)}^h \left(t_{ij^h(t)}^h \right). \tag{9}$$

We note a property of $x_i^h(t)$ that obviously follows from either (7) or (9), but might be unintuitive and is crucial for understanding the behavior of SM-MCTS(-A). Consider an action $i \neq i^h(t)$ that has not been selected at time t . Then, assuming the opponent keeps playing $j^h(t)$, the random variable $x_i^h(t)$ will *not change* until player 1 selects i and $t_{ij^h(t)}^h$ increases. (Because nodes that do not get visited remain inactive and none of their variables change.) This setting where rewards come from “reward pools” and stay around until they get “used up” is in a direct contrast with the non-adaptive MAB setting where even the non-selected rewards disappear. However, we argue that this behavior is inherent to the presented version of SM-MCTS(-A), and might cause its non-averaged variant to malfunction (as demonstrated in Sect. 4).

The MAB problem (\bar{P}^h) is defined analogously to (P^h), except that we denote the rewards as $\bar{x}_i^h(t)$ instead of $x_i^h(t)$, and define $\bar{x}_i^h(t)$ as the number \bar{x} (rather than x) from line 6 from Algorithm 2. This \bar{x} is, by definition, equal to the average reward $g_{ij^h(t)}^h(\cdot)$ from the corresponding child node. It follows that $\bar{x}_i^h(t)$ coincides with $g_{ij^h(t)}^h(t_{ij^h(t)}^h)$ (more precisely, $g_{ij^h(t)}^h(t_{ij^h(t)}^h)$ is a realization of the random variable $\bar{x}_i^h(t)$).

Lastly, we define the *empirical* and *average strategies* corresponding to a specific run of SM-MCTS(-A). By $t^h(T)$ we denote the number of visits of h up to the T -th iteration of SM-MCTS(-A). By *empirical frequencies* we mean the strategy profile $\hat{\sigma}(T) = (\hat{\sigma}_p^h(t^h(T)))_{p=1,2, h \in \mathcal{H}}$ defined as $\hat{\sigma}_1^h(t)(i) := t_i^h/t$ (resp. t_j/t for player 2). The average strategy $\bar{\sigma}(T)$ is defined analogously, with

$$\bar{\sigma}_p^h(t)(i) := \frac{1}{t} \sum_{s=1}^t \sigma_p^h(s)(i) \tag{10}$$

in place of $\hat{\sigma}_p^h(t)(i)$. The following lemma says the two strategies can be used interchangeably. The proof consists of an application of the strong law of large numbers and can be found in the appendix.

Lemma 1 *In the limit, the empirical frequencies and average strategies will almost surely be equal. That is, $\limsup_{t \rightarrow \infty} \max_{i \in \mathcal{A}_1} |\hat{\sigma}_1(t)(i) - \bar{\sigma}_1(t)(i)| = 0$ holds with probability 1.*

4 Insufficiency of local regret minimization for global convergence

One might hope that any HC selection policy in SM-MCTS would guarantee that the average strategy converges to a NE. Unfortunately, this is not the case—the goal of this section is to present a corresponding counterexample. The behavior of the counterexample is summarized by the following theorem:

Theorem 1 *There exists a simultaneous-move zero-sum game G with perfect information and a HC algorithm A , such that when A is used as a selection policy for SM-MCTS, then the average strategy $\bar{\sigma}(t)$ almost surely converges outside of the set of $\frac{1}{5}$ -Nash equilibria.*

How is such a pathological behavior possible? Essentially, it is because the sampling in SM-MCTS is closely related to the observed payoffs. By synchronizing the sampling algorithms in different nodes in a particular way, we will introduce a bias to the payoff observations. This will make the optimal strategy look worse than it actually is, leading the players to adopt a different strategy.

Note that the algorithm from Theorem 1 does have the guaranteed exploration property defined in Sect. 5, which rules out some trivial counterexamples where parts of the game tree are never visited.

4.1 Simplifying remarks

We present two observations regarding the proof of Theorem 1.

Firstly, instead of a HC algorithm, it is enough to construct an ε -HC algorithm A_ε with the prescribed behavior for arbitrary $\varepsilon > 0$. The desired 0-HC algorithm can then be constructed in a standard way—that is, by using a 1-Hannan consistent algorithm A_1 for some period t_1 , then a $\frac{1}{2}$ -HC algorithm $A_{1/2}$ for a longer period t_2 and so on. By choosing a sequence $(t_n)_n$, which increases quickly enough, we can guarantee that the resulting combination of algorithms $(A_{1/n})$ is 0-Hannan consistent.

Furthermore, we can assume without loss of generality that the algorithm A knows if it is playing as the first or the second player and that in each node of the game, we can actually use a different algorithm A . This is true because the algorithm always accepts the number of available actions as input. Therefore we could define the algorithm differently based on this number, and modify our game in some trivial way which would not affect our example (such as duplicating rows or columns).

4.2 The counterexample

The structure of the proof of Theorem 1 is now as follows. First, we introduce the game G (Fig. 2) and a sequence of joint actions in G which leads to

$$\sum_{h \in \mathcal{H}} r_1^h(t^h(T)) = \sum_{h \in \mathcal{H}} r_2^h(t^h(T)) = 0 \quad \& \quad u_1(\bar{\sigma}(T)) \leq v^G - \frac{1}{4}.$$

This behavior will serve as a basis for our counterexample. However, the “algorithms” generating this sequence of actions will be oblivious to the actions of the opponent, which means that they will not be ε -HC. In the second step of our proof (Lemma 2), we modify these algorithms in such a way that the resulting sequence of joint actions stays similar to the original sequence, but the new algorithms are ε -HC. In combination with the simplifying remarks from Sect. 4.1, this gives Theorem 1.

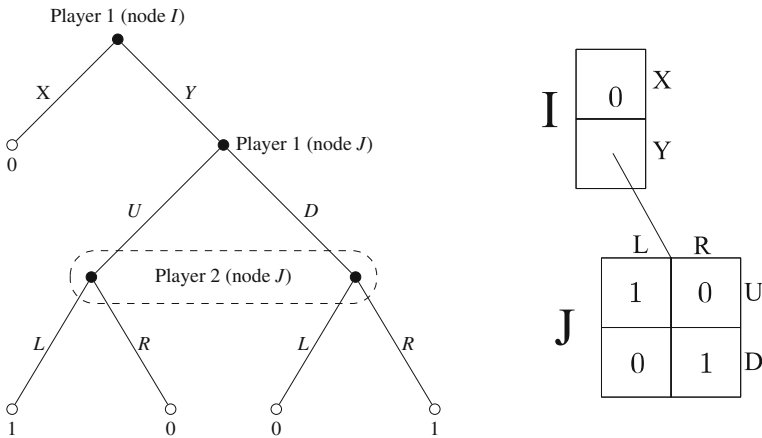


Fig. 2 Example of a game in which it is possible to minimize regret at each of the nodes while producing highly sub-optimal average strategy. The extensive form representation of the game is on the left side, the SMG representation is on the right side. Since the game is zero-sum, only the utilities for Player 1 are displayed

4.2.1 Deterministic version of the counterexample

Let G be the game from Fig. 2. First, we will note what the Nash equilibrium strategy π in G looks like. Having done that, we will describe a sequence of actions in G which leads to a different (non-NE) strategy $\bar{\sigma}$. We will then analyze the properties of this sequence, showing that utility of $\bar{\sigma}$ is sub-optimal, even though the regrets r_p^h for every $h \in \mathcal{H} = \{I, J\}$ will be equal to zero. The counter-intuitive part of the example is the fact that this pathological sequence satisfies $r_1^I = 0$.

When it comes to the optimal solution of G , we see that J is the well known game of matching pennies. The equilibrium strategy in J is $\pi_1^J = \pi_2^J = (\frac{1}{2}, \frac{1}{2})$ and the value of this subgame is $v^J = \frac{1}{2}$. Consequently, player 1 always wants to play Y at I , meaning that the NE strategy at I is $\pi_1^I = (0, 1)$ and the value of the whole game is $v^G = \frac{1}{2}$.

We define the “pathological” action sequence as follows. Let $h \in I, J$ and $t \in \mathbb{N}$. We set⁸

$$\begin{aligned} (i^I(1), i^I(2), \dots) &= (Y, X, X, Y, Y, X, X, Y, \dots), \\ ((i^J(1), j^J(1)), (i^J(2), j^J(2)), \dots) &= ((U, L), (U, R), (D, R), (D, L), \dots), \end{aligned}$$

where the dots mean that both sequences are 4-periodic. The resulting average strategy converges to $\bar{\sigma}$, where $\bar{\sigma}^I = \bar{\sigma}_1^J = \bar{\sigma}_2^J = (\frac{1}{2}, \frac{1}{2})$.

We now calculate the observed rewards which correspond to the behavior described above. We will use the notation from Sect. 3 to describe the events at I and J as MAB problems (P^I) and (P^J). The action sequence at J is defined in such a way that we have

$$(x^J(1), x^J(2), \dots) = (1, 0, 1, 0, \dots). \tag{11}$$

⁸ Recall that $i^h(t)$ and $j^h(t)$ are the actions of players 1 and 2 in t -th visit of a node h .

At the root I of G , we clearly have $x_X^I(t) = 0$ for all $t \in \mathbb{N}$. Since for $t = 2, 3, 6, 7, 10, 11, \dots$ we have $i^I(t) = X$, the rewards received at I during these iterations will be $x^I(t) = 0$. Combining this with (11), we see that the rewards received at I are

$$\begin{aligned} (x^I(1), x^I(2), \dots) &= (x^J(1), 0, 0, x^J(2), x^J(3), 0, 0, x^J(4), \dots) \\ &= (1, 0, 0, 0, 1, 0, 0, 0, \dots). \end{aligned} \tag{12}$$

In particular, the average reward in G converges to the above mentioned $\frac{1}{4}$.

We claim that the limit average strategy $\bar{\sigma}$ is far from being optimal. Indeed, as we observed earlier, the average payoff at I converges to $\frac{1}{4}$, which is strictly less than the game value $v^G = \frac{1}{2}$. Clearly, player 1 is playing sub-optimally by achieving only $u_1(\bar{\sigma}) = \frac{1}{4}$. By changing their action sequence at I to (Y, Y, Y, \dots) , they increase their utility to $\frac{1}{2}$.

However, none of the players observes any *local* regret. First, we check that no regret is observed at J . We see that the average payoff from J converges to $\frac{1}{2}$. In the node J , each player takes each action exactly⁹ half the time. Since we are in the matching pennies game, this means that neither of the players can improve their payoff at J by changing all their actions to any single action. Thus for both players $p = 1, 2$, we have $r_p^J(t) = 0$.

Next, we prove that player 1 observes no regret at I . We start by giving an intuitive explanation of why this is so. When computing regret, the player compares the rewards they received by playing as they did with the rewards they would receive if they changed their actions to Y, Y, \dots . If, at time $t = 1$, they asks themselves: ‘‘What would I receive if I played Y ?’’, the answer is ‘‘1’’. They *do* play Y , and so the strategy in J changes and when they next asks themselves the same question at time $t = 2$, the answer is ‘‘0’’. However this time, they *do not* play Y until $t = 4$, and so the answer remains ‘‘0’’ for $t = 3$ and $t = 4$. Only at $t = 5$ does the answer change to ‘‘1’’ and the whole process repeats. This way, they is ‘‘tricked’’ into thinking that the average reward coming from J is $\frac{1+0+0+0}{4} = \frac{1}{4}$, rather than $\frac{1}{2}$.

To formalize the above idea, we use the notation introduced in Sect. 3. In order to get $r_1^I(t) \rightarrow 0$, it suffices to show that for $t = 4k$, the $\frac{1}{t} \sum_{s=1}^t x_Y^I(s) = \frac{1}{4}$. Since the whole pattern is clearly 4-periodic, this reduces to showing that

$$x_Y^I(1) + x_Y^I(2) + x_Y^I(3) + x_Y^I(4) = 1 + 0 + 0 + 0 = 1.$$

Since for $t = 1$, the action Y actually got chosen, we have $Y = i^I(1)$ and

$$x_Y^I(1) = x_{i^I(1)}^I(1) \stackrel{\text{def.}}{=} x^I(1) \stackrel{(11)}{=} 1.$$

The same argument yields $x_Y^I(4) = 0$. For $t = 2$ and $t = 3$, the first iteration s satisfying $s \geq t$ and $i^I(s) = Y$ is $s = 4$. By definition in (8), it follows that both $x_Y^I(2)$ and $x_Y^I(3)$ are equal to $x^I(4)$, which is zero by (12).

4.2.2 Hannan consistent version of the counterexample

The following lemma states that the (deterministic) algorithms of players 1 and 2 described in Sect. 4.2.1 can be modified into ε -Hannan consistent algorithms in such a way that when facing each other, their behavior remains similar to the original algorithms.

Lemma 2 *Let G be the game from Fig. 2. Then for each $\varepsilon > 0$ there exist ε -HC algorithms A^I, A_1^J, A_2^J , such that when these algorithms are used for SM-MCTS in G , the resulting average strategy $\bar{\sigma}(t)$ converges to $\bar{\sigma}$, where $\bar{\sigma}^I = \bar{\sigma}_1^J = \bar{\sigma}_2^J = (\frac{1}{2}, \frac{1}{2})$.*

⁹ To be precise, the ratio is exactly 50% for iterations divisible by 4, and slightly different for the rest.

The strategy $\bar{\sigma}$ satisfies $u_1(\bar{\sigma}) = \frac{1}{4}$, while the equilibrium strategy π (where $\pi^I = (0, 1)$, $\pi_1^J = \pi_2^J = (\frac{1}{2}, \frac{1}{2})$), as shown in Sect. 4.2.1, gives utility $u_1(\pi) = \frac{1}{2}$. Therefore the existence of algorithms from Lemma 2 proves Theorem 1.

The key idea behind Lemma 2 is the following: both players repeat the pattern from Sect. 4.2.1, but we let them perform random checks which detect any major deviations from this pattern. If both players do this, then by Sect. 4.2.1 they observe no regret at any of the nodes. On the other hand, if one of them deviates enough to cause a non-negligible regret, they will be detected by the other player, who then switches to a “safe” ε -HC algorithm, leading again to a low regret. The definition of the modified algorithms used in Lemma 2, along with the proof of their properties, can be found in Appendix B.2.

4.3 Breaking the counterexample: SM-MCTS-A

We now discuss the pathological behavior of SM-MCTS above and explain how these issues are avoided by SM-MCTS-A and some other algorithms.

Firstly, what makes the counterexample work? The SM-MCTS algorithm repeatedly observes parts of the game tree to estimate the value of available decisions. When sampling the node J , half of the rewards propagated upwards to I have value 0, and half have value 1. SM-MCTS uses these payoffs directly and our rigged sampling scheme abuses this by introducing a payoff observation bias—that is, by making the zeros ‘stay’ three times longer than the ones.¹⁰ This causes the algorithm sitting at I to estimate the value v^J of J as $\frac{1}{4}$ (when in reality, it is $\frac{1}{2}$). Since $\frac{1}{4}$ is the average payoff at I , the algorithm at I suffers no regret, and it can keep on working in this pathological manner.

The culprit here is the *payoff observation bias*,¹¹ made possible by the synchronization of the selection algorithm used at I with the rewards coming from J .

So, why does SM-MCTS-A work in the counterexample above, when SM-MCTS did not? SM-MCTS-A also uses the biased payoffs, but instead of directly using the most recent sample, it works with the average of samples observed thus far. Since these averages converge to $v^J = \frac{1}{2}$, the estimates of v^J will also converge to this value, and no amount of ‘rigged weighting’ can ruin this. The only way to obtain the average reward of $\frac{1}{2}$ is to (almost) always play Y at I and because the algorithm used at I is HC, this is exactly what will happen.

