

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/173598>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Online Markov Decision Processes with Non-oblivious Strategic Adversary

Le Cong Dinh¹, David Henry Mguni², Long Tran-Thanh³, Jun Wang⁴, Yaodong Yang⁵

¹University of Southampton, ²Huawei R&D UK, ³University of Warwick

⁴University College London, ⁵King's College London

Corresponding to: <yaodong.yang@kcl.ac.uk, lcd1u18@soton.ac.uk>

Abstract

We study a novel setting in Online Markov Decision Processes (OMDPs) where the loss function is chosen by a *non-oblivious* strategic adversary who follows a no-external regret algorithm. In this setting, we first demonstrate that MDP-Expert, an existing algorithm that works well with oblivious adversaries can still apply and achieve a policy regret bound of $O(\sqrt{T \log(L)} + \tau^2 \sqrt{T \log(|A|)})$ where L is the size of adversary's pure strategy set and $|A|$ denotes the size of agent's action space. Considering real-world games where the support size of a NE is small, we further propose a new algorithm: *MDP-Online Oracle Expert* (MDP-OOE), that achieves a policy regret bound of $O(\sqrt{T \log(L)} + \tau^2 \sqrt{T k \log(k)})$ where k depends only on the support size of the NE. MDP-OOE leverages the key benefit of Double Oracle in game theory and thus can solve games with prohibitively large action space. Finally, to better understand the learning dynamics of no-regret methods, under the same setting of no-external regret adversary in OMDPs, we introduce an algorithm that achieves last-round convergence result to a NE. To our best knowledge, this is first work leading to the last iteration result in OMDPs.¹

Introduction

Reinforcement Learning (RL) (Sutton and Barto 2018) provides a general solution framework for optimal decision making under uncertainty, where the agent aims to minimise its cumulative loss while interacting with the environment. While RL algorithms have shown empirical and theoretical successes in stationary environments, it is an open challenge to deal with non-stationary environments in which the loss function and/or the transition dynamics change over time (Laurent et al. 2011). In tackling non-stationary environments, we are interested in designing learning algorithms that can achieve no-regret guarantee (Even-Dar, Kakade, and Mansour 2009; Dick, Gyorgy, and Szepesvari 2014), where the regret is defined as the difference between the accumulated total loss and the total loss of the best fixed stationary policy in hindsight.

There are online learning algorithms that can achieve no-external regret property with changing loss func-

Preprint.

¹Accepted at Autonomous Agents and Multi-Agent Systems (2023): <https://doi.org/10.1007/s10458-023-09599-5>.

tion (but not changing transition dynamics), either in the full-information (Even-Dar, Kakade, and Mansour 2009; Dick, Gyorgy, and Szepesvari 2014) or the bandit (Neu et al. 2010; Neu and Olkhovskaya 2020) settings. However, most existing solutions are established based on the key assumption that the adversary is *oblivious*, meaning the changes in loss functions do not depend on the historical trajectories of the agent. This crucial assumption limits the applicability of no-regret algorithms to many RL fields, particularly multi-agent reinforcement learning (MARL) (Yang and Wang 2020). In a multi-agent system, since all agents are learning simultaneously, one agent's adaption on its strategy will make the environment *non-oblivious* from others agents' perspective. Therefore, to find the optimal strategy for each player, one must consider the strategic reactions from others rather than regarding them as purely oblivious. As such, studying no-regret algorithms against a non-oblivious adversary is a pivotal step in adapting existing online learning techniques into MARL settings.

Another challenge in online learning is the non-convergence dynamics in a system. When agents apply no-regret algorithms such as Multiplicative Weights Update (MWU) (Freund and Schapire 1999) or Follow the Regularized Leader (FTRL) (Shalev-Shwartz et al. 2011) to play against each other, the system demonstrates behaviours that are *Poincaré recurrent* (Mertikopoulos, Papadimitriou, and Piliouras 2018), meaning the last-round convergence can never be achieved (Bailey and Piliouras 2018). Recent works (Dinh et al. 2021a; Daskalakis and Panageas 2019) have focused on different learning dynamics in normal-form game that can lead to last-round convergence to a Nash equilibrium (NE) while maintaining the no-regret property. Yet, when it comes to OMDPs, it still remains an open challenge of how the no-regret property and the last-round convergence can be both achieved, especially with changing loss functions.

In this paper, we relax the assumption of the oblivious adversary in OMDPs and study a new setting where the loss function is chosen by a strategic agent that follows a no-external regret algorithm. This setting can be used in applications within economics to model systems and firms (Filar and Vrieze 1997), for example, an oligopoly with a dominant player, or ongoing interactions between industry players and an authority (e.g., a government that acts as an

order-setting body). Under this setting, we study how the agent can achieve different goals such as no-policy regret and last-round convergence.

Our contributions are at three folds:

- We prove that the well-known MDP-Expert (MDP-E) algorithm (Even-Dar, Kakade, and Mansour 2009) can still apply by achieving a policy regret bound of $O(\sqrt{T \log(L)} + \tau^2 \sqrt{T \log(|A|)})$, and the average strategies of the agents will converge to a NE of the game.
- For many real-world applications where the support size of NE is small (McMahan, Gordon, and Blum 2003; Dinh et al. 2021b), we introduce an efficient no-regret algorithm, *MDP-Online Oracle Expert (MDP-OOE)*, which achieves the policy regret bound of $O(\tau^2 \sqrt{T k \log(k)} + \sqrt{T \log(L)})$ against non-oblivious adversary, where k depends on the support size of the NE. MDP-OOE inherits the key benefits of both Double Oracle (McMahan, Gordon, and Blum 2003) and MDP-E (Even-Dar, Kakade, and Mansour 2009); it can solve games with large action space while maintaining the no-regret property.
- To achieve last-round convergence guarantee for no-external regret algorithms, we introduce the algorithm of *Last-Round Convergence in OMDPs (LRC-OMDPs)* such that in cases where the adversary follows a no-external regret algorithm, the dynamics will lead to the last-round convergence to a NE. To the best of our knowledge, this is first last-iteration convergence result in OMDPs.

Related Work

The setting of OMDPs with no-external regret adversary, though novel, shares certain aspects in common with existing literatures in online learning and stochastic games domains. Here we review each of them.

Many researchers have considered OMDPs with an oblivious environment, where the loss function can be set arbitrarily. The performance of the algorithm is measured by external regret: the difference between the total loss and the best stationary policy in hindsight. In this setting with stationary transition dynamics, MDP-E (Even-Dar, Kakade, and Mansour 2009) proved that if the agent bounds the “local” regret in each state, then the “global” regret will be bounded. Neu et al. (2010, 2013) considered the same problem with the bandit reward feedback and provided no-external regret algorithms in this setting. Dick, Gyorgy, and Szepesvari (2014) studied a new approach for OMDPs where the problem can be transformed into an online linear optimization form, from which no-external regret algorithms can be derived. Cheung, Simchi-Levi, and Zhu (2019) proposed a no-external regret algorithm in the case of non-stationary transition distribution, given that the variation of the loss and transition distributions do not exceed certain variation budgets.

In non-oblivious environment, Yu, Mannor, and Shimkin (2009) provided an example demonstrating that no algo-

rithms can guarantee sublinear external regret against a non-oblivious adversary. Thus, in OMDPs with non-oblivious opponents (e.g., agents using adaptive algorithms), the focus is often on finding stationary points of the system rather than finding a no-external regret algorithm (Leslie, Perkins, and Xu 2020). In this paper, we study cases where the adversary follows an adaptive no-regret algorithm, and tackle the hardness result of non-oblivious environments in OMDPs.

The problem of non-oblivious adversary has also been studied in the multi-armed bandit setting, a special case of OMDPs. In this setting, Arora, Dekel, and Tewari (2012) considered m -memory bounded adversary and provided an algorithm with policy regret bound that depends linearly on m , where the policy regret includes the adversary’s adaptive behaviour (i.e., see Equation (1)). Compared to their work, our paper considers strategic adversary which turns out to be ∞ -memory bounded adversary. Thus the algorithm suggested in Arora, Dekel, and Tewari (2012) can not be applied. Recently, Dinh et al. (2021a) studied the same strategic adversary in full information normal-form setting and provided an algorithm that leads to last round convergence. However, both of the above works only studied the simplified version of OMDPs, thus they do not capture the complexity of the problem. We argue that since strategic-adversary setting has many applications due to the popularity of no-regret algorithms (Cesa-Bianchi and Lugosi 2006; Zinkevich et al. 2007; Daskalakis et al. 2017), it is important to study no-regret methods in more practical settings such as OMDPs.

Stochastic games (SGs) (Shapley 1953; Deng et al. 2021) offer a multi-player game framework where agents jointly decide the loss and the state transition. Compared to OMDPs, the main difference is that SGs allow each player to have representation of states, actions and rewards, thus players can learn the representations over time and find the NE of the stochastic games (Wei, Hong, and Lu 2017; Tian et al. 2020). The performance in SGs is often measured by the difference between the average loss and the value of the game (i.e., the value when both players play a NE), which is a weaker notion of regret compared to the best fixed policy in hindsight in OMDPs. Intuitively, the player can learn the structure of the game (i.e., transition model, reward function) over time, thus on average, the player can calculate and compete with the value of the game. In non-episodic settings, the Upper Confidence Stochastic Game algorithm (UCSG) (Wei, Hong, and Lu 2017) guarantees the regret of $\text{Reg}_T = \tilde{O}(D^3 |S|^5 |A| + D |S| \sqrt{|A| T})$ with high probability, given that the opponent’s action is observable. However, to compete with the best stationary policy, knowing the game structure does not guarantee a good performance (i.e., the performance will heavily depend on the strategic behaviour of opponents). Tian et al. (2020) proved that in the SG setting, achieving no-regret with respect to best stationary policy in hindsight is statistically hard. Our settings can be considered as a sub-class of SGs where only the agent controls the transition model (i.e., single controller SGs), based on this, we try to overcome the above challenge.

We summarise the difference between our setting and

Table 1: The scope of our contribution in this work.

	Non-oblivious adversary within a two-player game framework	Oblivious adversary in Markov Decision Process
Regret <i>w.r.t</i> best policy in hindsight	MDP-OOE (our contribution) $O(\tau^2 \sqrt{Tk \log(k)} + \sqrt{T \log(L)})$	OMDPs: (MDP-E) (Even-Dar, Kakade, and Mansour 2009) $\text{Reg}_T = O(\tau^2 \sqrt{T \log(A)})$
Regret <i>w.r.t.</i> value of the game	SGs:(UCSG) (Wei, Hong, and Lu 2017) $\text{Reg}_T = \tilde{O}(D^3 S ^5 A + D S \sqrt{ A T})$	OMDPs

OMDPs and SGs in Table 1. Compared to OMDPs, we relax the assumption about oblivious environment and study a non-oblivious counterpart with strategic adversary. Compared to SGs, we relax the assumption of knowing opponent’s action in non-episodic setting and our results only require observing the loss functions. Furthermore, the performance measurement is with respect to the best stationary policy in hindsight, which is proved to be statistically hard in SGs (Tian et al. 2020). Intuitively, since we considers the problem of single controller SGs, it can overcome the hardness result. Guan et al. (2016) studied a similar setting to our paper, where only one player affects the transition kernel of the game. By viewing the game as an online linear optimisation, it can derive the minimax equilibrium of the game. There are two main challenges of the algorithm. Firstly, it requires both players to pre-calculate the minimax equilibrium of the game and fixes to this strategy during repeated game. Thus, in the situation where the adversary is an independent agent (i.e., it follows a different learning dynamic), the proposed algorithm can not be applied. Secondly, and most importantly, the no regret analysis is not provided for the algorithm in Guan et al. (2016), thus the algorithm can not be applied in an adversary environment. We fully address both challenges in this paper.

Problem Formulations & Preliminaries

We consider OMDPs where at each round $t \in \mathbb{N}$, an adversary can choose the loss function I_t based on the agent’s history $\{\pi_1, \pi_2, \dots, \pi_{t-1}\}$. Formally, we have OMDPs with finite state space S ; finite action set at each state A ; and a fixed transition model P . The agent’s starting state, x_1 , is distributed according to some distribution μ_0 over S . At time t , given state $x_t \in S$, the agent chooses an action $a_t \in A$, then the agent moves to a new random state x_{t+1} which is determined by the fixed transition model $P(x_{t+1}|x_t, a_t)$. Simultaneously, the agent receives an immediate loss $I_t(x_t, a_t)$, in which the loss function $I_t : S \times A \rightarrow R$ is bounded in $[0, 1]^{|A| \times |S|}$ and chosen by the adversary from a simplex $\Delta_L := \{I \in \mathbb{R}^{|S| \times |A|} | I = \sum_{i=1}^L x_i I_i, \sum_{i=1}^L x_i = 1, x_i \geq 0 \forall i\}$ where $\{I_1, I_2, \dots, I_L\}$ are the loss vectors of the adversary. We assume zero-sum game setting where the adversary receives the loss of $-I_t(x_t, a_t)$ at round t and consider popular full information feedback (Even-Dar, Kakade, and Mansour 2009; Dick, Gyorgy, and Szepesvari 2014), meaning the agent can observe the loss function I_t after each round t .