Thus when using SM-MCTS-A, we will successfully find a NE of the game G , even when the observations made by the selection algorithms are very much biased. This argument can be generalized for an arbitrary game and a HC algorithm—we will do so in Theorem 2.

We conjecture that guaranteed convergence of SM-MCTS might still be possible, provided that the algorithms used as selection policies were HC and that the strategies prescribed by them changed slowly enough - such as is the case with Exp3 (where the strategies change slower and slower).¹²

¹⁰ For an explanation of the payoff observation bias, see “No-regret-at- I ” paragraph in Sect. 4.2.1.

¹¹ More on this in Sect. 7.2.

¹² When the the sampling strategies are constant, no synchronization like above is possible, which makes the payoff observations unbiased. We believe that when the strategies change slowly enough, the situation might be similar.

4.4 Breaking the counterexample: CFR

In Sect. 2.5 we described the CFR algorithm, which provably converges in our setting. We now explain how CFR deals with the counterexample.

It is a feature of MCTS that at each iteration, we only ‘care’ about the nodes we visited and we ignore the rest. The downside is that we never realize that every time we do *not* visit J , the strategy suggested for J by our algorithm performs extremely poorly. On the other hand, CFR visits every node in the game tree in every iteration. This makes it immune to our counterexample—indeed, if we are forced to care about what happens at J during every iteration, we can no longer keep on getting three zero payoffs for every 1 while still being Hannan consistent. The same holds for MCCFR, a Monte Carlo variant of CFR which no longer traverses the whole tree each iteration, but instead only samples a small portion of it.

5 Convergence of SM-MCTS and SM-MCTS-A

In this section, we present the main positive results—Theorems 2 and 3. Apart from a few cases, we only present the key ideas of the proofs here, while the full proofs can be found in the appendix. For an overview of the notation we use, see Table 1.

To ensure that the SM-MCTS(-A) algorithm will eventually visit each node, we need the selection policy to satisfy the following property.

Definition 4 We say that A is an *algorithm with guaranteed exploration* if, for any simultaneous-move zero-sum game¹³ G (as further specified in Sect. 2.1) where A is used by both players as a selection policy for SM-MCTS(-A), and for any game node $h \in \mathcal{H}$, $\lim_{t \rightarrow \infty} t_{ij}^h = \infty$ holds almost surely for every joint action $(i, j) \in \mathcal{A}_1(h) \times \mathcal{A}_2(h)$.

It is an immediate consequence of this definition that when an algorithm with guaranteed exploration is used in SM-MCTS(-A), every node of the game tree will be visited infinitely many times. From now on, we will therefore assume that, at the start of our analysis, the full game tree is already built—we do this because it will always happen after a finite number of iterations and, in most cases, we are only interested in the limit behavior of SM-MCTS(-A) (which is not affected by the events in the first finitely many steps).

Note that most of the HC algorithms, namely RM and Exp3, guarantee exploration without the need for any modifications, but there exist some HC algorithms, which do not have this property. However, they can always be adjusted in the following way:

Definition 5 Let A be a bandit algorithm. For fixed *exploration parameter* $\gamma \in (0, 1)$ we define a modified algorithm A^γ as follows. For time $t = 1, 2, \dots$ either:

- (a) explore with probability γ , or
- (b) run one iteration of A with probability $1 - \gamma$

(where “explore” means we choose the action randomly uniformly over available actions, without updating any of the variables belonging to A).

We define an algorithm $A^{\sqrt{\cdot}}$ analogously, except that at time t , the probability of exploration is $1/\sqrt{t}$ rather than γ .

¹³ We define the guaranteed exploration property this way (using extensive form games) to avoid further technicalities. Alternatively, we could define “two-player adversarial MAB problem” analogously to how adversarial MAB problem is defined in Definition 1, except that it would use a setting similar to the first part of Definition 6. We would then say that A is an algorithm with guaranteed exploration if, for any reward assignment $(a_{ij}(t))$, the limit is almost surely infinity for all joint actions.

Fortunately, ε -Hannan consistency is not substantially influenced by the additional exploration:

Lemma 3 For $\varepsilon \geq 0$, let A be an ε -Hannan consistent algorithm.

- (i) For any $\gamma > 0$, A^γ is an $(\varepsilon + \gamma)$ -HC algorithm with guaranteed exploration.
- (ii) $A^{\sqrt{\cdot}}$ is an ε -HC algorithm with guaranteed exploration.

The proof of this lemma can be found in the appendix.

5.1 Asymptotic convergence of SM-MCTS-A

A HC selection will always minimize the regret with respect to the values used as an input. But as we have seen in Sect. 4, using the observed values with no modification might introduce a bias, so that we end up minimizing the wrong quantity. One possible solution is to first modify the input by taking the averages, as in SM-MCTS-A. The high-level idea behind averaging is that it forces each pair of selection policies to optimize with respect to the subgame values v_{ij}^h , which leads to the following result:

Theorem 2 Let $\varepsilon \geq 0$ and let G be a zero-sum game with perfect information and simultaneous moves with maximal depth D and let A be an ε -Hannan consistent algorithm with guaranteed exploration, which we use as a selection policy for SM-MCTS-A.

Then almost surely, the empirical frequencies $(\hat{\sigma}_1(t), \hat{\sigma}_2(t))$ will eventually get arbitrarily close to a subgame-perfect $C\varepsilon$ -equilibrium, where $C = 2D(D + 1)$.

For $\varepsilon = 0$, Theorem 2 gives the following:

Corollary 1 If the algorithm from Theorem 2 is Hannan-consistent, the resulting strategy will eventually get arbitrarily close to a Nash equilibrium.

To simplify the proofs, we assume that $\varepsilon > 0$ - the variant with $\varepsilon = 0$ can be obtained by sending ε to zero, or by minor modifications of the proofs. We will first state two preliminary results, then we use an extension of the later one to prove Theorem 2 by backward induction. Firstly, we recall the following well-known fact, which relates the quality of the best responses available to the players with the concept of an equilibrium.

Lemma 4 In a zero-sum game with value v the following holds:

$$\begin{aligned} & \left(u_1(br, \hat{\sigma}_2) < v + \frac{\varepsilon}{2} \quad u_1(\hat{\sigma}_1, br) > v - \frac{\varepsilon}{2} \right) \\ & \implies \left(u_1(br, \hat{\sigma}_2) - u_1(\hat{\sigma}_1, \hat{\sigma}_2) < \varepsilon \quad u_2(\hat{\sigma}_1, br) - u_2(\hat{\sigma}_1, \hat{\sigma}_2) < \varepsilon \right) \\ & \stackrel{\text{def}}{\iff} (\hat{\sigma}_1, \hat{\sigma}_2) \text{ is an } \varepsilon\text{-equilibrium.} \end{aligned}$$

In order to start the backward induction, we first need to study what happens on the lowest level of the game tree, where the nodes consist of matrix games. It is well-known that in a zero sum matrix game, the average strategies of two Hannan consistent players will eventually get arbitrarily close to a Nash equilibrium—see Waugh (2009) and Blum and Mansour (2007). We prove a similar result for the approximate versions of the notions.

Lemma 5 *Let $\varepsilon \geq 0$ be a real number. If both players in a matrix game M are ε -Hannan consistent, then the following inequalities hold almost surely:*

$$v^M - \varepsilon \leq \liminf_{t \rightarrow \infty} g^M(t) \leq \limsup_{t \rightarrow \infty} g^M(t) \leq v^M + \varepsilon, \tag{13}$$

$$v^M - 2\varepsilon \leq \liminf_{t \rightarrow \infty} u_1^M(\hat{\sigma}_1(t), br) \quad \& \quad \limsup_{t \rightarrow \infty} u_1^M(br, \hat{\sigma}_2(t)) \leq v^M + 2\varepsilon. \tag{14}$$

The inequalities (13) are a consequence of the definition of ε -HC and the game value v^M . The proof of inequality (14) then shows that if the value caused by the empirical frequencies was outside of the interval infinitely many times with positive probability, it would be in contradiction with the definition of ε -HC.

Now, assume that the players are repeatedly playing some matrix game M , but they receive slightly distorted information about their payoffs. This is exactly the situation which will arise when a stacked matrix game is being solved by SM-MCTS-A and the “players” are selection algorithms deployed at a non-terminal node h (where the average payoffs $g^{hij}(t)$ coming from subgames rooted in h_{ij} are never exactly equal to v_{ij}^h). Such a setting can be formalized as follows:

Definition 6 (*Repeated matrix game with bounded distortion*) Consider the following general problem (\tilde{M}) . Let $m, n \in \mathbb{N}$. For each $t \in \mathbb{N}$, the following happens

1. Nature chooses a matrix $\tilde{M}(t) = (v_{ij}^{\tilde{M}}(t)) \in [0, 1]^{m \times n}$;
2. Players 1 and 2 choose actions $i \leq m$ and $j \leq n$ and observe the number $v_{ij}^{\tilde{M}}(t)$. This choice can be random and might depend on the previously observed rewards and selected actions (that is $v_{i(s)j(s)}^{\tilde{M}}(s)$ and $i(s)$, resp. $j(s)$ for player 2, for $s < t$).

Let $\delta > 0$. We call (\tilde{M}) a *repeated matrix game with δ -bounded distortion* if the following holds:

- For each i, j, t , $v_{ij}^{\tilde{M}}(t)$ is a random variable possibly depending on the choice of actions $i(s), j(s)$ for $s < t$.
- There exists a matrix $M = (v_{ij}^M) \in [0, 1]^{m \times n}$, such that almost surely:

$$\exists t_0 \in \mathbb{N} \forall t \geq t_0 \forall i, j : \left| v_{ij}^{\tilde{M}}(t) - v_{ij}^M \right| \leq \delta.$$

From first player’s point of view, each repeated matrix game with bounded distortion can be seen as an adversarial MAB problem $(P^{\tilde{M}})$ with reward assignment

$$x_i^{\tilde{M}}(t) := v_{ij(t)}^{\tilde{M}}(t).$$

When referring to the quantities related to $(P^{\tilde{M}})$, we add a superscript \tilde{M} —in particular, this concerns $R^{\tilde{M}}, g^{\tilde{M}}, g_{\max}^{\tilde{M}}$ (and so on) from (4) and (5) and the empirical frequencies $\hat{\sigma}^{\tilde{M}}(t)$ corresponding to the actions selected in $(P^{\tilde{M}})$. Recall that v^M and u_p^M for $p = 1, 2$ stand for the value and utility function corresponding to the matrix game M .

We formulate the following analogy of Lemma 5, which shows that ε -HC algorithms perform well even if they observe slightly perturbed rewards.

Proposition 1 *Let $c, \varepsilon > 0$ and let (\tilde{M}) be a repeated matrix game with $c\varepsilon$ -bounded distortion, played by two ε -HC players. Then the following inequalities hold almost surely:*

$$v^M - (c + 1)\varepsilon \leq \liminf_{t \rightarrow \infty} g^{\tilde{M}}(t) \leq \limsup_{t \rightarrow \infty} g^{\tilde{M}}(t) \leq v^M + (c + 1)\varepsilon, \tag{15}$$

$$v^M - 2(c + 1)\varepsilon \leq \liminf_{t \rightarrow \infty} u_1^M(\hat{\sigma}_1, br) \leq \limsup_{t \rightarrow \infty} u_1^M(br, \hat{\sigma}_2) \leq v^M + 2(c + 1)\varepsilon. \tag{16}$$

The proof is similar to the proof of Lemma 5. It needs an additional claim that if the algorithm is ε -HC with respect to the observed values with errors, it still has a bounded regret with respect to the exact values.

Next, we present the induction hypothesis around which the proof of Theorem 2 revolves. We consider the setting from Theorem 2 and let $\eta > 0$ be a small positive constant. For $d \in \{1, \dots, D\}$ we denote $C_d := d + O(\eta)$.¹⁴ If a node h in the game tree is terminal, we denote $d_h = 0$. For other nodes, we inductively define d_h , the depth of the sub-tree rooted at h , as the maximum of $d_{h_{ij}}$ over its children h_{ij} , increased by one. Recall that for $(i, j) \in \mathcal{A}_1(h) \times \mathcal{A}_2(h)$, $v_{ij}^h = v^{h_{ij}} \in [0, 1]$ is the value of the subgame rooted at the child node h_{ij} of h .

Induction hypothesis (IH_d) is the claim that for each node h with $d_h = d$, there almost surely exists t_0 such that for each $t \geq t_0$

- (i) the payoff $g^h(t)$ will fall into the interval $(v^h - C_d\varepsilon, v^h + C_d\varepsilon)$;
- (ii) the utilities $u_1^{M_h}(\hat{\sigma}_1(t), br)$ and $u_1^{M_h}(br, \hat{\sigma}_2(t))$ with respect to the matrix game $M_h := (v_{ij}^h)_{i,j}$ will fall into the interval $(v^h - 2C_d\varepsilon, v^h + 2C_d\varepsilon)$.

As we mentioned above, any node h with $d_h = 1$ is actually a matrix game. Therefore Lemma 5 ensures that (IH_1) holds. In order to get Theorem 2, we only need the property (ii) to hold for every $h \in \mathcal{H}$. The condition (i) is “only” required for the induction itself to work.

We now show that in the setting of Theorem 2, the induction step works. This is the part of the proof where we see the difference between SM-MCTS and SM-MCTS-A—in SM-MCTS, the rewards will be different over time, due to the randomness of used strategies. But the averaged rewards, which SM-MCTS-A uses, will eventually be approximately the same as values of the respective game nodes. This allows us to view the situation in every node as a repeated game with bounded distortion and apply Proposition 1.

Proposition 2 *In the setting from Theorem 2, the implication $(IH_d) \implies (IH_{d+1})$ holds for every $d \in \{1, \dots, D - 1\}$.*

Proof Assume that (IH_d) holds and let h be a node with $d_h = d + 1$. We will describe the situation in h as a repeated matrix game (\tilde{M}_h) with bounded distortion:

- The corresponding ‘non-distorted’ matrix is $M_h := (v_{ij}^h)_{i,j}$.
- The actions available to the players 1 and 2 are $i \in \mathcal{A}_1(h)$ and $j \in \mathcal{A}_2(h)$.
- The actions $i(t), j(t)$ chosen by the players are the actions $i^h(t), j^h(t)$.¹⁵

¹⁴ Our proofs use many inequalities that hold up to a depth-dependent noise, which we denote as $M_d\eta$. However, this noise can be made arbitrarily small by choosing the right η at the start of each proof. The exact value of M_d is therefore unimportant, and we hide it using the big O notation. Note that multiplying $O(\eta)$ by ε still yields $O(\eta)$, since we can always assume $\varepsilon \leq 1$ (the utilities are bounded by 1 by definition).

¹⁵ Recall that by definition in (8), $i^h(t)$ and $j^h(t)$ are the actions chosen by SM-MCTS-A during the t -th visit of the node h .

- The “distorted” rewards $v_{ij}^{\tilde{M}^h}(t)$ are the average rewards g_{ij}^h coming from the subgames rooted in the children of h . The exact correspondence is

$$v_{ij}^{\tilde{M}^h}(t) := g_{ij}^h(t_{ij}^h).$$

In other words, $v_{ij}^{\tilde{M}^h}(t)$ is the average of the first t_{ij}^h rewards obtained by visiting h_{ij} during SM-MCTS-A.

In particular, we have $x_i^{\tilde{M}^h}(t) = v_{ij(t)}^{\tilde{M}^h}(t) = g_{ij(t)}^h(t_{ij}^h)$. Comparing this with the definition of $\bar{x}_i^h(t)$ from Sect. 3, we see that the MAB problems (\bar{P}^h) and $(P^{\tilde{M}^h})$ coincide. As a result, Proposition 1 can be applied to the actions taken at h (assuming we can bound the distortion).