Against strategic adversary, the formal definition of no-external regret becomes inadequate since the adversary is allowed to adapt to the agent’s action. In this paper, we adopt the same approach in Arora, Dekel, and Tewari (2012) and consider policy regret. Formally, the goal of the agent is to have minimum policy regret with respect to the best fixed policy in hindsight:

$$R_T(\pi) = \mathbb{E}_{X,A} \left[\sum_{t=1}^T I_t^{\pi_t}(X_t, A_t) \right] - \mathbb{E}_{X,A} \left[\sum_{t=1}^T I_t^\pi(X_t^\pi, A_t^\pi) \right], \quad (1)$$

where $I_t^{\pi_t}$ denotes the loss function at time t while the agent follows π_1, \dots, π_T and I_t^π is the adaptive loss function against the fixed policy π of the agent. We say that the agent achieves sublinear policy regret (i.e., no-policy regret property) with respect to the best fixed strategy in hindsight if $R_T(\pi)$ satisfies: $\lim_{T \rightarrow \infty} \max_\pi \frac{R_T(\pi)}{T} = 0$.

In a general non-oblivious adversary, we prove by a counter example that it is impossible to achieve an algorithm with a sublinear policy regret². Suppose the agent faces an adversary such that it gives a very low loss for the agent if the action in the first round of the agent is a specific action (i.e., by fixing the loss function to $\mathbf{0}$), otherwise the adversary will give a high loss (i.e., by fixing the loss function to $\mathbf{1}$). Against this type of adversary, without knowing the specific action, the agent’s policy regret in Equation (1) will be $O(T)$. Thus, in general non-oblivious adversary case, we will have a hardness result in policy regret. To resolve the hardness result, we study strategic adversary in OMDPs.

Assumption 1 (Strategic Adversary). *The adversary flows a no-external regret algorithm such as for any sequence of π_t :*

$$\lim_{T \rightarrow \infty} \max_I \frac{R_T(I)}{T} = 0, \text{ where}$$

$$R_T(I) = \mathbb{E}_{X,A} \left[\sum_{t=1}^T I(X_t, A_t) \right] - \mathbb{E}_{X,A} \left[\sum_{t=1}^T I_t^{\pi_t}(X_t, A_t) \right].$$

The rationale of Assumption 1 comes from the vanilla property of no-external algorithms: without prior information, the adversary will not do worse than the best-fixed strat-

²In the multi-armed bandit setting, it is also impossible to achieve sublinear policy regret against all adaptive adversaries (see Theorem 1 in Arora, Dekel, and Tewari (2012)).

egy in hindsight (Dinh et al. 2021a). Thus, without the priority knowledge about the agent, the adversary will have incentive to follow a no-external regret algorithm. In the same way as full information feedback assumption for the agent, we assume that after each round t , the adversary observes the agent's stationary policy distribution \mathbf{d}_{π_t} .

For every policy π , we define $P(\pi)$ the state transition matrix induced by π such that $P(\pi)_{s,s'} = \sum_{a \in A} \pi(a|s) P_{s,s'}^a$. We assume through the paper that we have the mixing time assumption, which is a common assumption in OMDPs (Even-Dar, Kakade, and Mansour 2009; Dick, Gyorgy, and Szepesvari 2014; Neu et al. 2013):

Assumption 2 (Mixing time). *There exists a constant $\tau > 0$ such that for all distributions \mathbf{d} and \mathbf{d}' over the state space, any policy π ,*

$$\|\mathbf{d}P(\pi) - \mathbf{d}'P(\pi)\|_1 \leq e^{-1/\tau} \|\mathbf{d} - \mathbf{d}'\|_1,$$

where $\|\mathbf{x}\|$ denotes the l_1 norm of a vector \mathbf{x} .

Denote $\mathbf{v}_t^\pi(x, a)$ the probability of (state, action) pair (x, a) at time step t by following policy π with initial state x_1 . Following Assumption 2, for any initial states, \mathbf{v}_t^π will converge to a stationary distribution \mathbf{d}_π as t goes to infinity. Denote \mathbf{d}_Π the stationary distribution set from all agent's deterministic policies. With a slight abuse of notation, when an agent follows an algorithm A with use π_1, π_2, \dots at each time step, we denote $\mathbf{v}_t(x, a) = \mathbb{P}[X_t = x, A_t = a]$, $\mathbf{d}_t = \mathbf{d}_{\pi_t}$. Thus, the regret in Equation (1) can be expressed as

$$R_T(\pi) = \mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{v}_t \rangle \right] - \mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{l}_t^\pi, \mathbf{v}_t^\pi \rangle \right].$$

Assumption 2 allows us to define the average loss of policy π in an online MDP with loss \mathbf{l} as $\eta_\pi(\pi) = \langle \mathbf{l}, \mathbf{d}_\pi \rangle$ and the accumulated loss $Q_{\pi, \mathbf{l}}(s, a)$ is defined as

$$Q_{\pi, \mathbf{l}}(s, a) = E \left[\sum_{t=1}^{\infty} (\mathbf{l}(s_t, a_t) - \eta_\pi(\pi)) \middle| s_1 = s, a_1 = a, \pi \right].$$

As the dynamic between the agent and adversary is zero-sum, we can apply the minimax theorem (Neumann 1928):

$$\min_{\mathbf{d}_\pi \in \Delta_{\mathbf{d}_\Pi}} \max_{\mathbf{l} \in \Delta_{\mathbf{l}}} \langle \mathbf{l}, \mathbf{d}_\pi \rangle = \max_{\mathbf{l} \in \Delta_{\mathbf{l}}} \min_{\mathbf{d}_\pi \in \Delta_{\mathbf{d}_\Pi}} \langle \mathbf{l}, \mathbf{d}_\pi \rangle = v. \quad (2)$$

The saddle point $(\mathbf{l}, \mathbf{d}_\pi)$ that satisfies Equation (2) is the NE of the game (Nash et al. 1950) and v is called the value of the game. Our work is based on no-external regret algorithms in normal-form game such as Multiplicative Weights Update (Freund and Schapire 1999), which is described as

Definition 3 (Multiplicative Weights Update). *Let $\mathbf{k}_1, \mathbf{k}_2, \dots$ be a sequence of feedback received by the agent. The agent is said to follow the MWU if strategy $\tilde{\pi}_{t+1}$ is updated as follows*

$$\tilde{\pi}_{t+1}(i) = \tilde{\pi}_t(i) \frac{\exp(-\mu_t \mathbf{k}_t(\mathbf{a}^i))}{\sum_{i=1}^n \tilde{\pi}_t(i) \exp(-\mu_t \mathbf{k}_t(\mathbf{a}^i))}, \forall i \in [n], \quad (3)$$

where $\mu_t > 0$ is a parameter, n is the number of pure strategies (i.e., experts) and $\tilde{\pi}_0 = [1/n, \dots, 1/n]$.

We also consider ϵ -Nash equilibrium of the game:

Definition 4 (ϵ -Nash equilibrium). *Assume $\epsilon > 0$. We call a point $(\mathbf{l}, \mathbf{d}_\pi) \in \Delta_{\mathbf{l}} \times \Delta_{\mathbf{d}_\Pi}$ ϵ -NE if:*

$$\max_{\mathbf{l} \in \Delta_{\mathbf{l}}} \langle \mathbf{l}, \mathbf{d}_\pi \rangle - \epsilon \leq \langle \mathbf{l}, \mathbf{d}_\pi \rangle \leq \min_{\mathbf{d}_\pi \in \Delta_{\mathbf{d}_\Pi}} \langle \mathbf{l}, \mathbf{d}_\pi \rangle + \epsilon.$$

Under the setting of OMDPs against strategic adversary who aims to minimise the external regret (i.e., Assumption 1), we study several properties that the agent can achieve such as no-policy regret and last round convergence.

MDP-Expert against Strategic Adversary

Algorithm 1: MDP-Expert (MDP-E)

- 1: **Input:** Expert algorithm B_s (i.e., MWU) for each state
 - 2: **for** $t = 1$ to ∞ **do**
 - 3: Using algorithm B_s with set of expert A and the feedback $Q_{\pi_t, \mathbf{l}_t}(s, \cdot)$ for each state s
 - 4: Output π_{t+1} and observe \mathbf{l}_{t+1}
 - 5: **end for**
-

When the agent plays against a non-oblivious opponent, one challenge is that the best fixed policy π is not based on the current loss sequence $[l_1, l_2, \dots]$ of the agent but a different loss sequence $[l_1^\pi, l_2^\pi, \dots]$ induced by the policy π . Thus, to measure the regret in the case of a non-oblivious opponent, we need information on how the opponent will play against a fixed policy π . Under Assumption 1, we prove that existing MDP-E (Even-Dar, Kakade, and Mansour 2009) method, which is designed for oblivious adversary, will have no-policy regret property against the non-oblivious strategic adversary in our setting. Intuitively, MDP-E maintains a no-external regret algorithm (i.e., MWU) in each state to bound the local regret, thus the global regret can be bounded accordingly. The pseudocode of MDP-E is given in Algorithm 1.

The following lemma links the relationship between the external-regret of the adversary and the regret with respect to the policy stationary distribution:

Lemma 1. *Under MDP-E played by the agent, the external-regret of the adversary in Assumption 1 can be expressed as:*

$$\begin{aligned} R_T(\mathbf{l}) &= \mathbb{E}_{X, A} \left[\sum_{t=1}^T \mathbf{l}(X_t, A_t) \right] - \mathbb{E}_{X, A} \left[\sum_{t=1}^T \mathbf{l}_t^{\pi_t}(X_t, A_t) \right] \\ &= \sum_{t=1}^T \langle \mathbf{l}, \mathbf{d}_{\pi_t} \rangle - \sum_{t=1}^T \langle \mathbf{l}, \mathbf{d}_{\pi_t} \rangle + \mathcal{O}(\tau^2 \sqrt{T \log(|A|)}). \end{aligned}$$

(We provide the full proof in the Appendix A.1.)

Based on Lemma 1, we can tell that the sublinear regret will hold if and only if the adversary maintains a sublinear regret with respect to the agent's policy stationary distribution. As we assume that after each time t , the adversary can observe the stationary distribution \mathbf{d}_{π_t} , then by applying standard no-regret algorithm for online linear optimization against the feedback \mathbf{d}_{π_t} (i.e., MWU), the adversary can guarantee a good performance for himself. Thus, the Assumption 1 for the adversary is justifiable.

In the rest of the paper, with out loss of generality, we will study the case where the external-regret of the adversary with respect to agent’s policy stationary distribution has the following bound (i.e., the adversary follows optimal no-external regret algorithms such as MWU, FTRL with respect to policy stationary distribution of the agent ³):

$$\max_{l \in \Delta_L} \left(\sum_{t=1}^T \langle l, \mathbf{d}_{\pi_t} \rangle - \sum_{t=1}^T \langle l_t, \mathbf{d}_{\pi_t} \rangle \right) = \sqrt{\frac{T \log(L)}{2}}.$$

The next lemma provides a lower bound for performance of a fixed policy of the agent against a strategic adversary.

Lemma 2. *Suppose the agent follows a fixed stationary strategy π , then the adversary will converge to the best response to the fixed stationary strategy and*

$$\sum_{t=1}^T \langle l_t^\pi, \mathbf{d}_\pi \rangle \geq Tv - \sqrt{\frac{T \log(L)}{2}}.$$

The full proof is provided in the Appendix A.2. From Lemma 2, we can prove the following theorem:

Theorem 5. *Suppose the agent follows MDP-E Algorithm 1, then the regret with respect to the stationary distribution will be bounded by*

$$\sum_{t=1}^T \langle l_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \sum_{t=1}^T \langle l_t^\pi, \mathbf{d}_\pi \rangle \leq \sqrt{\frac{T \log(L)}{2}} + 3\tau \sqrt{\frac{T \log(|A|)}{2}}.$$

Proof. From Lemma 2, it is sufficient to show that

$$\sum_{t=1}^T \langle l_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle \leq Tv + 3\tau \sqrt{\frac{T \log(|A|)}{2}}.$$

Since the agent uses a no-regret algorithm with respect to the stationary distribution (e.g., MDP-E), following the same argument in Theorem 5.3 in (Even-Dar, Kakade, and Mansour 2009) we have:

$$\sum_{t=1}^T \langle l_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle \leq T \min_{\mathbf{d}_\pi} \langle \hat{l}, \mathbf{d}_\pi \rangle + 3\tau \sqrt{\frac{T \log(|A|)}{2}},$$

where $\hat{l} = \frac{1}{T} \sum_{t=1}^T l_t^{\pi_t}$. From the minimax equilibrium, we also have

$$\min_{\mathbf{d}_\pi} \langle \hat{l}, \mathbf{d}_\pi \rangle \leq \max_{l \in \Delta_L} \min_{\mathbf{d}_\pi \in \mathcal{d}_\Pi} \langle l, \mathbf{d}_\pi \rangle = v.$$

Thus, the proof is complete. \square

Now, we can make the link between the stationary regret and the regret of the agent in Equation (1).

Theorem 6. *Suppose the agent follows MDP-E Algorithm 1, then the agent’s regret in Equation (1) will be bounded by*

$$R_T(\pi) = O(\sqrt{T \log(L)} + \tau^2 \sqrt{T \log(|A|)}).$$

(We provide the full proof in the Appendix A.3.)

³If the adversary does not follow the optimal bound (i.e., irrational), then regret bound of the agent will change accordingly.