By (i) from (IH_d) , the distortion of (\tilde{M}^h) is eventually bounded by $C_d\varepsilon = d\varepsilon + O(\eta)$. Consequentially, we can apply Proposition 1—the inequality (15) shows that the average rewards g^h will eventually belong to the interval $(v^h - ((C_d + 1)\varepsilon + \eta), v^h + (C_d + 1)\varepsilon + \eta)$. Since $(C_d + 1)\varepsilon + \eta = C_{d+1}\varepsilon$, (i) from (IH_{d+1}) holds. The property (ii) follows from the inequality (16) in Proposition 1. □

We are now ready to give the proof of Theorem 2. Essentially, it consists of summing up the errors from the first part of the induction hypothesis and using Lemma 4.

Proof of Theorem 2 First, we observe that by Lemma 5, (IH_1) holds, and consequently by Proposition 2, (IH_d) holds for every $d = 1, \dots, D$.

Denote by $u^h(\sigma)$ (resp. $u_{ij}^h(\sigma)$) the expected player 1 payoff corresponding to the strategy σ used in the subgame rooted at node $h \in \mathcal{H}$ (resp. its child h_{ij}). By Lemma 4 we know that in order to prove Theorem 2, it is enough to show that for every $h \in \mathcal{H}$, the strategy $\hat{\sigma}(t)$ will eventually satisfy

$$u^h(br, \hat{\sigma}_2(t)) \leq v^h + (d_h + 1)d_h\varepsilon + O(\eta). \tag{17}$$

The role of both players in the whole text is symmetric, therefore (17) also implies the second inequality required by Lemma 4. We will do this by backward induction.

Since any node h with $d_h = 1$ is a matrix game, it satisfies $(P^h) = (P^{M_h})$ and $u^h = u_1^{M_h}$. In particular, (ii) from (IH_1) implies that the inequality (17) holds for any such h . Let $1 < d \leq D$, $h \in \mathcal{H}$ be such that $d_h = d$ and assume, as a hypothesis for backward induction, that the inequality (17) holds for each h' with $d_{h'} < d$. We observe that

$$\begin{aligned} u^h(br, \hat{\sigma}_2(t)) &= \max_{i \in \mathcal{A}_1(h)} \sum_{j \in \mathcal{A}_2(h)} \hat{\sigma}_2^h(t)(j) u_{ij}^h(br, \hat{\sigma}_2(t)) \\ &\leq v^h + \left(\max_{i \in \mathcal{A}_1(h)} \sum_{j \in \mathcal{A}_2(h)} \hat{\sigma}_2^h(t)(j) v_{ij}^h - v^h \right) \\ &\quad + \max_{i \in \mathcal{A}_1(h)} \sum_{j \in \mathcal{A}_2(h)} \hat{\sigma}_2^h(t)(j) \left(u_{ij}^h(br, \hat{\sigma}_2(t)) - v_{ij}^h \right). \end{aligned}$$

Since $v_{ij}^h = v_{ij}^{M_h}$, the first term in the brackets is equal to $u_1^{M_h}(br, \hat{\sigma}_2^h) - v^h$. By (ii) from (IH_d) , this is bounded by $2d\varepsilon + \eta$. By the backward induction hypothesis we have

$$u_{ij}^h(br, \hat{\sigma}_2(t)) - v_{ij}^h \leq d(d - 1)\varepsilon + O(\eta)$$

Therefore we have

$$u^h(br, \hat{\sigma}_2(t)) \leq v^h + 2d\varepsilon + d(d - 1)\varepsilon + O(\eta) = v^h + (d + 1)d\varepsilon + O(\eta).$$

For $d = D, h = \text{root}$, Lemma 4 implies that $(\hat{\sigma}_1(t), \hat{\sigma}_2(t))$ forms a $(2D(D+1)\varepsilon + O(\eta))$ -equilibrium of the whole game. \square

5.2 Asymptotic convergence of SM-MCTS

We would like to prove an analogous result to Theorem 2 for SM-MCTS. Unfortunately, such a goal is unattainable in general—in Sect. 4 we presented a counterexample, showing that a such a theorem with no additional assumptions does not hold. We show that if an ε -HC algorithm with guaranteed exploration has an additional property of ‘having ε -unbiased payoff observations’ (ε -UPO, see Definition 7), it can be used as a selection policy for SM-MCTS, and it will always find an approximate equilibrium (Theorem 3). While we were unable to prove that specific ε -HC algorithms have the ε -UPO property, in Sect. 7, we provide empirical evidence supporting our hypothesis that the ‘typical’ algorithms, such as regret matching or Exp3, indeed do have ε -unbiased payoff observations.

5.2.1 Definition of the UPO property

First, we introduce the notation required for the definition of the ε -UPO property. Recall that for $h \in \mathcal{H}$ and a joint action $(i, j) \in \mathcal{A}_1(h) \times \mathcal{A}_2(h)$, $x_{ij}^h(m) = x^{h_{ij}}(m)$ denotes the m -th reward¹⁶ received when visiting h ’s child node h_{ij} . We denote the arithmetic average of the sequence $(x_{ij}^h(1), \dots, x_{ij}^h(n))$ as

$$\bar{x}_{ij}^h(n) = \frac{1}{n} \sum_{m=1}^n x_{ij}^h(m). \quad (18)$$

The numbers \bar{x}_{ij}^h are the variables we would prefer to work with. Indeed, their definition is quite intuitive and when A is ε -HC and $d_{h_{ij}} = 1$, the averages $\bar{x}_{ij}^h(n)$ can be used to approximate v_{ij}^h :

$$\limsup_n \left| \bar{x}_{ij}^h(n) - v_{ij}^h \right| \leq \varepsilon.$$

Unfortunately, the variables $\bar{x}_{ij}^h(n)$ do not suffice for our analysis, because as we have seen in Sect. 4, even when the differences $\left| \bar{x}_{ij}^h(n) - v_{ij}^h \right|$ get arbitrarily small, SM-MCTS might still perform very poorly. In (19), we define differently weighted averages $\tilde{x}_{ij}^h(n)$ of $(x_{ij}^h(1), \dots, x_{ij}^h(n))$ which naturally appear in this setting and are crucial in the proof of the upcoming Theorem 3. Before we jump into details, let us expand on an idea that we already hinted at earlier:

Assume that a node h_{ij} has already been visited $(m-1)$ -times or, equivalently, that $t_{ij}^h = m$. We can compute the m -th reward $x_{ij}^h(m)$ in advance and when player 2 selects j for the first time at some $s \geq t$, we offer $x_{ij}^h(m)$ as a possible reward $x_i^h(s)$. However, if player 1 plays *something other than i* at s , $x_{ij}^h(m)$ does not “get used up”—instead it “waits” to be used in future. Therefore, each reward $x_{ij}^h(m)$ must be weighted proportionally to the number

¹⁶ Note that $x_{ij}^h(t)$ are random variables and their distribution depends on the game to which h belongs and on the selection policies applied at all of the nodes in the subgame rooted in h_{ij} .

$w_{ij}^h(m)$ of times it got offered as a possible reward (including the one time when it eventually got chosen), because this is the number of times it will show up when regret R^h is calculated.

We define the weights $w_{ij}(n)$ and the weighted averages $\tilde{x}_{ij}(n)$ as follows¹⁷:

$$w_{ij}^h(m) := \left| \left\{ t \in \mathbb{N} \mid t_{ij}^h = m \ \& \ j^h(t) = j \right\} \right|$$

$$\tilde{x}_{ij}^h(n) := \frac{1}{\sum_{m=1}^n w_{ij}^h(m)} \sum_{m=1}^n w_{ij}^h(m) * x_{ij}^h(m), \tag{19}$$

We are now ready to give the definition:

Definition 7 (UPO) We say that an algorithm A guarantees ε -unbiased payoff observations, if for every (simultaneous-move zero-sum perfect information) game G in which A is used as a selection policy for SM-MCTS, for its every node h and every joint action (i, j) , the arithmetic averages \bar{x}_{ij}^h and weighted averages \tilde{x}_{ij}^h almost surely satisfy

$$\limsup_{n \rightarrow \infty} \left| \bar{x}_{ij}^h(n) - \tilde{x}_{ij}^h(n) \right| \leq \varepsilon.$$

We will abbreviate this by saying that “ A is an ε -UPO algorithm”.

Observe in particular that for an ε -UPO algorithm and $c > 0$, we have,

$$\limsup_{n \rightarrow \infty} \left| \bar{x}_{ij}^h(n) - v_{ij}^h \right| \leq c\varepsilon \text{ a.s.} \implies \limsup_{n \rightarrow \infty} \left| \tilde{x}_{ij}^h(n) - v_{ij}^h \right| \leq (c + 1)\varepsilon \text{ a.s..}$$

The most relevant ‘examples’ related to the UPO property are

- (a) the pathological algorithms from Sect. 4.2.1, where the UPO property does not hold (and SM-MCTS fails to find a reasonable strategy),
- (b) Theorem 3 below, which states that when a HC algorithm is UPO, it finds an equilibrium strategy,
- (c) if at each node h , players chose their actions as independent samples of some probability distributions P_1^h and P_2^h over $\mathcal{A}_1(h)$ and $\mathcal{A}_2(h)$, then these selection ‘algorithms’ are UPO.

We agree that the definition of the UPO property, as presented in Definition 7, is quite impractical. Unfortunately, we were unable to find a replacement property which would be easier to check while still allowing the “HC & ? \implies SM-MCTS finds a NE” proof to go through. We at least provide some more examples and a further discussion in the appendix (Sect. B.1).

5.2.2 The convergence of SM-MCTS

The goal of this section is to prove the following theorem:

Theorem 3 For $\varepsilon \geq 0$, let A be an ε -HC algorithm with guaranteed exploration that is ε -UPO. If A is used as selection policy for SM-MCTS, then the average strategy of A will eventually get arbitrarily close to a $C'\varepsilon$ -NE of the whole game, where $C' = 12(2^D - 1) - 8D$.

¹⁷ By $w_{ij}^h(m) * x_{ij}^h(m)$ we mean the sum of $w_{ij}^h(m)$ independent copies of $x_{ij}^h(m)$.

Note that if we prove the statement for $\varepsilon > 0$, we get the variant with $\varepsilon = 0$ ‘for free’. The proof of Theorem 3 is similar to the proof of Theorem 2. The only major difference is Proposition 3, which serves as a replacement for Proposition 1 (which cannot be applied to SM-MCTS). Essentially, Proposition 3 shows that under the specified assumptions, having low “local” regret in some $h \in \mathcal{H}$ with respect to $x_i^h(t)$ is sufficient to bound the regret with respect to the rewards originating from the matrix game (v_{ij}^h) .

Proposition 3 For $\varepsilon, c \geq 0$ let an algorithm A be ε -HC and $h \in \mathcal{H}$. If A is ε -UPO and $\limsup_{n \rightarrow \infty} |\bar{x}_{ij}^h(n) - v_{ij}^h| \leq c\varepsilon$ holds a.s. for each i, j , then

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \left(\max_i \sum_{s=1}^t v_{ij(s)}^h - \sum_{s=1}^t v_{i(s)j(s)}^h \right) \leq 2(c + 1)\varepsilon \text{ holds a.s.} \tag{20}$$

In particular, the regret $r^{M_h}(t)$ of A with respect to the matrix game $M_h = (v_{ij}^h)_{ij}$ satisfies

$$\limsup_{t \rightarrow \infty} r^{M_h}(t) \leq 2(c + 1)\varepsilon.$$

The proof of this proposition consists of rewriting the sums in inequality (20) and using the fact that the weighted averages \tilde{x}_{ij} are close to the standard averages \bar{x}_{ij} .

Consider the setting from Theorem 3 and let $\eta > 0$ be a positive constant. For $d \in \{1, \dots, D\}$ we denote $C'_d := 3 \cdot 2^{d-1} - 2 + O(\eta)$.

Induction hypothesis (IH'_d) is the claim that for each node h with $d_h = d$, there almost surely exists t_0 such that for each $t \geq t_0$

- (i) the payoff $g^h(t)$ fall into the interval $(v^h - C'_d\varepsilon, v^h + C'_d\varepsilon)$;
- (ii) the utilities $u^{M_h}(\hat{\sigma}_1(t), br)$ and $u^{M_h}(br, \hat{\sigma}_2(t))$ with respect to the matrix game $M_h = (v_{ij}^h)_{i,j}$ fall into the interval $(v^h - 2C'_d\varepsilon, v^h + 2C'_d\varepsilon)$.

Proposition 4 In the setting from Theorem 3, (IH'_d) holds for every $d \in \{1, \dots, D\}$.

Proof We proceed by backward induction—since the algorithm A is ε -HC, we know that, by Lemma 5, (IH'_1) holds with $C'_1 = 1 + O(\eta)$. Assume that (IH'_d) holds for some $d \in \{1, \dots, D - 1\}$. By the “in particular” part of Proposition 3, the choice of actions in h is $2(C'_d + 1)$ -HC with respect to the matrix game (v_{ij}^h) . By Lemma 5, (IH'_{d+1}) holds with $C'_{d+1} := 2(C'_d + 1) + O(\eta)$. It remains to show that C'_{d+1} is equal to $3 \cdot 2^d - 2 + O(\eta)$:

$$C'_{d+1} = 2(C'_d + 1) + O(\eta) = 2(3 \cdot 2^{d-1} - 2 + 1) + O(\eta) = 3 \cdot 2^d - 2 + O(\eta).$$

□

Analogously to Proposition 2 from Sect. 5.1, this in turn implies the main theorem of this section, the proof of which is similar to the proof of Theorem 2.

Proof of Theorem 3 Using Proposition 4, the proof is identical to the proof of Theorem 2—it remains to determine the value of C' . As in the proof of Theorem 2 we have $C' = 2 \cdot 2 \sum_{d=1}^D C'_d$, and we need to calculate this sum:

$$\sum_{d=1}^D C'_d = \sum_{d=1}^D (3 \cdot 2^{d-1} - 2) = 3(1 + \dots + 2^{D-1}) - 2D = 3(2^D - 1) - 2D.$$

□




Fig. 3 A game tree of a single-player game where the quality of a strategy has linear dependence on the exploration parameter and on the game depth D . The numbers u_d satisfy $0 < u_1 < u_2 < \dots < u_D < 1$. The player begins in the leftmost node and the numbers correspond to utilities in terminal nodes

5.3 Dependence of the eventual NE distance on the game depth

In Theorems 2 and 3, the bound on the distance of the average strategy to a NE (as measured by the constant C) is quadratic, resp. exponential, in the game depth. We investigate how far can the dependence be improved, constructing a counterexample where we lower-bound C by some constant times the depth. However, the more general question of lower-bounds remains open and is posed as Problem 1.

Proposition 5 *There exists $k > 0$, such that none of the Theorems 2 and 3 hold if the constant C is replaced by $\tilde{C} = kD$. This holds even when the exploration is removed from $\hat{\sigma}$.*

The proposition above follows from Example 1.