We note that Theorem 6 will hold true for a larger set of adversary outside Assumption 1 (e.g., FP (Brown 1951)) satisfying the following property: for every fixed policy of the agent, the adversary’s policy converges to the best response with respect to this fixed policy. With this property, we can bound the performance of agent’s fixed policy in Lemma 2 and thus derive the regret bound of the algorithm. Note that the regret bound in Theorem 6 will depend on the rate of convergence to best response against agent’s fixed policy.

As we have shown in previous theorems, the dynamic of playing no-regret algorithm in OMDPs against strategic adversary can be interpreted as a two-player zero-sum game setting with the corresponding stationary distribution. From the classical saddle point theorem (Freund and Schapire 1999), if both player follows a no-regret algorithm then the average strategies will converge to the saddle point (i.e., a NE).

Theorem 7. *Suppose the agent follows MDP-E, then the average strategies of both the agent and the adversary will converge to the ϵ_t -Nash equilibrium of the game with:*

$$\epsilon_t = \sqrt{\frac{\log(L)}{2T}} + 3\tau \sqrt{\frac{\log(|A|)}{2T}}$$

(We provide the full proof in the Appendix A.4.)

With the sublinear convergence rate to an NE, the dynamic between MDP-E and no-regret adversary (i.e., MWU) provides an efficient method to solve the single-controller SGs.

MDP-Online Oracle Expert Algorithm

As shown in the previous section, we can bound the regret in Equation (1) by bounding the regret with respect to the stationary distribution. In MDP-E, the regret bound (i.e., $O(\sqrt{T \log(L)} + \tau^2 \sqrt{T \log(|A|)})$) depends on the size of pure strategy set (i.e., $|A|$) thus it becomes less efficient when the agent has a prohibitively large pure strategy set.

Interestingly, recent paper by Dinh et al. (2021b) suggested that on normal-form games, it is possible to achieve a better regret bound where it only depends on the support size of NE rather than $|A|$. Unfortunately, extending this finding for OMDPs is highly non-trivial. The method in Dinh et al. (2021b) is designed for normal-form games only; in the worst scenario, its regret bound will depend on the size of pure strategy set, which is huge under our settings (i.e., $|A|^{|S|}$).

In this section, we provide a no-policy regret algorithm: MDP-Online Oracle Expert (MDP-OOE). It achieves the regret bound that only depends on the size of NE support rather than the size of the game. We start from presenting the small NE support size assumption.

Assumption 8 (Small Support Size of NE). *Let $(\mathbf{d}_{\pi^*}, \mathbf{l}^*)$ be a Nash equilibrium of the game of size $|A|^{|S|} \times L$. We assume the support size of $(\mathbf{d}_{\pi^*}, \mathbf{l}^*)$ is smaller than the game size: $\max(|\text{supp}(\mathbf{d}_{\pi^*})|, |\text{supp}(\mathbf{l}^*)|) < \min(|A|^{|S|}, L)$.*

Note that the assumption of small support size of NE holds in many real-world games (Czarnecki et al. 2020; Dinh et al. 2021b; Perez-Nieves et al. 2021; Liu et al. 2021;

Yang et al. 2021). In addition, we prove that such an assumption also holds in cases where the loss vectors $[l_1, \dots, l_L]$ is sampled from a continuous distribution and the size of the loss vector set L is small compared to the agent's pure strategy set, that is, $|A|^{|S|} \gg L$, thus further justifying the generality of this assumption.

Lemma 3. *Suppose that all loss functions are sampled from a continuous distribution and the size of the loss function set is small compared to the agent's pure strategy set (i.e., $|A|^{|S|} \gg L$). Let $(\mathbf{d}_{\pi^*}, \mathbf{l}^*)$ be a Nash equilibrium of the game of size $|A|^{|S|} \times L$. Then we have:*

$$\max(|\text{supp}(\mathbf{d}_{\pi^*})|, |\text{supp}(\mathbf{l}^*)|) \leq L.$$

(We provide the full proof in the Appendix B.1.)

Since the pure strategy set of the adversary L is much smaller compared to pure strategy set of the agent $|A|^{|S|}$, the support size of NE will highly likely be smaller compared to the size of agent's strategy set. Thus the agent can exploit this extra information to achieve better performance.

We now present the MDP-Online Oracle Expert (MDP-OOE) algorithm as follow. MDP-OOE maintains a set of effective strategy A_t^s in each state. In each iteration, the best response with respect to the average loss function will be calculated. If all the action in the best response are included in the current effective strategy set A_t^s for each state, then the algorithm continues with the current set A_t^s in each state. Otherwise, the algorithm updates the set of effective strategy in step 8 and 9 of Algorithm 2. We define the period of consecutive iterations as one *time window* T_i in which the set of effective strategy A_t^s stays fixed, i.e., $T_i := \{t \mid |A_t^s| = i\}$. Intuitively, since both the agent and the adversary use a no-regret algorithm to play, the average strategy of both players will converge to the NE of the game. Under the small NE support size assumption, the size of the agent's effective strategy set is also small compared to the whole pure strategy set (i.e., $|A|^{|S|}$). MDP-OOE ignores the pure strategies with poor average performance and only considers ones with high average performance. The regret bound with respect to the agent's stationary distribution is given as follow:

Theorem 9. *Suppose the learning agent uses Algorithm 2, then the regret with respect to the stationary distribution will be bounded by:*

$$\sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi^*}, \mathbf{d}_{\pi^*} \rangle \leq 3\tau \left(\sqrt{2Tk \log(k)} + \frac{k \log(k)}{8} \right),$$

where k is the number of time window.

(We provide the full proof in the Appendix B.2.)

In Algorithm 2, each time the agent updates the effective strategy set A_t^s at state s , exactly one new pure strategy is added into the effective strategy set for each state, thus the number k will be at most $|A|$. Therefore, we have the regret w.r.t the stationary distribution in the worst case will be:

$$3\tau \left(\sqrt{2T|A| \log(|A|)} + \frac{|A| \log(|A|)}{8} \right).$$

However, as shown in Dinh et al. (2021b, Figure 1), the number of iteration in DO method (respectively the number of

Algorithm 2: MDP-Online Oracle Expert

- 1: **Initialise:** Sets A_0^1, \dots, A_0^S of effective strategy set in each state
 - 2: **for** $t = 1$ to ∞ **do**
 - 3: $\pi_t = BR(\bar{\mathbf{l}})$
 - 4: **if** $\pi_t(s, \cdot) \in A_{t-1}^s$ for all s **then**
 - 5: $A_t^s = A_{t-1}^s$ for all s
 - 6: Using the expert algorithm B_s with effective strategy set A_t^s and the feedback $Q_{\pi_t, \mathbf{l}_t}(s, \cdot)$
 - 7: **else if** there exists $\pi_t(s, \cdot) \notin A_{t-1}^s$ **then**
 - 8: $A_t^s = A_{t-1}^s \cup \pi_t(s, \cdot)$ if $\pi_t(s, \cdot) \notin A_{t-1}^s$
 - 9: $A_t^s = A_{t-1}^s \cup a$ if $\pi_t(s, \cdot) \in A_{t-1}^s$ where a is randomly selected from the set A/A_{t-1}^s .
 - 10: Reset the expert algorithm B_s with effective strategy set A_t^s and the feedback $Q_{\pi_t, \mathbf{l}_t}(s, \cdot)$
 - 11: **end if**
 - 12: $\bar{\mathbf{l}} = \sum_{i=\bar{T}_i}^T \mathbf{l}_i$
 - 13: **end for**
-

time window in our setting) is linearly dependent in the support size of the NE, thus with Assumption 8, Algorithm 2 will be highly efficient.

Remark 1. *The regret bound in Theorem 9 will still hold in the case we consider the total average lost instead of average lost in each time window when calculating the best response in Algorithm 2.*

(We provide the full proof in the Appendix B.3.)

Given the regret with respect to policy's stationary distribution in Theorem 9, we can now derive the regret bound of Algorithm 2 with respect to the true performance:

Theorem 10. *Suppose the agent uses Algorithm 2 in our online MDPs setting, then the regret in Equation 1 can be bounded by:*

$$R_T(\pi) = O(\tau^2 \sqrt{Tk \log(k)} + \sqrt{T \log(L)}).$$

(We provide the full proof in the Appendix B.4.)

Notably, Algorithm 2 will not only reduce the regret bound in the case the number of strategy set k is small, it also reduces the computational hardness of computing expert algorithm when the number of experts is prohibitively large.

MDP-Online Oracle Algorithm with ϵ -best response.

In Algorithm 2, in each iteration the agent needs to calculate the exact best response to the average loss function $\bar{\mathbf{l}}$. Since calculating the exact best response is computationally hard and maybe infeasible in many situations (Vinyals et al. 2019), an alternative way is to consider ϵ -best response. That is, in each iteration in Algorithm 2, the agent can only access to a ϵ -best response to the average loss function, where ϵ is a predefined parameter. In this situation, we provide the regret analysis for Algorithm 2 as follow.

Theorem 11. *Suppose the agent only accesses to ϵ -best response in each iteration when following Algorithm 2. If the adversary follows a no-regret algorithm then the average*

strategy of the agent and the adversary will converge to ϵ -Nash equilibrium. Furthermore, the algorithm has ϵ -regret.

(We provide the full proof in the Appendix B.5.)

Theorem 11 implies that by following MDP-OOE, the agent can optimise the accuracy level (in terms of ϵ) based on the data that it receives to obtain the convergence rate and regret bound accordingly.

Last-Round Convergence to NE in OMDPs

In this section, we investigate OMDPs where the agent not only aims to minimize the regret but also stabilize the strategies. This is motivated by the fact that changing strategies through repeated games may be undesirable (e.g., see Dinh et al. (2021a); Daskalakis and Panageas (2019)). In on-line learning literature, minimizing regret and achieving the system's stability are often two conflict goals. That is, if all player in a system follows a no-regret algorithm (e.g., MWU, FTRL) to minimise the regret, then the dynamic of the system will become chaotic and the strategies of players will not converge in the last round (Dinh et al. 2021a; Mertikopoulos, Papadimitriou, and Piliouras 2018).

To achieve the goal, we start from studying the scenarios where the agent knows its NE of the game π^* . We then propose an algorithm: Last-Round Convergence in OMDPs (LRC-OMDP) that leads to last-round convergence to NE of the game in our setting. This is the first algorithm to our knowledge that achieves last-round convergence in OMDPs where only the learning agent knows the NE of the game. Notably, this goal is non-trivial to achieve. For example, if the agent keeps following the same strategy (i.e., the NE), then while the system might be stabilised (i.e., the adversary converges to the best response), yet this is still not a no-regret algorithm. Moreover, we notice that understanding the learning dynamics even when the NE is known is still challenging in the multi-agent learning domain. The AWESOME (Conitzer and Sandholm 2007) and CMLeS (Chakraborty and Stone 2014) algorithms make significant efforts to achieve convergence to NE under the assumption that each agent has access to a precomputed NE strategy. Compared to these algorithms, LRC-OMDP enjoys the key benefit that it does not require the adversary know its NE. Importantly, the adversary in our setting can be any types of strategic agent who observes the history and applies a no-regret algorithm to play, rather than being a restricted opponent such as a stationary opponent in AWESOME or a memory-bounded opponent in CMLeS.

The LRC-OMDP algorithm can be described as follow. At each odd round, the agent follows the NE strategy π^* so that in the next round, the strategy of the adversary will not deviating from the current strategy. Then, at the following even round, the agent chooses a strategy such that \mathbf{d}_{π_t} is a direction towards the NE strategy of the adversary. Depending on the distance between the current strategy of the adversary and its NE (which is measured by $v - \eta_{t-1}(\hat{\pi}_t)$), the agent will chooses a step size α_t such that the strategy of adversary will approach the NE. Note here that β is constant parameter and depends on the specific no-regret algorithm adversary follows, there is different optimal value for β . In case where

Algorithm 3: Last-Round Convergence in OMDPs

```

1: Input: Current iteration  $t$ 
2: Output: Strategy  $\pi_t$  for the agent
3: for  $t = 1, 2, \dots, T$  do
4:   if  $t = 2k - 1, k \in \mathbb{N}$  then
5:      $\pi_t = \pi^*$ 
6:   else if  $t = 2k, k \in \mathbb{N}$  then
7:      $\hat{\pi}_t(s) = \operatorname{argmin}_{a \in A} Q_{\pi^*, \mathbf{l}_t}(s, a) \forall s \in S$ 
8:      $\alpha_t = \frac{v - \eta_{t-1}(\hat{\pi}_t)}{\beta}$ ;  $\mathbf{d}_{\pi_t} = (1 - \alpha_t)\mathbf{d}_{\pi^*} + \alpha_t \mathbf{d}_{\hat{\pi}_t}$ 
9:     Output  $\pi_t$  via  $\mathbf{d}_{\pi_t}$ 
10:   end if
11: end for

```

the adversary follows the MWU algorithm, we can set $\beta = 1$.

We first introduce the condition in which the system achieves stability through the following lemma:

Lemma 4. *Let π^* be the NE strategy of the agent. Then, \mathbf{l} is the Nash Equilibrium of the adversary if the two following conditions hold:*

$$Q_{\pi^*, \mathbf{l}}(s, \pi^*) = \operatorname{argmin}_{\pi \in \Pi} Q_{\pi^*, \mathbf{l}}(s, \pi) \quad \forall s \in S \quad \text{and} \quad \eta_{\mathbf{l}}(\pi^*) = v.$$

The above lemma implies that if there is no improvement in the Q-value function for every state and the value of the current loss function equals to the value of the game, then there is last-round convergence to the NE. In situations where there is an improvement in one state, the following lemma bounds the value of a new strategy:

Lemma 5. *Assume that $\forall \pi \in \Pi, \mathbf{d}_{\pi}(s) > 0$. Then if there exist $s \in S$ such that*

$$Q_{\pi^*, \mathbf{l}_t}(s, \pi^*) > \operatorname{argmin}_{\pi \in \Pi} Q_{\pi^*, \mathbf{l}_t}(s, \pi),$$

then for $\pi_{t+1}(s) = \operatorname{argmin}_{a \in A} Q_{\pi^*, \mathbf{l}_t}(s, a) \forall s \in S$:

$$\eta_{\mathbf{l}_t}(\pi_{t+1}) < v.$$

Based on the above lemmas, we finally reach the last-round convergence of LRC-MDP in Algorithm 3.

Theorem 12. *Assume that the adversary follows the MWU algorithm with non-increasing step size μ_t such that $\lim_{T \rightarrow \infty} \sum_{t=1}^T \mu_t = \infty$ and there exists $t' \in \mathbb{N}$ with $\mu_{t'} \leq \frac{1}{3}$. If the agent follows Algorithm 3 then there exists a Nash equilibrium \mathbf{l}^* for the adversary such that $\lim_{t \rightarrow \infty} \mathbf{l}_t = \mathbf{l}^*$ almost everywhere and $\lim_{t \rightarrow \infty} \pi_t = \pi^*$.*

Proof. The full proof is given the the Appendix C.3. The main idea of the proof is to consider the relative entropy distance between the NE and the current strategy of the adversary and derives that $\forall k \in \mathbb{N} : 2k \geq t'$:

$$\operatorname{RE}(\mathbf{l}^* \| \mathbf{l}_{2k-1}) - \operatorname{RE}(\mathbf{l}^* \| \mathbf{l}_{2k+1}) \geq \frac{1}{2} \mu_{2k} \alpha_{2k} (v - \eta_{2k-1}(\hat{\pi}_{2k})).$$

□

The Algorithm 3 also applies in the situations where the adversary follows different learning dynamic such as Follow

the Regularized Leader or linear MWU (Dinh et al. 2021a). In these situations, Algorithm 3 requires to adapt the constant parameter β so that the convergence result still holds. Since both the agent and the adversary converges to a NE, the NE is also the best fixed strategy in hindsight. Consequently, LRC-OMDP is also a no-regret algorithm where the regret bound depends on the convergence rate to the NE.

Conclusion

In this paper, we have studied a novel setting in Online Markov Decision Processes where the loss function is chosen by a non-oblivious strategic adversary who follows a no-external regret algorithm. In this setting, we then revisited the MDP-E algorithm and provided a sublinear regret bound for it. We suggested a new algorithm of MDP-OOE that achieves the policy regret of $\mathcal{O}(\sqrt{T \log(L)} + \sqrt{\tau^2 T k \log(k)})$ where the regret does not depend the size of strategy set $|A|$ but the support size of the NE k . Finally, in tackling non convergence property of no-regret algorithms in self-plays, we provided the LRC-OMDP algorithm for the agent that leads to the first-known result of the last-round convergence to a NE.

References

- Arora, R.; Dekel, O.; and Tewari, A. 2012. Online bandit learning against an adaptive adversary: from regret to policy regret. *arXiv preprint arXiv:1206.6400*.
- Bailey, J. P.; and Piliouras, G. 2018. Multiplicative weights update in zero-sum games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 321–338.
- Bohnenblust, H.; Karlin, S.; and Shapley, L. 1950. Solutions of discrete, two-person games. *Contributions to the Theory of Games*, 1: 51–72.
- Brown, G. W. 1951. Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1): 374–376.
- Cesa-Bianchi, N.; and Lugosi, G. 2006. *Prediction, learning, and games*. Cambridge university press.
- Chakraborty, D.; and Stone, P. 2014. Multiagent learning in the presence of memory-bounded agents. *Autonomous agents and multi-agent systems*, 28(2): 182–213.
- Cheung, W. C.; Simchi-Levi, D.; and Zhu, R. 2019. Non-stationary reinforcement learning: The blessing of (more) optimism. *Available at SSRN 3397818*.
- Conitzer, V.; and Sandholm, T. 2007. AWESOME: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning*, 67(1-2): 23–43.
- Czarnecki, W. M.; Gidel, G.; Tracey, B.; Tuyls, K.; Omidshafiei, S.; Balduzzi, D.; and Jaderberg, M. 2020. Real World Games Look Like Spinning Tops. *arXiv preprint arXiv:2004.09468*.
- Daskalakis, C.; Ilyas, A.; Syrgkanis, V.; and Zeng, H. 2017. Training gans with optimism. *arXiv preprint arXiv:1711.00141*.
- Daskalakis, C.; and Panageas, I. 2019. Last-Iterate Convergence: Zero-Sum Games and Constrained Min-Max Optimization. *10th Innovations in Theoretical Computer Science*.
- Deng, X.; Li, Y.; Mguni, D. H.; Wang, J.; and Yang, Y. 2021. On the Complexity of Computing Markov Perfect Equilibrium in General-Sum Stochastic Games. *arXiv preprint arXiv:2109.01795*.
- Dick, T.; Gyorgy, A.; and Szepesvari, C. 2014. Online learning in Markov decision processes with changing cost sequences. In *ICML*, 512–520.
- Dinh, L. C.; Nguyen, T.-D.; Zemhoho, A. B.; and Tran-Thanh, L. 2021a. Last Round Convergence and No-Dynamic Regret in Asymmetric Repeated Games. In *Algorithmic Learning Theory*, 553–577. PMLR.
- Dinh, L. C.; Yang, Y.; Tian, Z.; Nieves, N. P.; Slumbers, O.; Mguni, D. H.; and Wang, J. 2021b. Online Double Oracle. *arXiv preprint arXiv:2103.07780*.
- Even-Dar, E.; Kakade, S. M.; and Mansour, Y. 2009. Online Markov decision processes. *Mathematics of Operations Research*, 34(3): 726–736.
- Filar, J.; and Vrieze, K. 1997. Applications and Special Classes of Stochastic Games. In *Competitive Markov Decision Processes*, 301–341. Springer.
- Freund, Y.; and Schapire, R. E. 1999. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2): 79–103.
- Guan, P.; Raginsky, M.; Willett, R.; and Zois, D.-S. 2016. Regret minimization algorithms for single-controller zero-sum stochastic games. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, 7075–7080. IEEE.
- Laurent, G. J.; Maignon, L.; Fort-Piat, L.; et al. 2011. The world of independent learners is not Markovian. *International Journal of Knowledge-based and Intelligent Engineering Systems*, 15(1): 55–64.
- Leslie, D. S.; Perkins, S.; and Xu, Z. 2020. Best-response dynamics in zero-sum stochastic games. *Journal of Economic Theory*, 189: 105095.
- Liu, X.; Jia, H.; Wen, Y.; Yang, Y.; Hu, Y.; Chen, Y.; Fan, C.; and Hu, Z. 2021. Unifying Behavioral and Response Diversity for Open-ended Learning in Zero-sum Games. *arXiv preprint arXiv:2106.04958*.
- McMahan, H. B.; Gordon, G. J.; and Blum, A. 2003. Planning in the presence of cost functions controlled by an adversary. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, 536–543.
- Mertikopoulos, P.; Papadimitriou, C.; and Piliouras, G. 2018. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2703–2717. SIAM.
- Nash, J. F.; et al. 1950. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1): 48–49.
- Neu, G.; Antos, A.; György, A.; and Szepesvári, C. 2010. Online Markov decision processes under bandit feedback. In *NeurIPS*, 1804–1812.

- Neu, G.; György, A.; Szepesvári, C.; and Antos, A. 2013. Online Markov decision processes under bandit feedback. *IEEE Transactions on Automatic Control*, 59(3): 676–691.
- Neu, G.; and Olkhovskaya, J. 2020. Online learning in MDPs with linear function approximation and bandit feedback. *arXiv e-prints*, arXiv:2007.
- Neumann, J. v. 1928. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1): 295–320.
- Perez-Nieves, N.; Yang, Y.; Slumbers, O.; Mguni, D. H.; Wen, Y.; and Wang, J. 2021. Modelling Behavioural Diversity for Learning in Open-Ended Games. In *International Conference on Machine Learning*, 8514–8524. PMLR.
- Shalev-Shwartz, S.; et al. 2011. Online learning and online convex optimization. *Foundations and trends in Machine Learning*, 4(2): 107–194.
- Shapley, L. S. 1953. Stochastic games. *Proceedings of the national academy of sciences*, 39(10): 1095–1100.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Tian, Y.; Wang, Y.; Yu, T.; and Sra, S. 2020. Online Learning in Unknown Markov Games.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.
- Wei, C.-Y.; Hong, Y.-T.; and Lu, C.-J. 2017. Online reinforcement learning in stochastic games. *arXiv preprint arXiv:1712.00579*.
- Yang, Y.; Luo, J.; Wen, Y.; Slumbers, O.; Graves, D.; Bou Ammar, H.; Wang, J.; and Taylor, M. E. 2021. Diverse Auto-Curriculum is Critical for Successful Real-World Multiagent Learning Systems. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 51–56.
- Yang, Y.; and Wang, J. 2020. An Overview of Multi-Agent Reinforcement Learning from Game Theoretical Perspective. *arXiv preprint arXiv:2011.00583*.
- Yu, J. Y.; Mannor, S.; and Shimkin, N. 2009. Markov decision processes with arbitrary reward processes. *Mathematics of Operations Research*, 34(3): 737–757.
- Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2007. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20: 1729–1736.

Appendix

MDP-Expert against Strategic Adversary

First we provide the following lemmas and proposition:

Lemma (Lemma 3.3 in (Even-Dar, Kakade, and Mansour 2009)). *For all loss function l in $[0, 1]$ and policies π , $Q_{l,\pi}(s, a) \leq 3\tau$.*

Lemma (Lemma 1 from (Neu et al. 2013)). *Consider a uniformly ergodic OMDPs with mixing time τ with losses $l_t \in [0, 1]^d$. Then, for any $T > 1$ and policy π with stationary distribution \mathbf{d}_π , it holds that*

$$\sum_{t=1}^T |\langle l_t, \mathbf{d}_\pi - \mathbf{v}_t^\pi \rangle| \leq 2\tau + 2.$$

This lemma guarantees that the performance of a policy's stationary distribution is similar to the actual performance of the policy in the case of a fixed policy.

In the other case of non-fixed policy, the following lemma bound the performance of policy's stationary distribution of algorithm A with the actual performance:

Lemma (Lemma 5.2 in (Even-Dar, Kakade, and Mansour 2009)). *Let π_1, π_2, \dots be the policies played by MDP-E algorithm \mathcal{A} and let $\tilde{\mathbf{d}}_{\mathcal{A},t}, \tilde{\mathbf{d}}_{\pi_t} \in [0, 1]^{|S|}$ be the stationary state distribution. Then,*

$$\|\tilde{\mathbf{d}}_{\mathcal{A},t} - \tilde{\mathbf{d}}_{\pi_t}\|_1 \leq 2\tau^2 \sqrt{\frac{\log(|A|)}{t}} + 2e^{-t/\tau}.$$

From the above lemma, since the policy's stationary distribution is a combination of stationary state distribution and the policy's action in each state, it is easy to show that:

$$\|\mathbf{v}_t - \mathbf{d}_{\pi_t}\|_1 \leq \|\tilde{\mathbf{d}}_{\mathcal{A},t} - \tilde{\mathbf{d}}_{\pi_t}\|_1 \leq 2\tau^2 \sqrt{\frac{\log(|A|)}{t}} + 2e^{-t/\tau}.$$

Proposition 13. *For the MWU algorithm (Freund and Schapire 1999) with appropriate μ_t , we have:*

$$R_T(\pi) = \mathbb{E} \left[\sum_{t=1}^T l_t(\pi_t) \right] - \mathbb{E} \left[\sum_{t=1}^T l_t(\pi) \right] \leq M \sqrt{\frac{T \log(n)}{2}},$$

where $\|l_t(\cdot)\| \leq M$. Furthermore, the strategy π_t does not change quickly: $\|\pi_t - \pi_{t+1}\| \leq \sqrt{\frac{\log(n)}{t}}$.