Example 1 (Distance from NE) Let G be the single player game¹⁸ from Fig. 3, $\eta > 0$ some small number, and D the depth of the game tree. Let Exp3 with exploration parameter $\gamma = k\varepsilon$ be our ε -HC algorithm (for a suitable choice of k). We recall that this algorithm will eventually identify the optimal action and play it with frequency $1 - \gamma$, and it will choose randomly otherwise. Denote the available actions at each node as (up, right, down), resp. (right, down) at the rightmost inner node. We define each of the rewards $u_d, d = 1, \dots, D - 1$ in such a way that Exp3 will always prefer to go up, rather than right. By induction over d , we can see that the choice $u_1 = 1 - \gamma/2 + \eta, u_{d+1} = (1 - \gamma/3)u_d$ is sufficient and for η small enough, we have

$$u_{D-1} = \left(1 - \frac{\gamma}{2} + \eta\right) \left(1 - \frac{\gamma}{3}\right) \dots \left(1 - \frac{\gamma}{3}\right) \leq \left(1 - \frac{\gamma}{3}\right)^{D-1} \doteq 1 - \frac{D-1}{3}\gamma$$

(where by \doteq we mean that for small γ , in which we are interested, the difference between the two terms is negligible). Consequently in each of the nodes, Exp3 will converge to the strategy $(1 - \frac{2}{3}\gamma, \frac{1}{3}\gamma, \frac{1}{3}\gamma)$ (resp. $(1 - \frac{\gamma}{2}, \frac{\gamma}{2})$), which yields the payoff of approximately $(1 - \gamma/3)u_{D-1}$. Clearly, the expected utility of such a strategy is approximately

$$u = (1 - \gamma/3)^D \doteq 1 - \frac{D}{3}\gamma.$$

On the other hand, the optimal strategy of always going right leads to utility 1, and thus our strategy $\hat{\sigma}$ is $\frac{D}{3}\gamma$ -equilibrium.

Note that in this particular example, it makes no difference whether SM-MCTS or SM-MCTS-A is used. We also observe that when the exploration is removed, the new strategy is to go up at the first node with probability 1, which again leads to regret of approximately $\frac{D}{3}\gamma$.

¹⁸ The other player always has only a single *no-op* action.

By increasing the branching factor of the game in the previous example from 3 to b (adding more copies of the “0” nodes) and modifying the values of u_d accordingly, we could make the above example converge to $2\frac{b-2}{b}D\gamma$ -equilibrium (resp. $\frac{b-2}{b}D\gamma$ once the exploration is removed).

We were able to construct a game of depth D and ε -HC algorithms, such that the resulting strategy $\hat{\sigma}$ converged to $3D\varepsilon$ -equilibrium ($2D\varepsilon$ after removing the exploration). However, the ε -HC algorithms used in this example are non-standard and would require the introduction of more technical notation. Therefore, since in our main theorem we use quadratic dependence $C = kD^2$, we instead choose to highlight the following open question:

Problem 1 Does Theorem 2 (and Theorem 3) hold with $C = kD$ for some $k > 0$? If not, what is the optimal value of C ?

We hypothesize that the answer is affirmative (and possibly the values $k = 3$, resp. $k = 2$ after exploration removal, are optimal), but the proof of such proposition would require techniques different from the one used in the proof of Theorem 2. Moreover, it might be more interesting to consider Problem 1 restricted to some class of “reasonable” no-regret algorithms, as opposed to arbitrary HC algorithms such as those presented in Sect. 4.

6 Exploitability and exploration removal

In this section, we discuss which strategy should be considered the output of our algorithms. We also introduce the concept of exploitability—a common approach to measuring strategy strength that focuses on the worst-case performance (see for example Johanson et al. 2011):

Definition 8 *Exploitability* of strategy σ_1 of player 1 is the quantity

$$\text{expl}_1(\sigma_1) := v - u_1(\sigma_1, \text{br}),$$

where v is the value of the game and br is a second player’s best response strategy to σ_1 . Similarly, we denote $\text{expl}_2(\sigma_2) := u_1(\text{br}, \sigma_2) - v$ and $\text{expl}(\sigma) := \text{expl}_1(\sigma_1) + \text{expl}_2(\sigma_2)$.

By definition of the game value, exploitability is always non-negative. Exploitability is closely related Nash equilibria, since we apparently have $\text{expl}_1(\sigma_1) = \text{expl}_2(\sigma_2) = 0$ if and only if $\sigma = (\sigma_1, \sigma_2)$ forms a NE.

If SM-MCT(-A) is run for t iterations, the obvious candidates for its output are the strategy $\sigma(t)$ from the last iteration, the average strategy $\bar{\sigma}(t) = \frac{1}{t} \sum_{s=1}^t \sigma(s)$ (defined in Eq. 10), and the empirical frequencies $\hat{\sigma}(t)$ (which are equivalent to $\bar{\sigma}(t)$ by Lemma 1). Most theoretical results only give guarantees for $\bar{\sigma}(t)$, and indeed, using $\sigma(t)$ would be naive as it is often very exploitable. However, there is a better option than using $\bar{\sigma}(t)$.

To improve the learning rate and ensure that the important parts of the game tree are found quickly, the selection functions used by SM-MCTS(-A) often have a fixed exploration rate $\gamma > 0$. This can happen either naturally (like with Exp3) or because it has been added artificially (cf. Definition 5). Each strategy $\sigma(t)$ is then of the form

$$\sigma(s) = (1 - \gamma)\mu(s) + \gamma \cdot \text{rnd} \quad (21)$$

for some $\mu(t)$ (where rnd denotes the uniformly random strategy). Rewriting $\bar{\sigma}(t)$ as

$$\bar{\sigma}(s) = \frac{1}{t} \sum_{s=1}^t \sigma(s) = \frac{1}{t} \sum_{s=1}^t [(1 - \gamma)\mu(s) + \gamma \cdot \text{rnd}] = (1 - \gamma)\bar{\mu}(t) + \gamma \cdot \text{rnd}, \quad (22)$$

we see that it contains the exact same amount of the random noise caused by exploration. While the exploration speeds up learning, it also increases the exploitability of $\bar{\sigma}_1$ in three ways: (a) it introduces noise to each $\bar{\sigma}_1^h$, (b) it does the same for the opponent, giving us inaccurate expectations about how an optimal opponent looks like and (c) it introduces a noise to the rewards SM-MCTS(-A) propagates upward in the tree. As a heuristic, Teytaud and Flory (2011) suggest that the exploration should be removed. In our case, we can use (22) to obtain the “denoised” average strategy $\bar{\mu}(t)$:

$$\bar{\mu}(t) := \frac{1}{(1-\gamma)} (\bar{\sigma}(t) - \gamma \cdot \text{rnd}). \quad (23)$$

Using $\bar{\mu}(t)$ instead of $\bar{\sigma}(t)$ cannot help with (b) and (c), but it does serve as an effective remedy for (a). This intuition is captured by Proposition 6, which shows that a better exploitability bound can be obtained if the exploration is removed. The experiments in Sect. 7 verify that the benefit of removing the exploration in SM-MCTS is indeed large.

Proposition 6 *Let M be a matrix game, A a HC algorithm, $\gamma > 0$, and A^γ the “ γ -exploration modification” of A from Definition 5. Let $\bar{\sigma}(t)$ be the average strategy obtained if both players use A^γ in M , and let $\bar{\mu}(t)$ be as in (23). Then for each p we have*

- (i) $\limsup_t \text{expl}_p(\bar{\sigma}_p(t)) \leq 2\gamma$,
- (ii) $\limsup_t \text{expl}_p(\bar{\mu}_p(t)) \leq \gamma$.

Lemma 5 gives (i) from Proposition 6 as a special case, and, as shown in the last part of the appendix, going through its proof in more detail yields (ii).

7 Experimental evaluation


In this section, we present the experimental data related to our theoretical results. First, we empirically evaluate our hypothesis that Exp3 and regret matching algorithms ensure the ε -UPO property. Second, we test the empirical convergence rates of SM-MCTS and SM-MCTS-A on synthetic games as well as smaller variants of games played by people. We investigate the practical dependence of the convergence error based on the important parameters of the games and evaluate the effect of removing the samples due to exploration from the computed strategies. We show that SM-MCTS generally converges as close to the equilibrium as SM-MCTS-A, but does it substantially faster. Therefore, since the commonly used Hannan consistent algorithms seem to satisfy the ε -UPO property, SM-MCTS is the preferable algorithm for practical applications.

At the time of writing this paper, the best theoretically sound alternative to SM-MCTS known to us is a modification of MCCFR called Online Outcome Sampling (OOS, Lisy et al. 2015). For a comparison of SM-MCTS, this variant of MCCFR, and several exact algorithms, we refer the reader to Bošanský et al. (2016). While OOS converges the fastest out of the sampling algorithms in small games, it is often inferior to SM-MCTS in online game playing in larger games, because of the large variance caused by the importance sampling corrections.

7.1 Experimental domains

Goofspiel Goofspiel is a card game that appears in many works dedicated to simultaneous-move games (for example Ross 1971; Rhoads and Bartholdi 2012; Saffidine et al. 2012; Lanctot et al. 2014; Bošanský et al. 2013). There are 3 identical decks of cards with values

Fig. 4 The Anti game used for evaluation of the algorithms



$\{0, \dots, d - 1\}$ (one for nature and one for each player). Value of d is a parameter of the game. The deck for the nature is shuffled at the beginning of the game. In each round, nature reveals the top card from its deck. Each player selects any of their remaining cards and places it face-down on the table so that the opponent does not see the card. Afterward, the cards are turned face-up, and the player with the higher card wins the card revealed by nature. The card is discarded in case of a draw. At the end, the player with the higher sum of the nature cards wins the game or the game is a draw. People play the game with 13 cards, but we use smaller numbers to be able to compute the distance from the equilibrium (that is, exploitability) in a reasonable time. We further simplify the game by a common assumption that both players know the sequence of the nature’s cards in advance.

Oshi-Zumo Each player in Oshi-Zumo (for example, Buro 2004) starts with N coins, and a one-dimensional playing board with $2K + 1$ locations (indexed $0, \dots, 2K$) stretches between the players. At the beginning, there is a stone (or a wrestler) located in the center of the board (that is, at position K). During each move, both players simultaneously place their bid from the amount of coins they have (but at least one if they still have some coins). Afterward, the bids are revealed, the coins used for bids are removed from the game, and the highest bidder pushes the wrestler one location towards the opponent’s side. If the bids are the same, the wrestler does not move. The game proceeds until the money runs out for both players, or the wrestler is pushed out of the board. The player closer to the wrestler’s final position loses the game. If the final position of the wrestler is the center, the game is a draw. In our experiments, we use a version with $K = 2$ and $N = 5$.

Random Game To achieve more general results, we also use randomly generated games. The games are defined by the number of actions B available to each player in each decision point and a depth D ($D = 0$ for leaves), which is the same for all branches. Each joint action is associated with a uniformly-random reward from $\{-1, 0, 1\}$, and the corresponding accumulated utilities in leaves are integers between $-D$ and D .

Anti The last game we use in our evaluation is based on the well-known single player game which demonstrates the super-exponential convergence time of the UCT algorithm (Coquelin and Munos 2007). The game is depicted in Fig. 4. In each stage, it receives the MCTS algorithm to end the game while it is optimal to continue until the end.

7.2 ϵ -UPO property

To be able to apply Theorem 3 (that is, the convergence of SM-MCTS without averaging) to Exp3 and regret matching, the selection algorithms have to assure the ϵ -UPO property for some ϵ . So far, we were unable to prove this hypothesis. Instead, we support this claim by the following numerical experiments. Recall that having ϵ -UPO property is defined as the claim that for every game node $h \in \mathcal{H}$ and every joint action (i, j) available at h , the difference $|\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)|$ between the weighted and arithmetical averages decreases below ϵ , as the number n of uses of (i, j) at h increases to infinity. We can think of this difference as a bias between real and observed average payoffs.

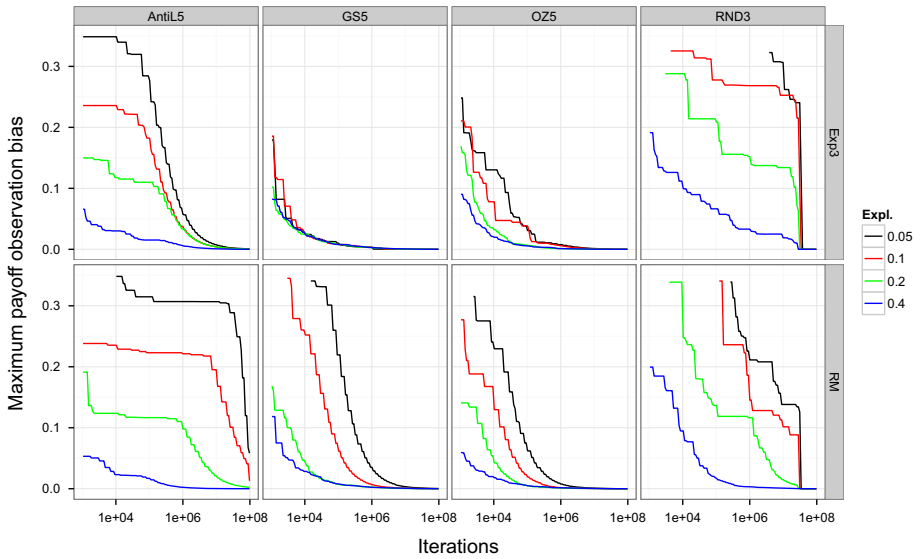


Fig. 5 The maximum of the bias in payoff observations in SM-MCTS without averaging the sample values (see the second paragraph in Sect. 7.2)

We measured the value of the difference in the sums in the root node of the four domains described above. Besides the random games, the depth of the game was set to 5. For the random games, the depth and the branching factor was $B = D = 3$. Figure 5 presents one graph for each domain and each algorithm. The x-axis is the number of iterations and the y-axis depicts the maximum value of the difference $|\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)|$ from the iteration on the x-axis to the end of the run of the algorithm ($\max_{n \in (x, \dots, 10^8), i, j} |\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)|$). The presented value is the maximum from 50 runs of the algorithms. For all games, the difference eventually converges to zero. Generally, larger exploration ensures that the difference goes to zero more quickly and the bias in payoff observation is smaller.

The main reason for the bias is easy to explain in the Anti game. Figure 6a presents the maximal values of the bias in small time windows during the convergence from all 50 runs. It is apparent that the bias during the convergence tends to jump very high (higher for smaller exploration) and then gradually decrease. This, however, happens only until a certain point in time. The reason for this behavior is that if the algorithm learns an action is good in a specific state, it will use it very often and will do the updates for the action with much smaller weight in $\tilde{x}_{ij}(n)$ than the updates for the other action. However, when the other action later proves to be substantially better, the value of that action starts increasing rather quickly. At the same time, its probability of being played starts increasing, and as a result, the weights used for the received rewards start decreasing. This will cause a strong dependence between the rewards and the weights, which causes the bias. With smaller exploration, it takes more time to identify the better alternative action; hence, when it happens, the wrong action has already accumulated larger reward and the discrepancy between the right values and the probability of playing the actions is even stronger.