Proof. For a fixed T , if the loss function satisfies $l_t(\cdot) \leq 1$ then by setting $\mu_t = \sqrt{\frac{8 \log(n)}{T}}$, following Theorem 2.2 in (Cesa-Bianchi and Lugosi 2006) we have:

$$R_T(\pi) = \mathbb{E} \left[\sum_{t=1}^T l_t(\pi_t) \right] - \mathbb{E} \left[\sum_{t=1}^T l_t(\pi) \right] \leq 1 \sqrt{\frac{T \log(n)}{2}}. \quad (4)$$

Thus, in the case where $l_t(\cdot) \leq M$, by scaling up both sides by M in Equation (4) we have the first result of the Proposition 13. For the second part, follow the updating rule of MWU we have:

$$\begin{aligned} \pi_{t+1}(i) - \pi_t(i) &= \pi_t(i) \left(\frac{\exp(-\mu_t l_t(\mathbf{a}^i))}{\sum_{i=1}^n \pi_t(i) \exp(-\mu_t l_t(\mathbf{a}^i))} - 1 \right) \\ &\approx \pi_t(i) \left(\frac{1 - \mu_t l_t(\mathbf{a}^i)}{1 - \mu_t l_t(\pi_t)} - 1 \right) \\ &= \mu_t \pi_t(i) \frac{l_t(\pi_t) - l_t(\mathbf{a}^i)}{1 - \mu_t l_t(\pi_t)} = \mathcal{O}(\mu_t), \end{aligned} \quad (5a)$$

where we use the approximation $e^x \approx 1 + x$ for small x in Equation (5a). Thus, the difference in two consecutive strategies π_t will be proportional to the learning rate μ_t , which is set to be $\mathcal{O}(\sqrt{\frac{\log(n)}{t}})$. Similar result can be found in Proposition 1 in (Even-Dar, Kakade, and Mansour 2009). \square

Now, we are ready to prove the lemmas and theorems in the paper:

Proof of Lemma 1

Lemma. Under MDP-E played by the agent, the external-regret of the adversary in Assumption 1 can be expressed as:

$$\begin{aligned} R_T(\mathbf{l}) &= \mathbb{E}_{X,A} \left[\sum_{t=1}^T \mathbf{l}(X_t, A_t) \right] - \mathbb{E}_{X,A} \left[\sum_{t=1}^T \mathbf{l}_t^{\pi_t}(X_t, A_t) \right] \\ &= \sum_{t=1}^T \langle \mathbf{l}, \mathbf{d}_{\pi_t} \rangle - \sum_{t=1}^T \langle \mathbf{l}_t, \mathbf{d}_{\pi_t} \rangle + O(\tau^2 \sqrt{T \log(|A|)}). \end{aligned}$$

Proof. It is sufficient to show that for any sequence of \mathbf{l}_t

$$\mathbb{E}_{X,A} \left[\sum_{t=1}^T \mathbf{l}_t(X_t, A_t) \right] - \sum_{t=1}^T \langle \mathbf{l}_t, \mathbf{d}_{\pi_t} \rangle = O(\tau^2 \sqrt{T \log(|A|)}),$$

where \mathbf{l}_t denotes the loss vector of the adversary when the agent follows π_1, π_2, \dots (i.e., the same as $\mathbf{l}_t^{\pi_t}$).

Using the consequence of Lemma 5.2 in (Even-Dar, Kakade, and Mansour 2009), for any sequence of \mathbf{l}_t we have:

$$\begin{aligned} &\mathbb{E}_{X,A} \left[\sum_{t=1}^T \mathbf{l}_t(X_t, A_t) \right] - \sum_{t=1}^T \langle \mathbf{l}_t, \mathbf{d}_{\pi_t} \rangle \\ &= \sum_{t=1}^T \langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle \leq \sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| \leq \sum_{t=1}^T \|\mathbf{v}_t - \mathbf{d}_{\pi_t}\|_1 \\ &\leq \sum_{t=1}^T 2\tau^2 \sqrt{\frac{\log(|A|)}{t}} + 2e^{-t/\tau} \\ &\leq 4\tau^2 \sqrt{T \log(|A|)} + 2(1 + \tau) = O(\tau^2 \sqrt{T \log(|A|)}). \end{aligned} \tag{6}$$

The proof is complete. \square

Proof of Lemma 2

Lemma. Suppose the agent follows a fixed stationary strategy π , then the adversary will converge to the best response to the fixed stationary strategy and

$$\sum_{t=1}^T \langle \mathbf{l}_t^{\pi}, \mathbf{d}_{\pi} \rangle \geq Tv - \sqrt{\frac{T \log(L)}{2}}.$$

Proof. From Lemma 1, if the adversary follows a no-regret algorithm to achieve good performance in Assumption 1, then the adversary must follow a no-regret algorithm with respect to the policy's stationary distribution. Without loss of generality, we can assume that the adversary follows the Multiplicative Weight Update with respect to the policy's stationary distribution \mathbf{d}_{π} . Then follow the property of Multiplicative Weight Update in online linear problem, we have:

$$\max_{\mathbf{l} \in L} \langle \mathbf{l}, \mathbf{d}_{\pi} \rangle - \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi}, \mathbf{d}_{\pi} \rangle \leq \sqrt{\frac{\log(L)}{2T}}.$$

From the famous minimax theorem (Neumann 1928) we also have:

$$\max_{\mathbf{l} \in L} \langle \mathbf{l}, \mathbf{d}_{\pi} \rangle \geq \min_{\mathbf{d}_{\pi} \in \mathbf{d}_{\Pi}} \max_{\mathbf{l} \in L} \langle \mathbf{l}, \mathbf{d}_{\pi} \rangle = v.$$

Thus we have:

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi}, \mathbf{d}_{\pi} \rangle &\geq T \max_{\mathbf{l} \in L} \langle \mathbf{l}, \mathbf{d}_{\pi} \rangle - \sqrt{\frac{T \log(L)}{2}} \\ &\geq Tv - \sqrt{\frac{T \log(L)}{2}}. \end{aligned} \tag{7}$$

\square

Proof of Theorem 6

Theorem. Suppose the agent follows MDP-E Algorithm 1, then the agent's regret in Equation (1) will be bounded by

$$R_T(\pi) = O(\sqrt{T \log(L)} + \tau^2 \sqrt{T \log(|A|)}).$$

Proof. Using the consequence of Lemma 5.2 in (Even-Dar, Kakade, and Mansour 2009), for any sequence of \mathbf{l}_t we have:

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle &\leq \sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| \leq \sum_{t=1}^T \|\mathbf{v}_t - \mathbf{d}_{\pi_t}\|_1 \\ &\leq \sum_{t=1}^T 2\tau^2 \sqrt{\frac{\log(|A|)}{t}} + 2e^{-t/\tau} \\ &\leq 4\tau^2 \sqrt{T \log(|A|)} + 2(1 + \tau) = O(\tau^2 \sqrt{T \log(|A|)}). \end{aligned} \quad (8)$$

Thus we have

$$\sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| \leq 2(1 + \tau) + 4\tau^2 \sqrt{T \log(|A|)}. \quad (9)$$

Furthermore, if the agent uses a fixed policy π then by Lemma 2, we have:

$$\left| \sum_{t=1}^T \langle \mathbf{l}_t, \mathbf{d}_{\pi} - \mathbf{v}_t^{\pi} \rangle \right| \leq 2\tau + 2.$$

Since the agent uses MDP-E, a no-external regret algorithm, following the same argument in Theorem 4.1 in (Even-Dar, Kakade, and Mansour 2009) we have:

$$\sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle \leq T \min_{\mathbf{d}_{\pi}} \langle \hat{\mathbf{l}}, \mathbf{d}_{\pi} \rangle + 3\tau \sqrt{\frac{T \log(|A|)}{2}} \leq T\nu + 3\tau \sqrt{\frac{T \log(|A|)}{2}}.$$

Along with Lemma 2, we have:

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \sum_{t=1}^T \langle \mathbf{l}_t^{\pi}, \mathbf{d}_{\pi} \rangle &\leq \left(T\nu + 3\tau \sqrt{\frac{T \log(|A|)}{2}} \right) - \left(T\nu - \sqrt{\frac{T \log(L)}{2}} \right) \\ &= 3\tau \sqrt{\frac{T \log(|A|)}{2}} + \sqrt{\frac{T \log(L)}{2}}. \end{aligned}$$

Using the above two inequalities, we can bound the regret of the agent with respect to the regret of the policy's stationary distribution:

$$\begin{aligned} R_T(\pi) &= \mathbb{E}_{x,a} \left[\sum_{t=1}^T \mathbf{l}_t^{\pi_t}(x_t, a_t) \right] - \mathbb{E}_{x,a} \left[\sum_{t=1}^T \mathbf{l}_t^{\pi}(x_t^{\pi}, a_t^{\pi}) \right] \\ &= \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{v}_t \rangle - \sum_{t=1}^T \langle \mathbf{l}_t^{\pi}, \mathbf{v}_t^{\pi} \rangle \\ &\leq \sum_{t=1}^T (\langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle + |\langle \mathbf{l}_t^{\pi_t}, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle|) - \sum_{t=1}^T (\langle \mathbf{l}_t^{\pi}, \mathbf{d}_{\pi} \rangle - |\langle \mathbf{l}_t^{\pi}, \mathbf{v}_t^{\pi} - \mathbf{d}_{\pi} \rangle|) \\ &\leq \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \sum_{t=1}^T \langle \mathbf{l}_t^{\pi}, \mathbf{d}_{\pi} \rangle + 2(1 + \tau) + 4\tau^2 \sqrt{T \log(|A|)} + 2 + 2\tau \\ &\leq \sqrt{\frac{T \log(L)}{2}} + 3\tau \sqrt{\frac{T \log(|A|)}{2}} + 4(1 + \tau) + 4\tau^2 \sqrt{T \log(|A|)} \\ &= O(\sqrt{T \log(L)} + \tau^2 \sqrt{T \log(|A|)}). \end{aligned} \quad (10)$$

The proof is complete. \square

Proof of Theorem 7

Theorem. Suppose the agent follows MDP-E, then the average strategies of both the agent and the adversary will converge to the ϵ_t -Nash equilibrium of the game with the rate:

$$\epsilon_t = \sqrt{\frac{\log(L)}{2T}} + 3\tau\sqrt{\frac{\log(|A|)}{2T}}$$

Proof. Since the agent and the adversary use no-regret algorithms with respect to the policy's stationary distribution, we can use the property of regret bound in normal-form game to apply. Thus we have:

$$\begin{aligned} \max_{\mathbf{l} \in L} \langle \mathbf{l}, \hat{\mathbf{d}}_\pi \rangle - \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle &\leq \sqrt{\frac{\log(L)}{2T}}, \\ \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \min_{\mathbf{d}_\pi} \langle \hat{\mathbf{l}}, \mathbf{d}_\pi \rangle &\leq 3\tau\sqrt{\frac{\log(|A|)}{2T}}, \end{aligned}$$

where $\hat{\mathbf{d}}_\pi = \frac{1}{T} \sum_{t=1}^T \mathbf{d}_{\pi_t}$ and $\hat{\mathbf{l}} = \frac{1}{T} \sum_{t=1}^T \mathbf{l}_t^{\pi_t}$. From this, we can prove that

$$\begin{aligned} \langle \hat{\mathbf{l}}, \hat{\mathbf{d}}_\pi \rangle &\geq \min_{\mathbf{d}_\pi} \langle \hat{\mathbf{l}}, \mathbf{d}_\pi \rangle \geq \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - 3\tau\sqrt{\frac{\log(|A|)}{2T}} \\ &\geq \max_{\mathbf{l} \in L} \langle \mathbf{l}, \hat{\mathbf{d}}_\pi \rangle - \sqrt{\frac{\log(L)}{2T}} - 3\tau\sqrt{\frac{\log(|A|)}{2T}}, \end{aligned}$$

and,

$$\begin{aligned} \langle \hat{\mathbf{l}}, \hat{\mathbf{d}}_\pi \rangle &\leq \max_{\mathbf{l} \in L} \langle \mathbf{l}, \hat{\mathbf{d}}_\pi \rangle \leq \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle + \sqrt{\frac{\log(L)}{2T}} \\ &\leq \min_{\mathbf{d}_\pi} \langle \hat{\mathbf{l}}, \mathbf{d}_\pi \rangle + 3\tau\sqrt{\frac{\log(|A|)}{2T}} + \sqrt{\frac{\log(L)}{2T}}. \end{aligned}$$

Thus, with $\epsilon_t = \sqrt{\frac{\log(L)}{2T}} + 3\tau\sqrt{\frac{\log(|A|)}{2T}}$, we derive

$$\max_{\mathbf{l} \in L} \langle \mathbf{l}, \hat{\mathbf{d}}_\pi \rangle - \epsilon_t \leq \langle \hat{\mathbf{l}}, \hat{\mathbf{d}}_\pi \rangle \leq \min_{\mathbf{d}_\pi} \langle \hat{\mathbf{l}}, \mathbf{d}_\pi \rangle + \epsilon_t.$$

By definition, $(\hat{\mathbf{l}}, \hat{\mathbf{d}}_\pi)$ is ϵ_t -Nash equilibrium. □

MDP-Online Oracle Expert Algorithm

Proof of Lemma 3

Lemma. Suppose that the loss function is sampled from a continuous distribution and the size of the loss function set is small compared to the agent's pure strategy set (i.e., $|A|^{|S|} \gg L$). Let $(\mathbf{d}_{\pi^*}, \mathbf{l}^*)$ be a Nash equilibrium of the game of size $|A|^{|S|} \times L$. Then we have:

$$\max(|\text{supp}(\mathbf{d}_{\pi^*})|, |\text{supp}(\mathbf{l}^*)|) \leq L.$$

Proof. Within the set of all zero-sum games, the set of zero-sum games with non-unique equilibrium has Lebesgue measure zero (Bailey and Piliouras 2018). Thus, if the loss function's entries are sampled from a continuous distribution, then with probability one, the game has a unique NE. Following the Theorem 1 in (Bohnenblust, Karlin, and Shapley 1950) for game with unique NE, we have:

$$|\text{supp}(\mathbf{d}_{\pi^*})| = |\text{supp}(\mathbf{l}^*)|.$$

We also note that the support size of the NE can not exceed the size of the game:

$$|\text{supp}(\mathbf{d}_{\pi^*})| \leq |A|^{|S|}; \quad |\text{supp}(\mathbf{l}^*)| \leq L.$$

Thus we have:

$$\max(|\text{supp}(\mathbf{d}_{\pi^*})|, |\text{supp}(\mathbf{l}^*)|) = |\text{supp}(\mathbf{l}^*)| \leq L.$$

□

Proof of Theorem 9

Theorem. Suppose the agent uses the Algorithm 2, then the regret with respect to the stationary distribution will be bounded by:

$$\sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi^*}, \mathbf{d}_{\pi^*} \rangle \leq 3\tau \left(\sqrt{2Tk \log(k)} + \frac{k \log(k)}{8} \right),$$

where k is the number of time window.

Proof. We first have:

$$\begin{aligned} \mathbb{E}_{s \sim \mathbf{d}_{\pi}} [\mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, \pi)] &= \mathbb{E}_{s \sim \mathbf{d}_{\pi}, a \sim \pi} [\mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, a)] \\ &= \mathbb{E}_{s \sim \mathbf{d}_{\pi}, a \sim \pi} [\mathbf{l}_t(s, a) - \eta_{\mathbf{l}_t}(\pi_t) + \mathbb{E}_{s' \sim P_{s,a}} [\mathcal{Q}_{\pi_t, \mathbf{l}_t}(s', \pi_t)]] \\ &= \mathbb{E}_{s \sim \mathbf{d}_{\pi}, a \sim \pi} [\mathbf{l}_t(s, a)] - \eta_{\mathbf{l}_t}(\pi_t) + \mathbb{E}_{s \sim \mathbf{d}_{\pi}} [\mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, \pi_t)] \\ &= \eta_{\mathbf{l}_t}(\pi) - \eta_{\mathbf{l}_t}(\pi_t) + \mathbb{E}_{s \sim \mathbf{d}_{\pi}} [\mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, \pi_t)]. \end{aligned}$$

Thus we have:

$$\langle \mathbf{l}_t^{\pi^*}, \mathbf{d}_{\pi^*} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle = \sum_{s \in S} \mathbf{d}_{\pi^*}(s) (\mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, \pi^*) - \mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, \pi_t)). \quad (11)$$

Let T_1, T_2, \dots, T_k be the time window that the $\text{BR}(\bar{\mathbf{l}})$ does not change. Then in that time window, the best response to the current $\bar{\mathbf{l}}$ is inside the current pure strategies set in each state. In each time window, following Equation (11) we have:

$$\sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi^*} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle = \sum_{s \in S} \mathbf{d}_{\pi^*}(s) \sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} (\mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, \pi^*) - \mathcal{Q}_{\pi_t, \mathbf{l}_t}(s, \pi_t)). \quad (12)$$

Since during each time window, the pure strategies A_t^s does not change, thus we have:

$$\min_{\pi \in \Pi} \sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi^*} \rangle = \min_{\pi \in A_{|\bar{T}_i|}^s} \sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi^*} \rangle.$$

Thus, in each state s of a time window, the agent only needs to minimize the loss with respect to the action in $A_{|\bar{T}_i|}^s$. Put it differently, the expert algorithm in each state does not need to consider all pure action in each state, but just the current effective strategy set. For a time window T_i , if the agent uses a no-regret algorithm with the current effective action set and the learning rate $\mu_t = \sqrt{8 \log(i)}/t$, then the regret in each state will be bounded by (Cesa-Bianchi and Lugosi 2006):

$$3\tau \left(\sqrt{2|T_i| \log(A_t^s)} + \frac{\log(A_t^s)}{8} \right) \leq 3\tau \left(\sqrt{2|T_i| \log(i)} + \frac{\log(i)}{8} \right).$$

Thus, the regret in this time interval will also be bounded by:

$$\sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi^*} \rangle \leq 3\tau \left(\sqrt{2|T_i| \log(i)} + \frac{\log(i)}{8} \right). \quad (13)$$

Sum up from $i = 1$ to k in Inequality (13) we have:

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi^*} \rangle &= \sum_{i=1}^k \sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi^*} \rangle \\ &\leq \sum_{i=1}^k 3\tau \left(\sqrt{2|T_i| \log(i)} + \frac{\log(i)}{8} \right) \leq 3\tau \left(\sqrt{2Tk \log(k)} + \frac{k \log(k)}{8} \right). \end{aligned} \quad (14)$$

The proof is complete. \square

Proof of Remark 1

Remark. The regret bound in Theorem 9 will still hold in the case we consider the total average lost instead of average lost in each time window when calculating the best response in Algorithm 2.

Proof. We prove by induction that

$$\min_{\pi \in \Pi} \sum_{t=1}^{\bar{T}_k} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle \leq \sum_{j=1}^k \left[\sum_{t=\bar{T}_{j-1}+1}^{\bar{T}_j} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_j} \rangle \right],$$

where \mathbf{d}_{π_j} denotes the best response in the interval $[1, \bar{T}_j]$.

For $k = 1$, the claim is obvious. Suppose the claim is true k . We then have:

$$\begin{aligned} & \min_{\pi \in \Pi} \sum_{t=1}^{\bar{T}_{k+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle = \sum_{t=1}^{\bar{T}_{k+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_{k+1}} \rangle \\ & = \sum_{t=1}^{\bar{T}_k} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_{k+1}} \rangle + \sum_{t=\bar{T}_k+1}^{\bar{T}_{k+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_{k+1}} \rangle \\ & \leq \min_{\pi \in \Pi} \sum_{t=1}^{\bar{T}_k} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle + \sum_{t=\bar{T}_k+1}^{\bar{T}_{k+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_{k+1}} \rangle \\ & \leq \sum_{j=1}^k \left[\sum_{t=\bar{T}_{j-1}+1}^{\bar{T}_j} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_j} \rangle \right] + \sum_{t=\bar{T}_k+1}^{\bar{T}_{k+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_{k+1}} \rangle \quad (15a) \\ & = \sum_{j=1}^{k+1} \left[\sum_{t=\bar{T}_{j-1}+1}^{\bar{T}_j} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_j} \rangle \right], \end{aligned}$$

where the inequality (15a) dues to the induction assumption. Thus, for all k we have:

$$\min_{\pi \in \Pi} \sum_{t=1}^{\bar{T}_k} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle \leq \sum_{j=1}^k \left[\sum_{t=\bar{T}_{j-1}+1}^{\bar{T}_j} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_j} \rangle \right].$$

In other words, the Algorithm 2 will have the same regret bound when using the best response with respect to the total average strategy of the adversary. \square

Proof of Theorem 10

Theorem. Suppose the agent uses the Algorithm 2 in our online MDPs setting, then the regret in Equation (1) can be bounded by:

$$R_T(\pi) = O(\tau^2 \sqrt{Tk \log(k)} + \sqrt{T \log(L)}).$$

Proof. First we bound the difference between the true loss and the loss with respect to the policy's stationary distribution. Following the Algorithm 2, at the start of each time interval T_i (i.e., the time interval in which the effective strategy set does not change), the learning rate needs to restart to $O(\sqrt{\log(i)/t_i})$, where i denotes the number of pure strategies in the effective strategy set in the time interval T_i and t_i is relative position of the current round in that interval. Thus, following Lemma 5.2 in (Even-Dar, Kakade, and Mansour 2009), in each time interval T_i , the difference between the true loss and the loss with respect to the policy's stationary distribution will be:

$$\begin{aligned} & \sum_{t=t_{i-1}+1}^{t_i} |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| \leq \sum_{t=t_{i-1}+1}^{t_i} \|\mathbf{v}_t - \mathbf{d}_{\pi_t}\|_1 \\ & \leq \sum_{t=1}^{T_i} 2\tau^2 \sqrt{\frac{\log(i)}{t}} + 2e^{-t/\tau} \\ & \leq 4\tau^2 \sqrt{T_i \log(i)} + 2(1 + \tau). \end{aligned}$$

From this we have:

$$\begin{aligned}
\sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| &= \sum_{i=1}^k \sum_{t=\tau_{i-1}+1}^{\tau_i} |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| \\
&\leq \sum_{i=1}^k \left(4\tau^2 \sqrt{T_i \log(i)} + 2(1 + \tau) \right) \\
&\leq 4\tau^2 \sqrt{Tk \log(k)} + 2k(1 + \tau).
\end{aligned}$$

Following Lemma 1 from (Neu et al. 2013), we also have:

$$\sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{d}_\pi - \mathbf{v}_t^\pi \rangle| \leq 2\tau + 2.$$

Thus the regret in Equation (1) can be bounded by:

$$\begin{aligned}
R_T(\pi) &\leq \left(\sum_{t=1}^T \langle \mathbf{d}_{\pi_t}, \mathbf{l}_t \rangle + \sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| \right) - \left(\sum_{t=1}^T \langle \mathbf{l}_t^\pi, \mathbf{d}_\pi \rangle - \sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{d}_\pi - \mathbf{v}_t^\pi \rangle| \right) \\
&= \left(\sum_{t=1}^T \langle \mathbf{d}_{\pi_t}, \mathbf{l}_t \rangle - \sum_{t=1}^T \langle \mathbf{l}_t^\pi, \mathbf{d}_\pi \rangle \right) + \sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{v}_t - \mathbf{d}_{\pi_t} \rangle| + \sum_{t=1}^T |\langle \mathbf{l}_t, \mathbf{d}_\pi - \mathbf{v}_t^\pi \rangle| \\
&\leq 3\tau \left(\sqrt{2Tk \log(k)} + \frac{k \log(k)}{8} \right) + \frac{\sqrt{T \log(L)}}{\sqrt{2}} + 4\tau^2 \sqrt{Tk \log(k)} + 2k(1 + \tau) + 2\tau + 2 \\
&= O(\tau^2 \sqrt{Tk \log(k)} + \sqrt{T \log(L)}).
\end{aligned} \tag{16}$$

The proof is complete. \square

Proof of Theorem 11

Theorem. Suppose the agent only accesses to ϵ -best response in each iteration when following Algorithm 2. If the adversary follows a no-external regret algorithm then the average strategy of the agent and the adversary will converge to ϵ -Nash equilibrium. Furthermore, the algorithm has ϵ -regret.

Proof. Suppose that the player uses the Multiplicative Weights Update in Algorithm 2 with ϵ -best response. Let T_1, T_2, \dots, T_k be the time window that the players does not add up a new strategy. Since we have a finite set of strategies A then k is finite. Furthermore,

$$\sum_{i=1}^k T_k = T.$$

In a time window T_i , the regret with respect to the best strategy in the set of strategy at time T_i is:

$$\sum_{t=\bar{T}_i}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \min_{\pi \in A_{\bar{T}_i+1}} \sum_{t=\bar{T}_i}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_\pi \rangle \leq 3\tau \left(\sqrt{2T_i \log(i)} + \frac{\log(i)}{8} \right), \tag{17}$$

where $\bar{T}_i = \sum_{j=1}^{i-1} T_j$. Since in the time window T_i , the ϵ -best response strategy stays in $\Pi_{\bar{T}_i+1}$ and therefore we have:

$$\min_{\pi \in A_{\bar{T}_i+1}} \sum_{t=\bar{T}_i}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_\pi \rangle - \min_{\pi \in \Pi} \sum_{t=\bar{T}_i}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_\pi \rangle \leq \epsilon T_i.$$

Then, from the Equation (17) we have:

$$\sum_{t=\bar{T}_i}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \min_{\pi \in \Pi} \sum_{t=\bar{T}_i}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_\pi \rangle \leq 3\tau \left(\sqrt{2T_i \log(i)} + \frac{\log(i)}{8} \right) + \epsilon T_i. \tag{18}$$

Sum up the Equation (18) for $i = 1, \dots, k$ we have:

$$\begin{aligned} & \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \sum_{i=1}^k \min_{\pi \in \Pi} \sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle \leq \sum_{i=1}^k 3\tau \left(\sqrt{2T_i \log(i)} + \frac{\log(i)}{8} \right) + \epsilon T_i \\ \Rightarrow & \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \min_{\pi \in \Pi} \sum_{i=1}^k \sum_{t=|\bar{T}_i|}^{\bar{T}_{i+1}} \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle \leq \epsilon T + \sum_{i=1}^k 3\tau \left(\sqrt{2T_i \log(i)} + \frac{\log(i)}{8} \right) \end{aligned} \quad (19a)$$

$$\begin{aligned} \Rightarrow & \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \min_{\pi \in \Pi} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle \leq \epsilon T + \sum_{i=1}^k 3\tau \left(\sqrt{2T_i \log(i)} + \frac{\log(i)}{8} \right) \\ \Rightarrow & \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - \min_{\pi \in \Pi} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi} \rangle \leq \epsilon T + 3\tau \left(\sqrt{2Tk \log(k)} + \frac{k \log(k)}{8} \right). \end{aligned} \quad (19b)$$

Inequality (19a) is due to $\sum \min \leq \min \sum$. Inequality (19b) comes from Cauchy-Schwarz inequality and Stirling' approximation. Using inequality (19b), we have:

$$\min_{\pi \in \Pi} \langle \bar{\mathbf{l}}, \mathbf{d}_{\pi} \rangle \geq \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - 3\tau \left(\sqrt{\frac{2k \log(k)}{T}} + \frac{k \log(k)}{8T} \right) - \epsilon. \quad (20)$$