We also tested satisfaction of the UPO property in the root node of depth 4 Goofspiel, using Exp3 algorithm and exploration $\epsilon = 0.001$. The results in Fig. 6b serve as evidence supporting

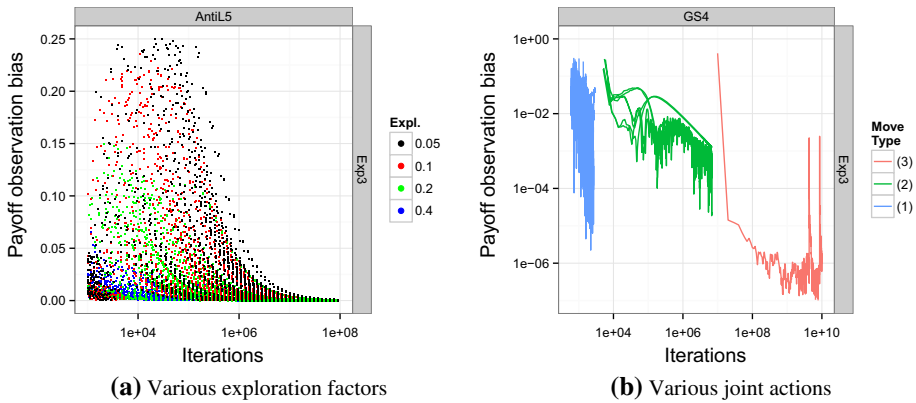


Fig. 6 The dependence of the current value of $|\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)|$ on the number of iterations that used the given joint action in **a** Anti game and **b** Goofspiel with 4 cards per deck

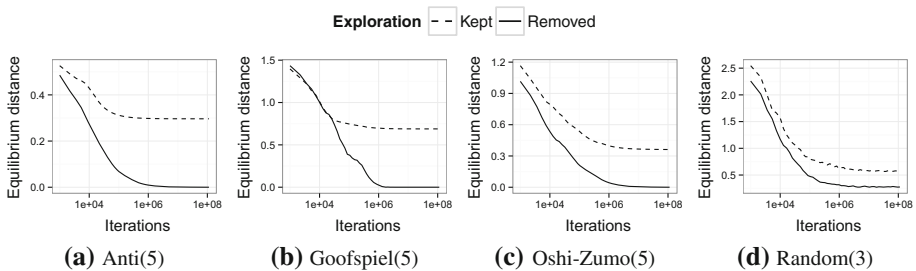


Fig. 7 The effect of removing exploration in SM-MCTS with Exp3 selection and $\epsilon = 0.2$

our conjecture that Exp3 with exploration 0.001 possesses the 0.001-UPO property (however this time, higher n_0 is required—around $5 \cdot 10^6$).

We can divide the joint actions (i, j) at the root into three groups: (1) the actions which both players play (nearly) only when exploring, (2) the actions which one of the players chooses only because of the exploration, and (3) the actions which none of the players uses only because of the exploration. In Fig. 6b, (1) is on the left, (2) in the middle and (3) on the right. The third type of actions easily satisfied $|\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)| \leq \epsilon$, while for the second type, this inequality seems to eventually hold as well. The shape of the graphs suggests that the difference between $\bar{x}_{ij}(n)$ and $\tilde{x}_{ij}(n)$ will eventually get below ϵ as well, however, the 10^9 iterations we used were not sufficient for this to happen. Luckily, even if the inequality did not hold for these cases, it does not prevent the convergence of SM-MCTS to an approximate equilibrium. In the proof of Proposition 3, the term $|\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)|$ is weighted by the empirical frequencies t_j/t , resp. even t_{ij}/t . If an action is only played when exploring, this number converges to $\epsilon / (\text{number of actions})$, resp. $(\epsilon / (\text{number of actions}))^2$, so even if we had $|\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)| = 1$, we could still bound the required term by ϵ , which is needed in the proof of the respective theorem.

7.2.1 Removing exploration

In Sect. 6, we show that the computed strategy should generally improve when we disregard the samples caused by exploration. Figure 7 shows that the practical benefit is large, and

suggests that the exploration should always be removed, and we do so in all the experiments below.

7.3 Empirical convergence rate of SM-MCTS(-A) algorithms

In this section, we investigate the empirical convergence rates of the analyzed algorithms. We first compare the speeds of convergence of SM-MCTS-A and SM-MCTS and then investigate the dependence of the error of the eventual solution of the algorithms on relevant parameters. Finally, we focus on the effect of removing the exploration samples discussed in Sect. 6. In all cases, “equilibrium distance” is measured by the exploitability $\text{expl}(\cdot)$ of the evaluated strategy.

7.3.1 SM-MCTS with and without averaging

Figure 8 presents the dependence of the exploitability of the strategies produced by the algorithms on the number of executed iterations. We compare the empirical speed and quality of convergence of SM-MCTS and SM-MCTS-A with different setting of the exploration parameter, however, the samples caused by exploration are removed from the strategy for both algorithms, as suggested in Sect. 6. All iterations are executed from the root of the game. The colors (and line types) in the graphs represent different settings of the exploration parameter. SM-MCTS-A (circles) seems to always converge to the same distance from the equilibrium as SM-MCTS (triangles), regardless of the used selection function and exploration setting. The convergence of the variant with averaging is generally slower. The difference is most visible in the Anti game with Exp3 selection function, where the averaging can cause the convergence to require even ten times more iterations to reach the same distance from a NE as the algorithm without averaging. The situation is similar also in Oshi-Zumo. However, the effect is much weaker with RM selection. With the exception of the Anti game, the variants with and without averaging converge at almost the same speed.

7.3.2 Distance from the equilibrium

Even though the depth of most games in Fig. 8 was 5, even with large exploration (0.4), the algorithm was often able to find the exact equilibrium. This indicates that in practical problems, even the linear bound on the distance from the equilibrium from the example in Sect. 8 is too pessimistic.

If the game contains pure Nash equilibria, as in Anti and the used setting of Oshi-Zumo, exact equilibrium can often be found. If (non-uniform) mixed equilibria are required, the gap between the final solution and the equilibrium increases both with the depth of the game and the amount of exploration. The effect of the amount of exploration is visible in Fig. 8c, where the largest exploration prevented the algorithm from converging to the exact equilibrium. A more gradual effect is visible in Fig. 8g, h, where the exploitability seems to increase linearly with increasing exploration. Note that in all cases, the exploitability (computed as the sum $\text{expl}_1 + \text{expl}_2$) was less than $2 \cdot \varepsilon M$, where ε is the amount of exploration and M is the maximum utility value.

Figure 9 presents the exploitability of SM-MCTS with Exp3 selection and exploration $\varepsilon = 0.2$ in random games with $B = 3$ and various depths. The eventual error increases with increasing depth, but even with the depth of 6, the eventual error was on average around 0.25 and always less than 0.3.

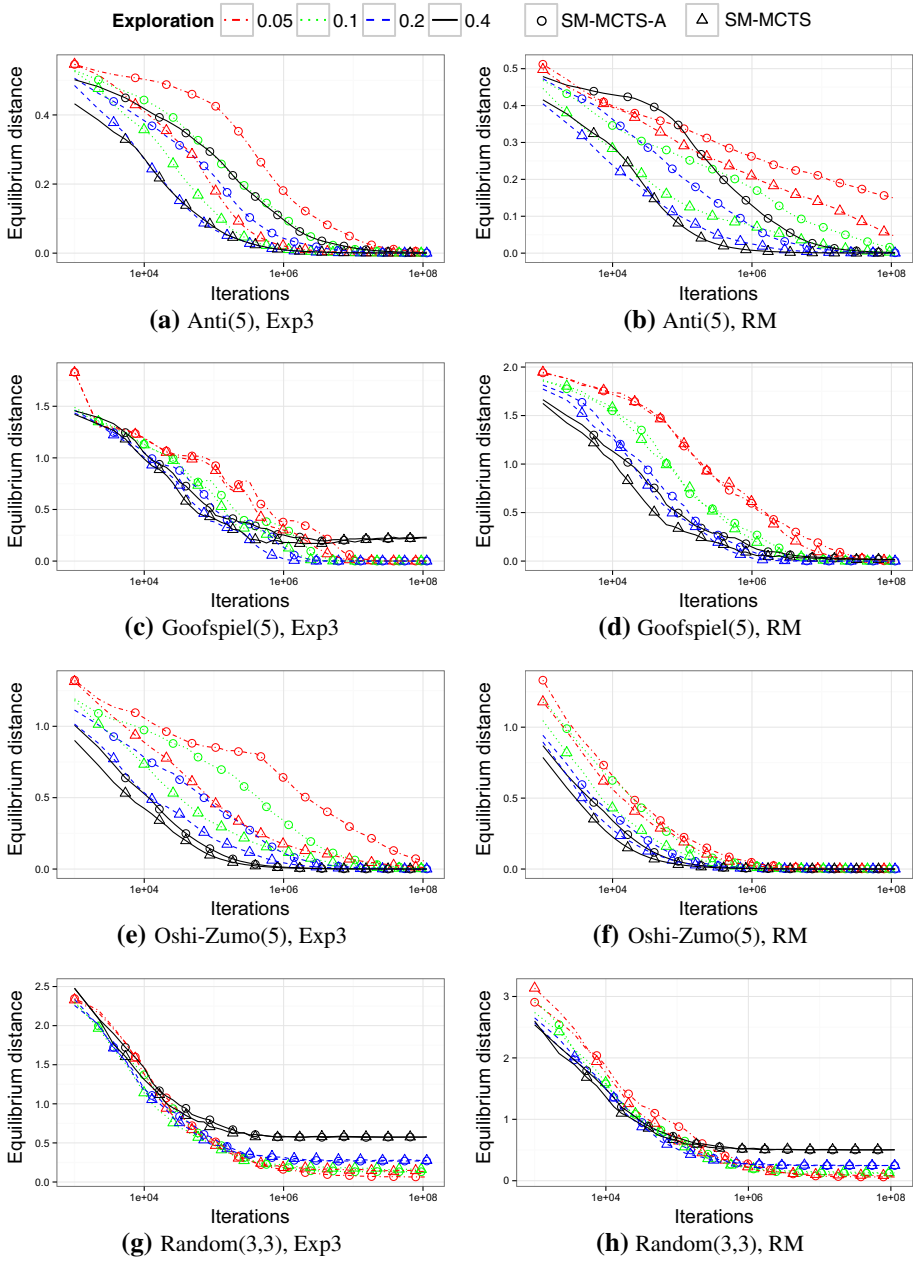



Fig. 8 Comparison of empirical convergence rates of SM-MCTS (triangles) and SM-MCTS-A (circles) with Exp3 and RM selection functions in various domains

Fig. 9 Convergence (measured by exploitability) of SM-MCTS with Exp3 selection on random games with three actions of each player in each stage and various depths



8 Conclusion

We studied the convergence properties of learning dynamics in zero-sum extensive form games with simultaneous moves, but otherwise perfect information (SMGs). These games generalize normal form games and are a special case of extensive form games. Hannan consistent action selection by each player is known to converge to a Nash equilibrium in normal form games. In general extensive form games, the convergence to a Nash equilibrium is guaranteed if both players use in each of their decision points a separate selection strategy that is Hannan consistent with respect to the counterfactual values. However, with a single simulation of the game play, the estimates of these values suffers from a large variance. This causes the simulation-based algorithms that use counterfactual values to converge slowly in larger games. We investigated whether the convergence can be achieved without using counterfactual values in SMGs. The studied learning dynamic directly corresponds to a class of Monte Carlo Tree Search algorithms commonly used for playing SMGs. There was very little pre-existing theory that would describe the behavior of the algorithms in these games and provide guarantees on their performance.

Our main results show that using a separate Hannan consistent algorithm to choose actions in each decision point of an SMG, which is the standard SM-MCTS, is not sufficient to guarantee the convergence of the learning dynamic to a Nash equilibrium. This guarantee is obtained if the used selection function, in addition to being HC, guarantees the Unbiased Payoff Observations property. We hypothesize that Exp3 and regret matching selection functions have this property, and provide experimental evidence supporting this claim. Alternatively, the convergence can be guaranteed with any HC selection, if one is willing to use SM-MCTS-A and accept the (empirically) slower convergence rate caused by averaging the back-propagated values (Sect. 7.3).

Our other results are as follows. In Sect. 6, we provide an analysis of a previously suggested improvement of SM-MCTS which proposes to remove the exploration samples from the algorithm's output. We provide formal and empirical evidence which suggests that this modification is sound and generally improves the strategy.

Table 2 summarizes some of the more detailed results. In Theorem 2 (3) we show that SM-MCTS-A (SM-MCTS) algorithm with an ε -HC selection function eventually converges at least to a $C\varepsilon$ -NE of a game, where for game depth D , C is of the order D^2 (2^D). In Sect. 4, we show that the worst-case dependence of C on D cannot be sublinear, even after the exploration is removed. This gives us both lower and upper bounds on the value of C , but

Table 2 Summary of the proven bounds on the worst-case eventual exploitability of the strategy produced by SM-MCTS(-A) run with an ε -Hannan consistent selection function

	SM-MCTS-A	SM-MCTS	
Assumptions	ε -HC	ε -HC, ε -UPO	ε -HC only
Upper bound	$2D(D + 1)\varepsilon$	$(12(2^D - 1) - 8D)\varepsilon$	Might not converge to an approx. NE at all
Lower bound	$2D\varepsilon$	$2D\varepsilon$	

it remains to determine whether these bounds are tight. The same lower bound holds for both SM-MCTS and SM-MCTS-A, which means they may in fact converge to the same distance from the equilibrium, as suggested by the experiments in Sect. 7.3.2. We hypothesize that the upper bound is not tight (and that after exploration-removal, the optimal value is $C = 2D$).

8.1 Open problems and future work

While this paper provides a significant step towards understanding learning and Monte Carlo tree search in simultaneous-move games, it also leaves some problems open. First of all, SM-MCTS is used because of the speed with which it can find a good strategy. However, a practically useful bound on this speed of convergence is missing. Furthermore, many of the guarantees presented in the paper have not been shown to be tight, so an obvious future research direction would be to improve the guarantees or show the tightness of these results (Problem 1). These could be further improved by allowing the exploration rate to depend on the position in the game tree. Finally, a better description of the properties (of the selection policy) which guarantee that SM-MCTS converges to an optimal strategy could be provided. In particular, it would be interesting to better understand the UPO property and, more generally, the propagation of information and uncertainty in game trees with simultaneous moves. This would in turn enable either proving that the common no-regret bandit algorithms indeed produce unbiased payoff observations, or finding some suitable alternative to the UPO property. Moreover, focusing on these specific (non-pathological) algorithms or their suitable modifications might simplify the derivation of tighter performance bounds.

Acknowledgements We are grateful to Tor Lattimore for helpful comments on earlier version of this paper, and to the anonymous reviewers for their patience and feedback. This research was supported by the Czech Science Foundation (Grants Nos. P202/12/2054 and 18-27483Y). Access to computing and storage facilities owned by parties and projects contributing to the National Grid Infrastructure MetaCentrum, provided under the programme “Projects of Large Infrastructure for Research, Development, and Innovations” (LM2010005), is appreciated.

Appendix

The appendix contains the proofs for those results which have not been already proven in the main text.

A Proofs of Lemma 1 and 3

We start with the proof of Lemma 1, which states that eventually, there is no difference between the empirical and the average strategies.

Lemma 1 *In the limit, the empirical frequencies and average strategies will almost surely be equal. That is, $\limsup_{t \rightarrow \infty} \max_{i \in \mathcal{S}_1} |\hat{\sigma}_1(t)(i) - \bar{\sigma}_1(t)(i)| = 0$ holds with probability 1.*

Proof It is enough to show that $\limsup_{t \rightarrow \infty} |\hat{\sigma}_1^h(t)(i) - \bar{\sigma}_1^h(t)(i)| = 0$ holds almost surely for any h and i . Without changing the limit behavior of $\hat{\sigma}$, we can assume that t_i^h is defined as the number of uses of i during iterations 1, 2, ..., t .¹⁹ Using the definitions of $\hat{\sigma}_1^h(t)(i)$ and $\bar{\sigma}_1^h(t)(i)$, we get

$$\hat{\sigma}_1^h(t)(i) - \bar{\sigma}_1^h(t)(i) = \frac{1}{t} \left(t_i - \sum_{s=1}^t \sigma_1^h(s)(i) \right) = \frac{1}{t} \sum_{s=1}^t \left(\delta_{i,i^h(s)} - \sigma_1^h(s)(i) \right),$$

where $\delta_{i,j}$ is the Kronecker delta. Using the (martingale version of) Central Limit Theorem on the sequence of random variables $X_t = \sum_{s=1}^t (\delta_{i,i^h(s)} - \sigma_1^h(s)(i))$ gives the result (the conditions clearly hold, since $\mathbf{E} [\delta_{i,i^h(t)} - \sigma_1^h(t)(i) | X_1, \dots, X_{t-1}] = 0$ implies that X_t is a martingale and $\delta_{i,i^h(t)} - \sigma_1^h(t)(i) \in [-1, 1]$ guarantees that all required moments are finite). \square

Next, we prove Lemma 3, which states that ε -Hannan consistency is not substantially affected by additional exploration.