Since the adversary follows a no-regret algorithm, we have:

$$\begin{aligned} & \max_{\mathbf{l} \in \Delta_L} \sum_{t=1}^T \langle \mathbf{l}, \mathbf{d}_{\pi_t} \rangle - \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle \leq \sqrt{\frac{T}{2}} \sqrt{\log(L)} \\ \Rightarrow & \max_{\mathbf{l} \in \Delta_L} \sum_{t=1}^T \langle \mathbf{l}, \bar{\mathbf{d}}_{\pi} \rangle \leq \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle + \sqrt{\frac{\log(L)}{2T}}. \end{aligned} \quad (21)$$

Using the Inequalities (20) and (21) we have:

$$\begin{aligned} \langle \bar{\mathbf{l}}, \bar{\mathbf{d}}_{\pi} \rangle & \geq \min_{\pi \in \Pi} \langle \bar{\mathbf{l}}, \mathbf{d}_{\pi} \rangle \geq \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle - 3\tau \left(\sqrt{\frac{2k \log(k)}{T}} + \frac{k \log(k)}{8T} \right) - \epsilon \\ & \geq \max_{\mathbf{l} \in \Delta_L} \sum_{t=1}^T \langle \mathbf{l}, \bar{\mathbf{d}}_{\pi} \rangle - \sqrt{\frac{\log(L)}{2T}} - 3\tau \left(\sqrt{\frac{2k \log(k)}{T}} + \frac{k \log(k)}{8T} \right) - \epsilon. \end{aligned}$$

Similarly, we also have:

$$\begin{aligned} \langle \bar{\mathbf{l}}, \bar{\mathbf{d}}_{\pi} \rangle & \leq \max_{\mathbf{l} \in \Delta_L} \sum_{t=1}^T \langle \mathbf{l}, \bar{\mathbf{d}}_{\pi} \rangle \leq \frac{1}{T} \sum_{t=1}^T \langle \mathbf{l}_t^{\pi_t}, \mathbf{d}_{\pi_t} \rangle + \sqrt{\frac{\log(L)}{2T}} \\ & \leq \min_{\pi \in \Pi} \langle \bar{\mathbf{l}}, \mathbf{d}_{\pi} \rangle + 3\tau \left(\sqrt{\frac{2k \log(k)}{T}} + \frac{k \log(k)}{8T} \right) + \epsilon. \end{aligned}$$

Take the limit $T \rightarrow \infty$, we then have:

$$\max_{\mathbf{l} \in \Delta_L} \sum_{t=1}^T \langle \mathbf{l}, \bar{\mathbf{d}}_{\pi} \rangle - \epsilon \leq \langle \bar{\mathbf{l}}, \bar{\mathbf{d}}_{\pi} \rangle \leq \min_{\pi \in \Pi} \langle \bar{\mathbf{l}}, \mathbf{d}_{\pi} \rangle + \epsilon.$$

Thus $(\bar{\mathbf{l}}, \bar{\mathbf{d}}_{\pi})$ is the ϵ -Nash equilibrium of the game. \square

Last-Round Convergence to NE in OMDPs

Proof of Lemma 4

Lemma. Let π^* be the NE strategy of the agent. Then, \mathbf{l} is the Nash Equilibrium of the adversary if the two following conditions hold:

$$Q_{\pi^*, \mathbf{l}}(s, \pi^*) = \operatorname{argmin}_{\pi \in \Pi} Q_{\pi^*, \mathbf{l}}(s, \pi) \quad \forall s \in S \text{ and } \eta_{\mathbf{l}}(\pi^*) = v.$$

Proof. Using the definition of accumulated loss function Q we have

$$\begin{aligned}
\mathbb{E}_{s \in d_\pi} [Q_{\pi^*, \mathbf{l}}(s, \pi)] &= \mathbb{E}_{s \in d_\pi, a \in \pi} [Q_{\pi^*, \mathbf{l}}(s, a)] \\
&= \mathbb{E}_{s \in d_\pi, a \in \pi} [\mathbf{l}(s, a) - \eta_{\mathbf{l}}(\pi^*) + \mathbb{E}_{s' \sim P_{s,a}} [Q_{\pi^*, \mathbf{l}}(s', \pi^*)]] \\
&= \mathbb{E}_{s \in d_\pi, a \in \pi} [\mathbf{l}(s, a) - \eta_{\mathbf{l}}(\pi^*)] + \mathbb{E}_{s \in d_\pi} [Q_{\pi^*, \mathbf{l}}(s, \pi^*)] \\
&= \eta_{\mathbf{l}}(\pi) - \eta_{\mathbf{l}}(\pi^*) + \mathbb{E}_{s \in d_\pi} [Q_{\pi^*, \mathbf{l}}(s, \pi^*)].
\end{aligned} \tag{22}$$

Thus we have

$$\eta_{\mathbf{l}}(\pi) - \eta_{\mathbf{l}}(\pi^*) = \mathbb{E}_{s \in d_\pi} [Q_{\pi^*, \mathbf{l}}(s, \pi) - Q_{\pi^*, \mathbf{l}}(s, \pi^*)]. \tag{23}$$

Since we assume that

$$Q_{\pi^*, \mathbf{l}}(s, \pi^*) = \operatorname{argmin}_{\pi \in \Pi} Q_{\pi^*, \mathbf{l}}(s, \pi) \quad \forall s \in S,$$

we have

$$Q_{\pi^*, \mathbf{l}}(s, \pi) \geq Q_{\pi^*, \mathbf{l}}(s, \pi^*) \quad \forall s \in S, \pi \in \Pi. \tag{24}$$

It implies that

$$\mathbb{E}_{s \in d_\pi} [Q_{\pi^*, \mathbf{l}}(s, \pi) - Q_{\pi^*, \mathbf{l}}(s, \pi^*)] \geq 0 \quad \forall \pi \in \Pi. \tag{25}$$

Therefore we have

$$\eta_{\mathbf{l}}(\pi) \geq \eta_{\mathbf{l}}(\pi^*) \quad \forall \pi \in \Pi. \tag{26}$$

Along with the assumption $\eta_{\mathbf{l}}(\pi^*) = v$, we have the following relationship:

$$\operatorname{argmin}_{\pi \in \Pi} \eta_{\mathbf{l}}(\pi) = \eta_{\mathbf{l}}(\pi^*) = v. \tag{27}$$

Now we prove that for the loss function \mathbf{l} that satisfies Equation (27), then \mathbf{l} is NE for the adversary. Let (π^*, \mathbf{l}^*) be one of the NE of the game. Since the game we are considering is zero-sum game, (π^*, \mathbf{l}^*) satisfies the famous minimax theorem:

$$\min_{\pi \in \Pi} \max_{\mathbf{l}_1 \in L} \langle \mathbf{l}_1, \mathbf{d}_\pi \rangle = \max_{\mathbf{l}_1 \in L} \min_{\pi \in \Pi} \langle \mathbf{l}_1, \mathbf{d}_\pi \rangle = v \quad \text{where } \langle \mathbf{l}, \mathbf{d}_\pi \rangle = \eta_{\mathbf{l}}(\pi). \tag{28}$$

From Equation (27) we have

$$v = \min_{\pi \in \Pi} \langle \mathbf{l}, \mathbf{d}_\pi \rangle \leq \langle \mathbf{l}, \mathbf{d}_{\pi^*} \rangle. \tag{29}$$

Further, since \mathbf{l}^* is the NE of the game, then we have

$$v = \langle \mathbf{l}^*, \mathbf{d}_{\pi^*} \rangle = \max_{\mathbf{l}_1 \in L} \langle \mathbf{l}_1, \mathbf{d}_{\pi^*} \rangle \geq \langle \mathbf{l}, \mathbf{d}_{\pi^*} \rangle. \tag{30}$$

From inequalities (29) and (30) we have

$$v = \langle \mathbf{l}, \mathbf{d}_{\pi^*} \rangle = \min_{\pi \in \Pi} \langle \mathbf{l}, \mathbf{d}_\pi \rangle = \max_{\mathbf{l}_1 \in L} \langle \mathbf{l}_1, \mathbf{d}_{\pi^*} \rangle. \tag{31}$$

Thus, by definition (\mathbf{l}, π^*) is the Nash equilibrium of the game. In other words, the loss function \mathbf{l} satisfies the above assumption is the NE of the adversary. \square

Proof of Lemma 5

Lemma. Assume that $\forall \pi \in \Pi, \mathbf{d}_\pi(s) > 0$. Then if there exists $s \in S$ such that

$$Q_{\pi^*, \mathbf{l}_t}(s, \pi^*) > \operatorname{argmin}_{\pi \in \Pi} Q_{\pi^*, \mathbf{l}_t}(s, \pi),$$

then with a new strategy $\pi_{t+1}(s) = \operatorname{argmin}_{a \in A} Q_{\pi^*, \mathbf{l}_t}(s, a) \quad \forall s \in S$, we have

$$\eta_{\mathbf{l}_t}(\pi_{t+1}) < v.$$

Proof. From the minimax theorem, we have:

$$\eta_{\mathbf{l}_t}(\pi^*) \leq \eta_{\mathbf{l}^*}(\pi^*) = v \quad \forall \mathbf{l} \in L.$$

From the proof of Lemma 4 we have:

$$\eta_{\mathbf{l}_t}(\pi) - \eta_{\mathbf{l}_t}(\pi^*) = \mathbb{E}_{s \in d_\pi} [Q_{\pi^*, \mathbf{l}_t}(s, \pi) - Q_{\pi^*, \mathbf{l}_t}(s, \pi^*)] \quad \forall \pi \in \Pi.$$

Since the construction of the new strategy π_{t+1} we have:

$$\mathbb{E}_{s \in d_{\pi_{t+1}}} [Q_{\pi^*, \mathbf{l}_t}(s, \pi_{t+1}) - Q_{\pi^*, \mathbf{l}_t}(s, \pi^*)] < 0,$$

thus we have:

$$\eta_{\mathbf{l}_t}(\pi) < \eta_{\mathbf{l}_t}(\pi^*) \leq 0.$$

The proof is complete. \square

Proof of Theorem 12

Theorem. Assume that the adversary follows the MWU algorithm with non-increasing step size μ_t such that $\lim_{T \rightarrow \infty} \sum_{t=1}^T \mu_t = \infty$ and there exists $t' \in \mathbb{N}$ with $\mu_{t'} \leq \frac{1}{3}$. If the agent follows the Algorithm 3 then there exists a Nash equilibrium \mathbf{l}^* for the adversary such that $\lim_{t \rightarrow \infty} \mathbf{l}_t = \mathbf{l}^*$ almost everywhere and $\lim_{t \rightarrow \infty} \boldsymbol{\pi}_t = \boldsymbol{\pi}^*$.

In order to prove the above theorem, we first need the following lemma:

Lemma.

$$RE(\mathbf{l}^* || \mathbf{l}_{2k-1}) - RE(\mathbf{l}^* || \mathbf{l}_{2k+1}) \geq \frac{1}{2} \mu_{2k} \alpha_{2k} (v - \eta_{\mathbf{l}_{2k-1}}(\hat{\boldsymbol{\pi}}_{2k})) \quad \forall k \in \mathbb{N} : 2k \geq t'.$$

Proof. Using the definition of relative entropy we have:

$$\begin{aligned} & RE(\mathbf{l}^* || \mathbf{l}_{2k-1}) - RE(\mathbf{l}^* || \mathbf{l}_{2k+1}) \\ &= (RE(\mathbf{l}^* || \mathbf{l}_{2k+1}) - RE(\mathbf{l}^* || \mathbf{l}_{2k})) + (RE(\mathbf{l}^* || \mathbf{l}_{2k}) - RE(\mathbf{l}^* || \mathbf{l}_{2k-1})) \\ &= \left(\sum_{i=1}^n \mathbf{l}^*(i) \log \left(\frac{\mathbf{l}^*(i)}{\mathbf{l}_{2k+1}(i)} \right) - \sum_{i=1}^n \mathbf{l}^*(i) \log \left(\frac{\mathbf{l}^*(i)}{\mathbf{l}_{2k}(i)} \right) \right) + \\ & \quad \left(\sum_{i=1}^n \mathbf{l}^*(i) \log \left(\frac{\mathbf{l}^*(i)}{\mathbf{l}_{2k}(i)} \right) - \sum_{i=1}^n \mathbf{l}^*(i) \log \left(\frac{\mathbf{l}^*(i)}{\mathbf{l}_{2k-1}(i)} \right) \right) \\ &= \left(\sum_{i=1}^n \mathbf{l}^*(i) \log \left(\frac{\mathbf{l}_{2k}(i)}{\mathbf{l}_{2k+1}(i)} \right) \right) + \left(\sum_{i=1}^n \mathbf{l}^*(i) \log \left(\frac{\mathbf{l}_{2k-1}(i)}{\mathbf{l}_{2k}(i)} \right) \right). \end{aligned}$$

Following the update rule of the Multiplicative Weights Update algorithm we have:

$$\begin{aligned} & RE(\mathbf{l}^* || \mathbf{l}_{2k+1}) - RE(\mathbf{l}^* || \mathbf{l}_{2k-1}) \\ &= (-\mu_{2k} \langle \mathbf{l}^*, \mathbf{d}_{\pi_{2k}} \rangle + \log(Z_{2k})) + (-\mu_{2k-1} \langle \mathbf{l}^*, \mathbf{d}_{\pi_{2k}} \rangle + \log(Z_{2k-1})) \\ &\leq \left(-\mu_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k}(i) e^{\mu_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\pi_{2k}} \rangle} \right) \right) + (-\mu_{2k-1} v + \log(Z_{2k-1})) \\ &= \left(-\mu_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k-1} \langle \mathbf{e}_i, \mathbf{d}_{\pi_{2k-1}} \rangle} e^{\mu_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\pi_{2k}} \rangle} \right) - \log(Z_{2k-1}) \right) \\ & \quad + (-\mu_{2k-1} v + \log(Z_{2k-1})), \end{aligned} \tag{32a}$$

where Inequality (32a) is due to the fact that $\langle \mathbf{l}^*, \mathbf{d}_{\pi} \rangle \geq v \forall \pi$. Thus,

$$\begin{aligned} & RE(\mathbf{l}^* || \mathbf{l}_{2k+1}) - RE(\mathbf{l}^* || \mathbf{l}_{2k-1}) \\ &\leq \left(-\mu_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k-1} \langle \mathbf{e}_i, \mathbf{d}_{\pi_{2k-1}} \rangle} e^{\mu_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\pi_{2k}} \rangle} \right) \right) - \mu_{2k-1} v \\ &\leq \left(-\mu_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k-1} v} e^{\mu_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\pi_{2k}} \rangle} \right) \right) - \mu_{2k-1} v \\ &= -\mu_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\pi_{2k}} \rangle} \right), \end{aligned} \tag{33a}$$

where Inequality (33a) is the result of the inequality:

$$\langle \mathbf{l}, \mathbf{d}_{\pi^*} \rangle \leq v \quad \forall \mathbf{l}.$$

Now, using the update rule of Algorithm 3

$$\mathbf{d}_{\pi_{2k}} = (1 - \alpha_{2k}) \mathbf{d}_{\pi^*} + \alpha_{2k} \mathbf{d}_{\hat{\pi}_{2k}},$$

we have

$$\begin{aligned} & RE(\mathbf{l}^* || \mathbf{l}_{2k+1}) - RE(\mathbf{l}^* || \mathbf{l}_{2k-1}) \\ &\leq -\mu_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k} ((1-\alpha_{2k}) \langle \mathbf{e}_i, \mathbf{d}_{\pi^*} \rangle + \alpha_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\hat{\pi}_{2k}} \rangle)} \right) \\ &\leq -\mu_{2k} \alpha_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k} \alpha_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\hat{\pi}_{2k}} \rangle} \right). \end{aligned}$$

Denote $f(\mathbf{l}_{2k-1}) = \langle \mathbf{l}_{2k-1}, \mathbf{d}_{\hat{\pi}_{2k}} \rangle$, we then have

$$\begin{aligned} & \text{RE}(\mathbf{l}^* || \mathbf{l}_{2k+1}) - \text{RE}(\mathbf{l}^* || \mathbf{l}_{2k-1}) \\ & \leq -\mu_{2k} \alpha_{2k} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k} \alpha_{2k} \langle \mathbf{e}_i, \mathbf{d}_{\hat{\pi}_{2k}} \rangle} \right) \\ & = \mu_{2k} \alpha_{2k} (1 - v) + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{-\mu_{2k} \alpha_{2k} (1 - \langle \mathbf{e}_i, \mathbf{d}_{\hat{\pi}_{2k}} \rangle)} \right) \end{aligned} \quad (35a)$$

$$\leq \mu_{2k} \alpha_{2k} (1 - v) + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) (1 - (1 - e^{-\mu_{2k} \alpha_{2k}}) (1 - \langle \mathbf{e}_i, \mathbf{d}_{\hat{\pi}_{2k}} \rangle)) \right) \quad (35b)$$

$$\begin{aligned} & = \mu_{2k} \alpha_{2k} (1 - v) + \log (1 - (1 - e^{-\mu_{2k} \alpha_{2k}}) (1 - \langle \mathbf{l}_{2k-1}, \mathbf{d}_{\hat{\pi}_{2k}} \rangle)) \\ & \leq \mu_{2k} \alpha_{2k} (1 - v) - (1 - e^{-\mu_{2k} \alpha_{2k}}) (1 - \langle \mathbf{l}_{2k-1}, \mathbf{d}_{\hat{\pi}_{2k}} \rangle) \\ & = \mu_{2k} \alpha_{2k} (1 - v) - (1 - e^{-\mu_{2k} \alpha_{2k}}) (1 - f(\mathbf{l}_{2k-1})), \end{aligned} \quad (35c)$$

Equation (35a) is created by adding and subtracting $\mu_{2k} \alpha_{2k}$ on the first and second terms.

Inequalities (35b, 35c) are due to

$$\beta^x \leq 1 - (1 - \beta)x \quad \forall \beta \geq 0, \mathbf{l} \in [0, 1] \text{ and } \log(1 - x) \leq -x \quad \forall x < 1.$$

We can develop Inequality (35c) further as

$$\begin{aligned} & \text{RE}(\mathbf{l}^* || \mathbf{l}_{2k+1}) - \text{RE}(\mathbf{l}^* || \mathbf{l}_{2k-1}) \\ & \leq \mu_{2k} \alpha_{2k} (1 - v) - (1 - e^{-\mu_{2k} \alpha_{2k}}) (1 - f(\mathbf{l}_{2k-1})) \\ & \leq \mu_{2k} \alpha_{2k} (1 - v) - \left(1 - \left(1 - \mu_{2k} \alpha_{2k} + \frac{1}{2} (\mu_{2k} \alpha_{2k})^2 \right) \right) (1 - f(\mathbf{l}_{2k-1})) \end{aligned} \quad (36a)$$

$$\begin{aligned} & = \mu_{2k} \alpha_{2k} (f(\mathbf{l}_{2k-1}) - v) + \frac{1}{2} (\mu_{2k} \alpha_{2k})^2 (1 - f(\mathbf{l}_{2k-1})) \\ & \leq \mu_{2k} \alpha_{2k} (f(\mathbf{l}_{2k-1}) - v) + \frac{1}{2} \mu_{2k} \alpha_{2k} \mu_{2k} \frac{v - f(\mathbf{l}_{2k-1})}{\beta} (1 - f(\mathbf{l}_{2k-1})) \end{aligned} \quad (36b)$$

$$\begin{aligned} & \leq \mu_{2k} \alpha_{2k} (f(\mathbf{l}_{2k-1}) - v) + \frac{1}{2} \mu_{2k} \alpha_{2k} (v - f(\mathbf{l}_{2k-1})) \\ & = -\frac{1}{2} \mu_{2k} \alpha_{2k} (v - f(\mathbf{l}_{2k-1})) \leq 0. \end{aligned} \quad (36c)$$

Here, Inequality (36a) is due to $e^x \leq 1 + x + \frac{1}{2}x^2 \quad \forall x \in [-\infty, 0]$, Inequality (36b) comes from the definition of α_t :

$$\alpha_t = \frac{v - f(\mathbf{l}_{2k-1})}{\beta}, \quad \beta \geq 1 - f(\mathbf{l}), \quad f(\mathbf{l}_{2k-1}) \leq 1.$$

Finally, Inequality (36c) comes from the choice of k at the beginning of the proof, i.e., $\mu_{2k} \leq 1$. \square

Now we can prove Theorem 12:

Proof. We focus on the regret analysis with respect to the stationary distribution \mathbf{d}_{π_t} . Let \mathbf{l}^* be a minimax equilibrium strategy of the adversary (\mathbf{l}^* may not be unique). Following the above Lemma, for all $k \in \mathbb{N}$ such that $2k \geq t'$, we have

$$\text{RE}(\mathbf{l}^* || \mathbf{l}_{2k+1}) - \text{RE}(\mathbf{l}^* || \mathbf{l}_{2k-1}) \leq -\frac{1}{2} \mu_{2k} \alpha_{2k} (v - f(\mathbf{l}_{2k-1})), \quad (37)$$

where we denote $f(\mathbf{l}_{2k-1}) = \langle \mathbf{l}_{2k-1}, \mathbf{d}_{\hat{\pi}_{2k}} \rangle$. Thus, the sequence of relative entropy $\text{RE}(\mathbf{l}^* || \mathbf{l}_{2k-1})$ is non-increasing for all $k \geq \frac{t'}{2}$. As the sequence is bounded below by 0, it has a limit for any minimax equilibrium strategy \mathbf{l}^* . Since t' is a finite number and $\sum_{t=1}^{\infty} \mu_t = \infty$, we have $\sum_{t=t'}^{\infty} \mu_t = \infty$. Thus,

$$\lim_{T \rightarrow \infty} \sum_{k=\lceil \frac{t'}{2} \rceil}^T \mu_{2k} = \infty.$$

We will prove that $\forall \epsilon > 0, \exists h \in \mathbb{N}$ such that when the agent follows Algorithm 3 and the adversary follows MWU algorithm, the adversary will play strategy \mathbf{l}_h at round h and $v - f(\mathbf{l}_h) \leq \epsilon$. In particular, we prove this by contradiction. That is, suppose that $\exists \epsilon > 0$ such that $\forall h \in \mathbb{N}, v - f(\mathbf{l}_h) > \epsilon$. Then $\forall k \in \mathbb{N}$,

$$\alpha_{2k}(v - f(\mathbf{l}_{2k-1})) = \frac{(v - f(\mathbf{l}_{2k-1}))^2}{\beta} > \frac{\epsilon^2}{\beta}.$$

Let k vary from $\lceil \frac{t'}{2} \rceil$ to T in Equation (37). By summing over k , we obtain:

$$\begin{aligned} \text{RE}(\mathbf{l}^* \|\mathbf{l}_{2T+1}) &\leq \text{RE}(\mathbf{l}^* \|\mathbf{l}_{t'}) - \frac{1}{2} \sum_{k=\lceil \frac{t'}{2} \rceil}^T \mu_{2k} \alpha_{2k}(v - f(\mathbf{l}_{2k-1})) \\ &\leq \text{RE}(\mathbf{l}^* \|\mathbf{l}_{t'}) - \frac{1}{2} \frac{e^2}{\beta} \sum_{k=\lceil \frac{t'}{2} \rceil}^T \mu_{2k}. \end{aligned}$$

Since $\lim_{T \rightarrow \infty} \sum_{k=\lceil \frac{t'}{2} \rceil}^T \mu_{2k} = \infty$ and $\text{RE}(\mathbf{l}^* \|\mathbf{l}_{T+1}) \geq 0$, it contradicts our assumption about $\forall h \in \mathbb{N}, v - f(\mathbf{l}_h) > \epsilon$.

Now, we take a sequence of $\epsilon_k > 0$ such that $\lim_{k \rightarrow \infty} \epsilon_k = 0$. Then for each k , there exists $\mathbf{l}_{t_k} \in \Delta_n$ such that $v - \epsilon_k \leq f(\mathbf{l}_{t_k}) \leq v$. As Δ_n is a compact set and \mathbf{l}_{t_k} is bounded then following the Bolzano-Weierstrass theorem, there is a convergence subsequence $\mathbf{l}_{\bar{t}_k}$. The limit of that sequence, $\bar{\mathbf{l}}^*$, is a minimax equilibrium strategy of the row player (since $f(\bar{\mathbf{l}}^*) = f(\lim_{k \rightarrow \infty} \mathbf{l}_{\bar{t}_k}) = \lim_{k \rightarrow \infty} f(\mathbf{l}_{\bar{t}_k}) = v$). Combining with the fact that $\text{RE}(\bar{\mathbf{l}}^* \|\mathbf{l}_{2k-1})$ is non-increasing for $k \geq \lceil \frac{t'}{2} \rceil$ and $\text{RE}(\bar{\mathbf{l}}^* \|\bar{\mathbf{l}}^*) = 0$, we have $\lim_{k \rightarrow \infty} \text{RE}(\bar{\mathbf{l}}^* \|\mathbf{l}_{2k-1}) = 0$. We also note that

$$\begin{aligned} \text{RE}(\bar{\mathbf{l}}^* \|\mathbf{l}_{2k}) - \text{RE}(\bar{\mathbf{l}}^* \|\mathbf{l}_{2k-1}) &= -\mu_{2k-1} \langle \bar{\mathbf{l}}^*, \mathbf{d}_{\pi_{2k-1}} \rangle + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k-1} \langle \mathbf{e}_i, \mathbf{d}_{\pi^*} \rangle} \right) \\ &\leq -\mu_{2k-1} v + \log \left(\sum_{i=1}^n \mathbf{l}_{2k-1}(i) e^{\mu_{2k-1} v} \right) = 0, \end{aligned}$$

following the fact that $\langle \bar{\mathbf{l}}^*, \mathbf{d}_{\pi} \rangle \geq v$ for all $\pi \in \Pi$ and $\langle \mathbf{l}, \mathbf{d}_{\pi^*} \rangle \leq v$ for all \mathbf{l} . Thus, we have $\lim_{k \rightarrow \infty} \text{RE}(\bar{\mathbf{l}}^* \|\mathbf{l}_{2k}) = 0$ as well. Subsequently, $\lim_{t \rightarrow \infty} \text{RE}(\bar{\mathbf{l}}^* \|\mathbf{l}_t) = 0$, which concludes the proof. \square