Lemma 3 *For $\varepsilon \geq 0$, let A be an ε -Hannan consistent algorithm.*

- (i) *For any $\gamma > 0$, A^γ is an $(\varepsilon + \gamma)$ -HC algorithm with guaranteed exploration.*
- (ii) *$A^{\sqrt{\cdot}}$ is an ε -HC algorithm with guaranteed exploration.*

Proof (i): The guaranteed exploration property of A^γ is trivial. Denoting by $*$ the variables corresponding to the algorithm A^γ we get

$$\begin{aligned} r^*(t) &= \frac{1}{t} R^*(t) \leq \frac{1}{t} (1 \cdot t_{\text{ex}} + R(t - t_{\text{ex}})) \\ &= \frac{t_{\text{ex}}}{t} + \frac{R(t - t_{\text{ex}})}{t - t_{\text{ex}}} \cdot \frac{t - t_{\text{ex}}}{t}, \end{aligned}$$

where, for given $t \in \mathbb{N}$, t_{ex} denotes the number of times A^γ explored up to t -th iteration. By the strong law of large numbers we have that $\lim_{t \rightarrow \infty} t_{\text{ex}}/t = \gamma$ holds almost surely. This implies

$$\begin{aligned} \limsup_{t \rightarrow \infty} r^*(t) &\leq \limsup_{t \rightarrow \infty} \frac{t_{\text{ex}}}{t} + \limsup_{t - t_{\text{ex}} \rightarrow \infty} \frac{R(t - t_{\text{ex}})}{t - t_{\text{ex}}} \cdot \limsup_{t \rightarrow \infty} \frac{t - t_{\text{ex}}}{t} \\ &\leq \gamma + \varepsilon (1 - \gamma) \\ &\leq \gamma + \varepsilon, \end{aligned}$$

which means that A^γ is $(\varepsilon + \gamma)$ -Hannan consistent.

(ii): The guaranteed exploration follows from the fact that if the second player also uses $A^{\sqrt{\cdot}}$, the probability that both players explore at once is at least $\frac{1}{\sqrt{t}} \cdot \frac{1}{\sqrt{t}} = \frac{1}{t}$. Since $\sum \frac{1}{t} = \infty$, this will happen infinitely many times, so each joint action will be sampled infinitely many times. The proof of ε -Hannan consistency is similar to (i) and we omit it. \square

¹⁹ See the definition of t_i^h and footnote 7.

B The counterexample for Theorem 3

We first present some details related to the UPO property, and then proceed to give a proof of Lemma 6 from Sect. 4.2.2.

B.1 Examples and a sufficient condition for the UPO property

We present a few examples which are related to the UPO property and support the discussion that follows. We first give the negative examples, and then continue with the positive ones.

Example 2 (Examples related to the UPO property)

- (i) Application to Sect. 4.2.1: Assume that $(x(n))_{n=1}^\infty = (1, 0, 1, 0, 1, \dots)$ and $(w(n))_{n=1}^\infty = (1, 3, 1, 3, 1, \dots)$. Then we have $\bar{x}(n) \rightarrow \frac{1}{2}$, but $\tilde{x}(n) \rightarrow \frac{1}{4}$.
- (ii) In Sect. 4 we constructed a HC algorithm which produces a pattern nearly identical to (i) and leads to $\limsup_{n \rightarrow \infty} |\bar{x}_{ij}^{h_0}(n) - \tilde{x}_{ij}^{h_0}(n)| \geq \frac{1}{4}$ (in a certain game).
- (iii) Suppose that $w(n), x(n), n \in \mathbb{N}$ do not necessarily originate from SM-MCTS, but assume they satisfy (a) and (b):
 - (a) $w(n) \in [0, C]$ and $x(n) \in [0, 1]$ are independent random variables.
 - (b) $x(n)$ have expectations approximately constant over $n \in \mathbb{N}$:

$$\exists v \in [0, 1] \forall n \in \mathbb{N} : \mathbf{E}[x(n)] \in \left[v - \frac{\varepsilon}{2}, v + \frac{\varepsilon}{2} \right]. \tag{24}$$

Then, by strong law of large numbers, we almost surely have

$$\limsup_{n \rightarrow \infty} |\bar{x}(n) - v| \leq \frac{\varepsilon}{2} \text{ and } \limsup_{n \rightarrow \infty} |\tilde{x}(n) - v| \leq \frac{\varepsilon}{2},$$

which leads to

$$\limsup_{n \rightarrow \infty} |\bar{x}(n) - \tilde{x}(n)| \leq \varepsilon. \tag{25}$$

- (iv) The previous case can be generalized in many ways—for example it is sufficient to replace bounded weights $w(n)$ by ones satisfying

$$\exists q \in (0, 1) \forall n \forall i, j : \mathbf{Pr}[w(n) \geq k] \leq q^k$$

(an assumption which holds with $q = \gamma / |\mathcal{A}_1(h)|$ when $w(n), x(n)$ originate from SM-MCTS with fixed exploration). Note that the weights still need to be sufficiently independent, to avoid the pathological behavior from Sect. 4.

- (v) In Sect. 7 we provide empirical evidence which suggests that when the variables $x(n), w(n)$ originate from SM-MCTS with Exp3 or RM selection policy, then the assertion (25) holds as well (and thus these two ε -HC algorithms are likely ε -UPO).

Parts (i) and (ii) from Example 2 show that the implication (A is ε -HC \implies A is $C\varepsilon$ -UPO) does not hold, no matter how huge is the constant $C > 0$. On the other hand, (iii) and (iv) suggest that it is possible to prove that specific ε -HC algorithms are ε -UPO. However, the guarantees we have about the behavior of, for example, Exp3 are much weaker than the assumptions made in (iii)—there is no independence between $w_{ij}(n), x_{ij}(m), m, n \in \mathbb{N}$, at best we can use some martingale theory. Moreover, even in nodes $h \in \mathcal{H}$ with $d_h = 1$, we only have $\limsup_{n \rightarrow \infty} |\bar{x}_{ij}^h(n) - v_{ij}^h| \leq \varepsilon$, instead of assumption (24) from (iii) (b).

The observations above suggest that the proof that specific ε -HC algorithms are ε -UPO will not be trivial. However, we can at least state a sufficient condition for an algorithm to be ε -UPO, which is not dependent on the tree structure of simultaneous-move games.

Remark 1 (*Sufficient condition for ε -UPO*) Let A be a bandit algorithm, $K, L \in \mathbb{N}$ and for each $i \leq K, j \leq L$, let $(x_{ij}(n))_{n=1}^\infty$ be an arbitrary sequence in $[0, 1]$. Assume that player 1 uses A to choose actions $i \in \{1, \dots, K\}$ and player 2 uses a different instance of A to choose actions $j \in \{1, \dots, L\}$, where the reward for playing actions (i, j) at time t is $x_{ij}(t_{ij})$ (resp. $1 - x_{ij}(t_{ij})$ for player 2). As before, denote

$$\begin{aligned} \bar{x}_{ij}(n) &:= \frac{1}{t} \sum_{m=1}^n x_{ij}(m), \\ w_{ij}(m) &:= \left| \{t \in \mathbb{N} \mid t_{ij} = m \ \& \ j(t) = j\} \right|, \\ \tilde{x}_{ij}(n) &:= \frac{1}{\sum_{m=1}^n w_{ij}(m)} \sum_{m=1}^n w_{ij}(m) x_{ij}(m). \end{aligned}$$

Clearly, if

$$\forall i \leq K \ \forall j \leq L : \limsup_{n \rightarrow \infty} |\bar{x}_{ij}(n) - \tilde{x}_{ij}(n)| \leq \varepsilon \text{ a.s.}$$

holds for all such sequences $(x_{ij}(n))_{n=1}^\infty$, then A is ε -UPO.

This condition should be easier to verify in practice than the original ε -UPO definition. On the other hand, the quantification over all possible sequences $(x_{ij}(n))_{n=1}^\infty$ might be a too strong requirement. Moreover, it would be even better to have an equivalent—or at least sufficient, yet not too strong—condition for an algorithm to be ε -UPO which only operates with one player.

B.2 Details related to the counterexample for Theorem 3

In Sect. 4 (Lemma 2) we postulated the existence of algorithms, which behave similarly to those from Sect. 4.2.1, but unlike those from Sect. 4.2.1, the new algorithms are ε -HC:

Lemma 2 *Let G be the game from Fig. 2. Then for each $\varepsilon > 0$ there exist ε -HC algorithms A^I, A_1^J, A_2^J , such that when these algorithms are used for SM-MCTS in G , the resulting average strategy $\bar{\sigma}(t)$ converges to $\bar{\sigma}$, where $\bar{\sigma}^I = \bar{\sigma}_1^J = \bar{\sigma}_2^J = (\frac{1}{2}, \frac{1}{2})$.*

First, we define these algorithms, and then we prove their properties in Lemma 6. Lemma 2, then follows directly from Lemma 6.

Remark 2 The intention behind the notation which follows is that ch denotes how many times the other player “cheated”, while \bar{ch} is the average ratio of cheating in following the cooperation pattern. The variables \tilde{ch} and \hat{ch} then serve as estimates of ch and \bar{ch} . We present the precise definitions below. The nodes I, J, actions X, Y, L, R, U and D and the respective payoffs refer to the game G from Fig. 2.

Definition of the algorithm A_1^J :

Fix an increasing sequence of integers b_n and repeat for $n \in \mathbb{N}$:

1. *Buffer building B_n* Play according to some ε -HC algorithm for b_n iterations (continuing with where we left of in the $(n - 1)$ -th buffer building phase).
2. *Cooperation C_n* Repeat U, U, D, D for $t = 1, 2, \dots$ and expect the other player to repeat L, R, R, L . At each iteration t , with probability ε , check whether the other player is cooperating—that is play the opposite of the pattern-prescribed action and if the payoff does not correspond to the expected pattern the second player should be following, set $\tilde{c}h(t) = \frac{1}{\varepsilon}$. If the other player passes this check, or if we did not perform it, set $\tilde{c}h(t) = 0$.
3. *End of cooperation (might not happen)* While executing step 2, denote $\tilde{c}h(t) := \frac{1}{t} \sum_{s=1}^t \tilde{c}h(s)$. Once t satisfies

$$\varepsilon \cdot \frac{b_n}{b_n + t} + 1 \cdot \frac{t}{b_n + t} \geq 2\varepsilon, \tag{26}$$

we check at each iteration whether the estimate $\tilde{c}h(t)$ threatens to exceed 2ε during the next iteration or not. If it does, we end the cooperation phase, set $n := n + 1$ and continue by the next buffer building phase.

4. *Simulation of the other player* While repeating steps 1, 2 and 3, we simulate the other player’s algorithm A_2^J (this is possible since from the knowledge of our action and the received payoff we can recover the adversary’s action). If it ends the cooperation phase and starts the next buffer building phase, we do the same.
5. Unless the cooperation phase is terminated, we stay in phase C_n indefinitely.

Definition of A_2^J : The algorithm A_2^J is identical to A_1^J , except for the fact that it repeats the pattern U, D, D, U instead of L, R, R, L and expects the other player to repeat L, R, R, L .

Definition of A^I : The algorithm A^I is a straightforward modification of A_1^J (replacing 2ε and ε by 2ε and 3ε in Eq. (26))—it repeats the sequence Y, X, X, Y and expects to receive payoffs $0, 0, 0, \dots$ whenever playing X and payoffs $1, 0, 1, 0, \dots$ when playing Y . However, whenever it deviates from the Y, X, X, Y pattern to check whether these expectations are met, it plays the same action once again (to avoid disturbing the payoff pattern for Y).

Remark 3 The steps 2 and 3 from the algorithm description are correctly defined because the opponent’s action choice can be recovered from the knowledge of our action choice and the resulting payoff. Regarding step 3, we note that the condition here is trivially satisfied for $t = 1, \dots, \varepsilon b_n$, so the length t_n of the cooperation phases C_n tends to infinity as $n \rightarrow \infty$, regardless of the opponent’s actions.

Lemma 6 (1) *When facing each other, the average strategies of algorithms A^I, A_1^J and A_2^J will converge to the sub-optimal strategy $\sigma^I = \sigma_1^I = \sigma_2^I = (\frac{1}{2}, \frac{1}{2})$. However the algorithms will suffer regret no higher than $C\varepsilon$ for some $C > 0$ (where C is independent of ε).*

(2) *There exists a sequence b_n (controlling the length of phases B_n), such that when facing a different adversary, the algorithms suffer regret at most $C\varepsilon$. Consequently A^I, A_1^J and A_2^J are $C\varepsilon$ -Hannan consistent.*

Proof Part (1): Note that, disregarding the checks made by the algorithms, the average strategies in cooperation phases converge to $(\frac{1}{2}, \frac{1}{2})$. Furthermore, the probability of making the checks is the same, regardless on which action is supposed to be played next. Therefore

if the algorithms eventually settle in cooperative phase, the average strategies converge to $(\frac{1}{2}, \frac{1}{2})$.

Step (i): The conclusion of (2) holds for $A_p^J, p = 1, 2$.

We claim that both algorithms $A_p^J, p = 1, 2$ will eventually settle in the cooperative mode, thus generating the payoff sequence 1, 0, 1, 0 during at least $(1 - \varepsilon)^2$ -fraction of the iterations and generating something else at the remaining $1 - (1 - \varepsilon)^2 < 2\varepsilon$ iterations. It is then immediate that the algorithms A_p^J suffer regret at most 2ε (since as observed in Sect. 4.2.1, the undisturbed pattern produces no regret).

Proof of (i): If the other player uses the same algorithm, we have $\mathbf{E}[\tilde{c}h(t)] = \varepsilon$ and $\mathbf{Var}[\tilde{c}h(t)] \leq \frac{1}{\varepsilon} < \infty$, thus by the Strong Law of Large Numbers $\tilde{c}h(t) \rightarrow \varepsilon$ holds almost surely. In particular, there exists $t_0 \in \mathbb{N}$ such that

$$\Pr[\forall t \geq t_0 : \tilde{c}h(t) \leq 2\varepsilon] > 0.$$

In Remark 3 we observed that the cooperative phase always lasts at least εb_n steps. Therefore once we have $b_n \geq \frac{1}{\varepsilon}t_0$, there is non-zero probability both players will stay in n -th cooperative phase forever. If they do not, they have the same positive probability of staying in the next cooperative phase and so on—by Borel–Cantelli lemma, they will almost surely stay in C_{n_0} for some $n_0 \in \mathbb{N}$.

Step (ii): A^I will eventually settle in the cooperative mode.

Proof of (ii): This statement can be proven by a similar argument as (i). The only difference is that instead of checking whether the other player is cheating, we check whether the payoff sequence is the one expected by A^I . The fact that A^I settles in cooperative mode then immediately implies that A^I will repeat the Y, X, X, Y pattern during approximately a $(1 - 2\varepsilon)$ -fraction of the iterations (we check with probability ε , and we always check twice). It is then also immediate that A^I will receive the expected payoffs during at least $(1 - \varepsilon)$ -fraction of the iterations. (The payoff is always 0, as expected, when X is played, and it is correct in at least $1 - 2\varepsilon$ cases when playing Y. X and Y are both played with same frequency, which gives the result.)

Step (iii) : If $A_p^J, p = 1, 2$ and A^I settle in cooperative mode, then A^I suffers regret at most $C\varepsilon$.

Proof of (iii): Recall the definition (8) of $x_i^h(t)$, definition of t_i^h which follows it, and the notation introduced in (18). We denote by $i(t) := i^I(t)$ the t -th action chosen by A^I and by $(i^*(s))_s$ the 4-periodic of sequence of “ideal case actions” starting with Y, X, X, Y. Recall that $x^J(n)$ is the payoff obtained during the n -th visit of J and $x_Y^I(t) = x^J(t_Y)$ is the payoff we would receive for playing Y in I at time t (and $x_X^I(t) = 0$ is always zero). Finally we denote by $(x_Y^*(s))_s$ the 4-periodic of sequence of “ideal case payoffs” starting with 1, 0, 0, 0. Our goal is to show that the average regret

$$r^I(t) = \frac{1}{t} \left(\sum_{s=1}^t x_Y^I(s) - \sum_{s=1}^t x_{i(s)}^I(s) \right)$$

is small. As we already observed in Sect. 4.2.1, if neither the payoff pattern $(X_Y^*(s))_s = (1, 0, 1, 0, \dots)$ coming from J , nor the action pattern $(i^*(s))_s = (Y, X, X, Y, \dots)$ at I are disturbed, then we have $(i(s)) = (i^*(s))$ and $(x_Y^I(s)) = (x_Y^*(s))$, and therefore A^I suffers

no regret. As long as the algorithms stay in the cooperative mode, the only way the pattern can be disturbed is when the random checks occur. The checks at J occur with probability ε for each player, which increases the average asymptotic regret by no more than 2ε (these checks will actually cancel out in expectation, but we ignore this). Moreover, the combination of the 2-periodicity of the reward-pattern at J and the fact that A^I checks twice in a row guarantees that any check performed at I only affects 2 the two iterations during which it occurs. In the worst case, this increases the average asymptotic regret by no more than 2ε . Hence the total average regret is at most 4ε in the limit.

Part (2): Algorithms A_p^J , $p = 1, 2$, and A^I are $C\varepsilon$ -HC.

a) Firstly, we assume that both players stick to their assigned patterns during at least $(1 - 2\varepsilon)$ -fraction of iterations (in other words, assume that $\limsup \bar{c}h(t) \leq 2\varepsilon$). By the same argument as in (1), we can show that the algorithms will then suffer regret at most $C\varepsilon$.

b) On the other hand, if $\limsup \bar{c}h(t) > 2\varepsilon$, the algorithm will almost surely keep switching between B_n and C_n (consequence of Strong Law of Large Numbers). Denote by $r_{b,n}$ and $r_{c,n}$ the regret from phases B_n and C_n , recall that b_n, t_n are the lengths of these phases and set

$$r_n := \frac{b_n}{b_n + t_n} r_{b,n} + \frac{t_n}{b_n + t_n} r_{c,n} = \text{overall regret in } B_n \text{ and } C_n \text{ together.}$$

Finally, let $r = \limsup r^J(t)$ denote the bound on the limit of regret of A_p^J . We need to prove that $r \leq C\varepsilon$. To do this, it is sufficient to show that $\limsup_n r_n$ is small—thus our goal will be to prove that if the sequence b_n increases quickly enough, then $\limsup r_n \leq C\varepsilon$ holds almost surely. Denote by (F_n) the formula

$$\forall t \geq \frac{1}{\varepsilon} b_n : \tilde{c}h(t) \leq 2\varepsilon \implies \bar{c}h(t) \leq 3\varepsilon.$$

We know that $|\tilde{c}h - \bar{c}h| \rightarrow 0$ a.s., therefore we can choose b_n such that

$$\Pr [(F_n) \text{ holds}] \geq 1 - 2^{-n}$$

holds. Since $\sum 2^{-n} < \infty$, Borel–Cantelli lemma gives that (F_n) will hold for all but finitely many $n \in \mathbb{N}$. Note that if (F_n) holds for both players, then their empirical strategy is at most 3ε away from the NE strategy, and thus $r_{c,n} \leq C\varepsilon$. Since r_n is a convex combination of $r_{b,n}$ and $r_{c,n}$ and in B_n we play ε -consistently, we can compute

$$r \leq \limsup r_n \leq \max \{ \limsup r_{b,n}, \limsup r_{c,n} \} \leq \max \{ \varepsilon, C\varepsilon \} = C\varepsilon.$$

The proof of $C\varepsilon$ -Hannan consistency of A^I is analogous. □

C Proofs related to the convergence of SM-MCTS(-A)

C.1 Convergence of SM-MCTS-A

In this section, we give the proofs for lemmas which were used to prove Theorem 2. We begin with Lemma 5, which established a connection between average payoff g and game value v for matrix games.

Lemma 5 *Let $\varepsilon \geq 0$ be a real number. If both players in a matrix game M are ε -Hannan consistent, then the following inequalities hold almost surely:*

$$v^M - \varepsilon \leq \liminf_{t \rightarrow \infty} g^M(t) \leq \limsup_{t \rightarrow \infty} g^M(t) \leq v^M + \varepsilon, \tag{13}$$

$$v^M - 2\varepsilon \leq \liminf_{t \rightarrow \infty} u_1^M(\hat{\sigma}_1(t), br) \quad \& \quad \limsup_{t \rightarrow \infty} u_1^M(br, \hat{\sigma}_2(t)) \leq v^M + 2\varepsilon. \tag{14}$$

Proof It is our goal to draw conclusion about the quality of the empirical strategy based on information about the performance of our ε -HC algorithm. Ideally, we would like to relate the utility $u_1^M(\hat{\sigma})$ to the average payoff $g^M(t)$. However, as this is generally impossible, we can do the next best thing:

$$\begin{aligned} u_1^M(br, \hat{\sigma}_2(t)) &= \max_i \sum_j \hat{\sigma}_2(t)(j) v_{ij}^M = \max_i \sum_j \frac{t_j}{t} v_{ij}^M = \frac{1}{t} \max_i \sum_j t_j v_{ij}^M \\ &= \frac{1}{t} \max_i \sum_{s=1}^t v_{ij(s)}^M = \frac{1}{t} G_{\max}(t) = g_{\max}(t). \end{aligned}$$

Step 1: Let $\eta > 0$. Using the ε -HC property gives us the existence of such t_0 that $g_{\max}^M(t) - g^M(t) < \varepsilon + \frac{\eta}{2}$ holds for all $t \geq t_0$, which is equivalent to $g^M(t) > g_{\max}^M - (\varepsilon + \frac{\eta}{2})$. However, in our zero-sum matrix game setting, g_{\max}^M is always at least v^M , which implies that $g^M(t) > v^M - (\varepsilon + \frac{\eta}{2})$. Using the same argument for player 2 gives us that $g^M(t) < v^M + \varepsilon + \frac{\eta}{2}$. Therefore we have the following statement, which proves the inequalities (13):

$$\forall t \geq t_0 : v^M - (\varepsilon + \frac{\eta}{2}) < g^M(t) < v^M + \varepsilon + \frac{\eta}{2} \text{ holds almost surely.}$$

Step 2: We assume, for contradiction with inequalities (14), that with non-zero probability, there exists an increasing sequence of time steps $t_n \nearrow \infty$, such that $u_1^M(br, \hat{\sigma}_2(t_n)) \geq v^M + 2\varepsilon + \eta$ for some $\eta > 0$. Combing this with the inequalities we proved above, we see that

$$\begin{aligned} \limsup_{t \rightarrow \infty} r^M(t) &\geq \limsup_{n \rightarrow \infty} r^M(t_n) = \limsup_{n \rightarrow \infty} (g_{\max}^M(t_n) - g^M(t_n)) \\ &= \limsup_{n \rightarrow \infty} (u_1^M(br, \hat{\sigma}_2(t_n)) - g^M(t_n)) \\ &\geq v^M + 2\varepsilon + \eta - (v^M + \varepsilon + \eta/2) = \varepsilon + \eta > \varepsilon \end{aligned}$$

holds with non-zero probability, which is in contradiction with ε -Hannan consistency. \square

Proposition 1 *Let $c, \varepsilon > 0$ and let (\tilde{M}) be a repeated matrix game with $c\varepsilon$ -bounded distortion, played by two ε -HC players. Then the following inequalities hold almost surely:*

$$v^M - (c + 1)\varepsilon \leq \liminf_{t \rightarrow \infty} g^{\tilde{M}}(t) \leq \limsup_{t \rightarrow \infty} g^{\tilde{M}}(t) \leq v^M + (c + 1)\varepsilon, \tag{15}$$

$$v^M - 2(c + 1)\varepsilon \leq \liminf_{t \rightarrow \infty} u_1^M(\hat{\sigma}_1, br) \leq \limsup_{t \rightarrow \infty} u_1^M(br, \hat{\sigma}_2) \leq v^M + 2(c + 1)\varepsilon. \tag{16}$$

Remark 4 In the following proof, and in the proof of Proposition 3, we will be working with regrets, average payoffs and other quantities related to matrix games with bounded distortion, in which we have two sets of rewards—the “precise” rewards v_{ij}^M corresponding to the matrix M and the “distorted” rewards $v_{ij}^{\tilde{M}}(t)$. We denote the variables related to the distorted rewards

$v_{ij}^{\tilde{M}}(t)$ with the superscript ‘ \tilde{M} ’ (for example $g_{\max}^{\tilde{M}}(t) = \max_i \frac{1}{t} \sum_{s=1}^t v_{ij(s)}^{\tilde{M}}(s)$) and use the superscript ‘ M ’ for variables related to v_{ij}^M (for example $g_{\max}^M(t) = \max_i \frac{1}{t} \sum_{s=1}^t v_{ij(s)}^M$).

Proof of Proposition 1 This proposition strengthens the result of Lemma 5 and its proof will also be similar. The only additional technical ingredient is the inequality (27).

Since (\tilde{M}) is a repeated game with $c\varepsilon$ -bounded distortion, there almost surely exists t_0 , such that for all $t \geq t_0$, $|v_{ij}^{\tilde{M}}(t) - v_{ij}^M| \leq c\varepsilon$ holds. This leads to

$$|g_{\max}^{\tilde{M}}(t) - g_{\max}^M(t)| \leq \max_i \left| \frac{1}{t} \sum_{s=1}^t (v_{ij(s)}^{\tilde{M}}(s) - v_{ij(s)}^M) \right| \leq \frac{t_0}{t} + c\varepsilon \cdot \frac{t - t_0}{t} \xrightarrow{t \rightarrow \infty} c\varepsilon. \tag{27}$$

The remainder of the proof contains no new ideas, and it is nearly exactly the same as the proof of Lemma 5, therefore we just note what the two main steps are:

Step 1: Hannan consistency gives us that

$$\forall \eta > 0 \exists t_0 \in \mathbb{N} \forall t \geq t_0 : g^{\tilde{M}}(t) \geq v^M - (\varepsilon + \eta) - |g_{\max}^{\tilde{M}}(t) - g_{\max}^M(t)| \text{ holds a.s.,}$$

from which we deduce the inequalities

$$v^M - (\varepsilon + c\varepsilon) \leq \liminf_{t \rightarrow \infty} g^{\tilde{M}}(t) \leq \limsup_{t \rightarrow \infty} g^{\tilde{M}}(t) \leq v^M + \varepsilon + c\varepsilon.$$

Step 2: For contradiction we assume that there exists an increasing sequence of time steps $t_n \nearrow \infty$, such that

$$u_1^M(br, \hat{\sigma}_2(t_n)) \geq v^M + 2(c + 1)\varepsilon + \eta \tag{28}$$

holds for some $\eta > 0$. Using the identity $g_{\max}^M(t_n) = u_1^M(br, \hat{\sigma}_2(t_n))$ and inequalities (27) and (28), we then compute that the regret $r^M(t_n)$ is too high, which completes the proof:

$$\begin{aligned} \varepsilon &\geq \limsup_{t \rightarrow \infty} r^{\tilde{M}}(t) \geq \limsup_{n \rightarrow \infty} r^{\tilde{M}}(t_n) = \limsup_{n \rightarrow \infty} (g_{\max}^{\tilde{M}}(t_n) - g^{\tilde{M}}(t_n)) \\ &= \limsup_{n \rightarrow \infty} (g_{\max}^M(t_n) - (g_{\max}^M(t_n) - g_{\max}^{\tilde{M}}(t_n)) - g^{\tilde{M}}(t_n)) \\ &\geq \limsup_{n \rightarrow \infty} u_1^M(br, \hat{\sigma}_2(t_n)) - \limsup_{n \rightarrow \infty} |g_{\max}^M(t_n) - g_{\max}^{\tilde{M}}(t_n)| \\ &\quad - \limsup_{n \rightarrow \infty} g^{\tilde{M}}(t_n) \\ &\stackrel{(27),(28)}{\geq} (v^M + 2(c + 1)\varepsilon + \eta) - c\varepsilon - (v^M + \varepsilon + c\varepsilon) \\ &= \varepsilon + \eta > \varepsilon. \end{aligned}$$

□

C.2 Convergence of SM-MCTS

In Sect. 5.2, we stated Theorem 3 which claimed that if a SM-MCTS algorithm uses a selection function which is both ε -UPO and ε -HC, it will find an approximate equilibrium. To finish the proof of Theorem 3, it remains to prove Proposition 3, which establishes a connection

between regrets R^h and R^{M_h} of the selection function with respect to rewards $x_i^h(t)$, resp. w.r.t. the matrix game $M_h = (v_{ij}^h)_{ij}$. The goal is to show that if $R^h(T)$ is small and algorithm A is ε -UPO, then the regret $R^{M_h}(T)$ is small as well.

Proposition 3 For $\varepsilon, c \geq 0$ let an algorithm A be ε -HC and $h \in \mathcal{H}$. If A is ε -UPO and $\limsup_{n \rightarrow \infty} |\bar{x}_{ij}^h(n) - v_{ij}^h| \leq c\varepsilon$ holds a.s. for each i, j , then

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \left(\max_i \sum_{s=1}^t v_{ij(s)}^h - \sum_{s=1}^t v_{i(s)j(s)}^h \right) \leq 2(c+1)\varepsilon \text{ holds a.s.} \tag{20}$$

In particular, the regret $r^{M_h}(t)$ of A with respect to the matrix game $M_h = (v_{ij}^h)_{ij}$ satisfies

$$\limsup_{t \rightarrow \infty} r^{M_h}(t) \leq 2(c+1)\varepsilon.$$

Proof Let ε, c, A and h be as in the proposition. Since in this proof we will only deal with actions taken in h , we will temporarily omit the upper index ‘ h ’ in variables $t_i^h, t_j^h, t_{ij}^h, i^h(t)$ and $j^h(t)$. As in Remark 4, R^{M_h} is the regret of player 1 with respect to the matrix game $M_h = (v_{ij}^h)_{ij}$:

$$R^{M_h}(T) = \max_{i^*} \sum_{t=1}^T v_{i^*j(t)}^h - \sum_{t=1}^T v_{i(t)j(t)}^h =: \max_{i^*} S_{i^*}^{M_h}(T).$$

Recall that by (9), we have $x_i^h(t) = x_{i(j(t))}^h(t_{ij(t)})$. This allows us to rewrite the regret $R^{\tilde{M}_h} = R^h$ that player 1 is attempting to minimize:

$$\begin{aligned} R^h(T) &\stackrel{\text{def.}}{=} \max_{i^*} \sum_{t=1}^T x_{i^*}^h(t) - \sum_{t=1}^T x^h(t) \\ &\stackrel{(9)}{=} \max_{i^*} \sum_{t=1}^T x_{i^*j(t)}^h(t_{i^*j(t)}) - \sum_{t=1}^T x_{i(t)j(t)}^h(t_{i(t)j(t)}) \\ &=: \max_{i^*} S_{i^*}^h(T). \end{aligned}$$

Since A is ε -HC, we have $\limsup_T R^h(T)/T \leq \varepsilon$ a.s. and our goal infer that $\limsup_T R^{M_h}(T)/T \leq 2(c+1)\varepsilon$ holds a.s. as well. Let i^* be an action of player 1. Note the following two facts:

$$T_j = \sum_{m=1}^{T_{i^*j}} w_{i^*j}^h(m) \quad \dots \text{ by definition of } w_{ij}^h(n), \tag{a}$$

$$(j(t) = j \ \& \ t_{i^*j(t)} = m) \implies x_{i^*j(t)}^h(t_{i^*j(t)}) = x_{i^*j}^h(m) \quad \dots \text{ by (9)} \tag{b}$$

We can rewrite $S_{i^*}^h(T)$ as follows:

$$\begin{aligned}
 S_{i^*}^h(T) &= \sum_{t=1}^T x_{i^*j(t)}^h(t_{i^*j(t)}) - \sum_{t=1}^T x_{i(t)j(t)}^h(t_{i(t)j(t)}) \\
 &\stackrel{(b)}{=} \sum_j \sum_{m=1}^{T_{i^*j}} x_{i^*j}^h(m) |\{t \leq T \mid t_{i^*j(t)} = m \ \& \ j(t) = j\}| - \sum_{i,j} \sum_{m=1}^{T_{ij}} x_{ij}^h(m) \\
 &\stackrel{(a)}{=} \sum_j \frac{T_j}{\sum_{m=1}^{T_{i^*j}} w_{i^*j}(m)} \sum_{m=1}^{T_{i^*j}} x_{i^*j}^h(m) * w_{ij}(m) - \sum_{i,j} \frac{T_{ij}}{T_{ij}} \sum_{m=1}^{T_{ij}} x_{ij}^h(m) \\
 &= \sum_j T_j \tilde{x}_{i^*j}^h(T_{i^*j}) - \sum_{i,j} T_{ij} \bar{x}_{ij}^h(T_{ij}) \\
 &= \sum_j T_j (v_{i^*j}^h + \tilde{x}_{i^*j}^h(T_{i^*j}) - v_{i^*j}^h) - \sum_{i,j} T_{ij} (v_{ij}^h + \bar{x}_{ij}^h(T_{ij}) - v_{ij}^h) \\
 &=: \sum_j \sum_{m=1}^{T_j} v_{i^*j}^h - \sum_{i,j} \sum_{m=1}^{T_{ij}} v_{ij}^h + X_{i^*}^h(T) \\
 &= \sum_{t=1}^T v_{i^*j(t)}^h - \sum_{t=1}^T v_{i(t)j(t)}^h + X_{i^*}^h(T), \\
 &= S_{i^*}^{M_h}(T) + X_{i^*}^h(T)
 \end{aligned}$$

where

$$X_{i^*}^h(T) := \sum_j T_j (\tilde{x}_{i^*j}^h(T_{i^*j}) - v_{i^*j}^h) - \sum_{i,j} T_{ij} (\bar{x}_{ij}^h(T_{ij}) - v_{ij}^h).$$

In particular, we can bind the regret $R^h(T)$ as

$$R^{M_h} = \max_{i^*} S_{i^*}^{M_h}(T) = \max_{i^*} (S_{i^*}^h(T) - X_{i^*}^h(T)) \leq R^h(T) + \max_{i^*} |X_{i^*}^h(T)|.$$

Clearly $X_{i^*}^h(T)$ satisfies

$$\left| \frac{X_{i^*}^h(T)}{T} \right| \leq \sum_j \frac{T_j}{T} |\tilde{x}_{i^*j}^h(T_{i^*j}) - v_{i^*j}^h| + \sum_{i,j} \frac{T_{ij}}{T} |\bar{x}_{ij}^h(T_{ij}) - v_{ij}^h|.$$

Using the ε -UPO property and the assumption that $\limsup_n |\bar{x}_{ij}^h(n) - v_{ij}^h| \leq c\varepsilon$ holds a.s. for each i, j , we get

$$\begin{aligned}
 \limsup_{T \rightarrow \infty} \left| \frac{X_{i^*}^h(T)}{T} \right| &\leq \limsup_{T \rightarrow \infty} \sum_j \frac{T_j}{T} |\tilde{x}_{i^*j}^h(T_{i^*j}) - v_{i^*j}^h| \\
 &\quad + \sum_{i,j} \frac{T_{ij}}{T} |\bar{x}_{ij}^h(T_{ij}) - v_{ij}^h|
 \end{aligned}$$

$$\begin{aligned} &\leq (c + 1) \varepsilon \limsup_{T \rightarrow \infty} \sum_j \frac{T_j}{T} \\ &\quad + c\varepsilon \limsup_{T \rightarrow \infty} \sum_{i,j} \frac{T_{ij}}{T} = (2c + 1) \varepsilon. \end{aligned}$$

Consequently, this implies that

$$\begin{aligned} \limsup_{T \rightarrow \infty} R^{M_h}(T) / T &\leq \limsup_{T \rightarrow \infty} R^h(T) / T + \limsup_{T \rightarrow \infty} \max_{i^*} X_{i^*}^h(T) / T \\ &\leq \varepsilon + (2c + 1) \varepsilon = 2(c + 1) \varepsilon \end{aligned}$$

holds almost surely, which is what we wanted to prove. □

D The proof of Proposition 6

Proposition 6 *Let M be a matrix game, A a HC algorithm, $\gamma > 0$, and A^γ the “ γ -exploration modification” of A from Definition 5. Let $\bar{\sigma}(t)$ be the average strategy obtained if both players use A^γ in M , and let $\bar{\mu}(t)$ be as in (23). Then for each p we have*

- (i) $\limsup_t \text{expl}_p(\bar{\sigma}_p(t)) \leq 2\gamma$,
- (ii) $\limsup_t \text{expl}_p(\bar{\mu}_p(t)) \leq \gamma$.

Proof (Sketch) By Lemma 1, it suffices to prove the result for the empirical frequencies $\hat{\sigma}(t)$ and $\hat{\mu}(t)$. By definition of exploitability, the first case follows from (14) in Lemma 5 (using $\varepsilon = 0$). Since the proof of the second case is similar to that of (14) in Lemma 5, we only give the main ideas. We focus on bounding the second player’s exploitability.

The first ingredient is that $u_1^M(br, \hat{\mu}_2(t))$ is equal to $g_{\max}^{M,A}(t)$, the analogy of $g_{\max}^M(t)$ where we only consider the iterations where the second player didn’t explore. Secondly, using the γ -Hannan-consistency of player 1’s A^γ , we bound the asymptotic difference between $g_{\max}^M(t)$ and $g^M(t)$. Since the second player explores independently of first player’s algorithm, the same bound holds for $g_{\max}^{M,A}(t)$ and $g^{M,A}(t)$. Thirdly, the Hannan-consistency of player 2’s A implies that $g^{M,A}(t) \geq v^M$ holds in the limit.

Finally, we combine these observations and show that if the exploitability of $\hat{\mu}_2(t)$ was higher than γ in the limit—that is, if $u_1^M(br, \hat{\mu}_2(t))$ was higher than $v^M + \gamma$ —then the regret $g_{\max}^{M,A}(t) - g^{M,A}(t) = u_1^M(br, \hat{\mu}_2(t)) - g^{M,A}(t)$ of player 1 would be higher than γ . □

References

Abbasi, Y., Bartlett, P. L., Kanade, V., Seldin, Y., & Szepesvári, C. (2013). Online learning in Markov decision processes with adversarially chosen transition probability distributions. In *Advances in neural information processing systems* (pp. 2508–2516).

Auer, P., Cesa-Bianchi, N., Freund, Y., & Schapire, R. E. (1995). Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of 36th annual symposium on foundations of computer science* (pp. 322–331). IEEE.

Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2–3), 235–256.

Auer, P., Cesa-Bianchi, N., Freund, Y., & Schapire, R. E. (2003). The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1), 48–77.

Blum, A., & Mansour, Y. (2007). Learning, regret minimization, and equilibria. In *Algorithmic game theory, chapter 4*. Cambridge University Press.

- Božanský, B., Lisý, V., Čermák, J., Vitek, R., & Pěchouček, M. (2013). Using double-oracle method and serialized alpha–beta search for pruning in simultaneous move games. In *Proceedings of the twenty-third international joint conference on artificial intelligence (IJCAI)* (pp. 48–54). AAAI Press.
- Božanský, B., Lisý, V., Lanctot, M., Čermák, J., & Winands, M. H. M. (2016). Algorithms for computing strategies in two-player simultaneous move games. *Artificial Intelligence*, 237, 1–40.
- Bowling, M., & Veloso, M. (2002). Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136(2), 215–250.
- Brown, N., & Sandholm, T. (2018). Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374), 418–424.
- Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., et al. (2012). A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1), 1–43.
- Bubeck, S., & Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Machine Learning*, 5(1), 1–122.
- Buro, Michael. (2004). Solving the Oshi-zumo game. In Van Den Herik H. J., H. Iida, & E. A. Heinz (Eds.) *Advances in computer games, volume 135 of IFIP—The International Federation for Information Processing* (pp. 361–366). Springer US, 2004. ISBN 978-1-4757-4424-8. https://doi.org/10.1007/978-0-387-35706-5_23.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. New York: Cambridge University Press. (ISBN 978-0521841085).
- Coquelin, P.-A., & Munos, R. (2007). Bandit algorithms for tree search. In *Proceedings of the twenty-third conference annual conference on uncertainty in artificial intelligence (UAI-07)* (pp. 67–74). AUAI Press.
- Coulom, R. (2007). Efficient selectivity and backup operators in Monte-Carlo tree search. In H. J. van den Herik, P. Ciancarini, & H. H. L. M. Donkers (Eds.), *Computers and games. CG 2006. Lecture Notes in Computer Science* (Vol. 4630). Berlin, Heidelberg: Springer.
- Finnsón, H., & Björnsson, Y. (2008). Simulation-based approach to general game playing. In *Proceedings of the 23rd national conference on artificial intelligence (AAAI)* (pp. 259–264). AAAI Press.
- Flaxman, A. D., Kalai, A. T., & McMahan, H. B. (2005). Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the sixteenth annual ACM-SIAM symposium on discrete algorithms* (pp. 385–394). Society for Industrial and Applied Mathematics.
- Fudenberg, D., Levine, D. K., et al. (1998). *The theory of learning in games* (Vol. 1). Cambridge: MIT Press Books.
- Gelly, S., & Silver, D. (2011). Monte-Carlo tree search and rapid action value estimation in computer Go. *Artificial Intelligence*, 175(11), 1856–1875.
- Gibson, R. G., Lanctot, M., Burch, N., Szafron, D., & Bowling, M. (2012). Generalized sampling and variance in counterfactual regret minimization. In *Proceedings of the twenty-sixth conference on artificial intelligence (AAAI-12)* (pp. 1355–1361).
- Hannan, J. (1957). Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, 3, 97–139.
- Hart, S., & Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5), 1127–1150.
- Hart, S., & Mas-Colell, A. (2001). A reinforcement procedure leading to correlated equilibrium. In G. Debreu, W. Neufeind, & W. Trockel (Eds.), *Economics essays* (pp. 181–200). Berlin: Springer. ISBN 978-3-642-07539-1. https://doi.org/10.1007/978-3-662-04623-4_12.
- Johanson, M., Waugh, K., Bowling, M., & Zinkevich, M. (2011). Accelerating best response calculation in large extensive games. In *Proceedings of the twenty-second international joint conference on artificial intelligence (IJCAI)* (pp. 258–265). AAAI Press.
- Kocsis, L., & Szepesvári, C. (2006). Bandit based Monte-Carlo planning. In J. Frnkranz, T. Scheffer, & M. Spiliopoulou (Eds.), *Machine learning: ECML 2006, volume 4212 of Lecture Notes in Computer Science* (pp. 282–293). Berlin: Springer. ISBN 978-3-540-45375-8.
- Lanctot, M., Waugh, K., Zinkevich, M., Bowling, M. (2009). Monte Carlo sampling for regret minimization in extensive games. In *Advances in neural information processing systems* (pp. 1078–1086).
- Lanctot, M., Lisý, V., & Winands, M. H. M. (2014). Monte Carlo tree search in simultaneous move games with applications to Goofspiel. In *Proceeding of the IJCAI computer games workshop, volume 408 of communications in computer and information science (CCIS)* (pp. 28–43). Springer.
- Lisý, V., Kovářik, V., Lanctot, M., & Božanský, B. (2013). Convergence of Monte Carlo tree search in simultaneous move games. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems* (Vol. 26, pp. 2112–2120). New York: Curran Associates Inc.

- Lisy, V., Lanctot, M., & Bowling, M. (2015). Online Monte Carlo counterfactual regret minimization for search in imperfect information games. In *Proceedings of the 2015 international conference on autonomous agents and multiagent systems* (pp. 27–36). International Foundation for Autonomous Agents and Multiagent Systems.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the eleventh international conference on machine learning*, vol. 157 (pp. 157–163).
- Moravčík, M., Schmid, M., Burch, N., Lisý, V., Morrill, D., Bard, N., et al. (2017). Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337), 508–513. <https://doi.org/10.1126/science.aam6960>.
- Neu, G., György, A., & Szepesvári, C. (2010). The online loop-free stochastic shortest-path problem. In *Conference on Learning Theory (COLT)* (pp. 231–243). Citeseer.
- Pandey, S., Chakrabarti, D., & Agarwal, D. (2007). Multi-armed bandit problems with dependent arms. In *Proceedings of the 24th international conference on machine learning (ICML)* (pp. 721–728). New York: ACM.
- Perick, P., St-Pierre, D. L., Maes, F., & Ernst, D. (2012). Comparison of different selection strategies in Monte-Carlo tree search for the game of Tron. In *Proceedings of the IEEE conference on computational intelligence and games (CIG)* (pp. 242–249).
- Puterman, M. L. (2014). *Markov decision processes: Discrete stochastic dynamic programming*. Hoboken: Wiley.
- Rhoads, G. C., & Bartholdi, L. (2012). Computer solution to the game of pure strategy. *Games*, 3(4), 150–156.
- Ross, S. M. (1971). Goofspiel: The game of pure strategy. *Journal of Applied Probability*, 8(3), 621–625.
- Saffidine, A., Finnsson, H., & Buro, M. (2012). Alpha-beta pruning for games with simultaneous moves. In *Proceedings of the twenty-sixth AAAI conference on artificial intelligence (AAAI)* (pp. 22–26).
- Shafiei, M., Sturtevant, N., & Schaeffer, J. (2009). Comparing UCT versus CFR in simultaneous games. In *Proceeding of the IJCAI workshop on general game-playing (GIGA)* (pp. 75–82).
- Stoltz, G. (2005). *Incomplete information and internal regret in prediction of individual sequences*. Ph.D. Thesis, Department of Mathematics, University Paris XI, Orsay.
- Tak, M. J. W., Winands, M. H. M., & Lanctot, M. (2014). Monte Carlo tree search variants for simultaneous move games. In *Proceedings of the IEEE conference on computational intelligence and games (CIG)* (pp. 232–239).
- Teytaud, O., & Flory, S. (2011). Upper confidence trees with short term partial information. In *Applications of evolutionary computation, volume 6624 of Lecture Notes in Computer Science* (pp. 153–162). Berlin: Springer.
- Waugh, K. (2009). *Abstraction in large extensive games*. Master's Thesis, University of Alberta.
- Yu, J. Y., & Mannor, S. (2009). Arbitrarily modulated Markov decision processes. In *Proceedings of the 48th IEEE Conference on decision and control, 2009 held jointly with the 2009 28th Chinese control conference. CDC/CCC 2009* (pp. 2946–2953). IEEE.
- Zinkevich, M., Johanson, M., Bowling, M., & Piccione, C. (2007). Regret minimization in games with incomplete information. In *Advances in neural information processing systems* (pp. 1729–1736).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.