
Policy Regret in Repeated Games

Raman Arora
Dept. of Computer Science
Johns Hopkins University
Baltimore, MD 21204
arora@cs.jhu.edu

Michael Dinitz
Dept. of Computer Science
Johns Hopkins University
Baltimore, MD 21204
mdinitz@cs.jhu.edu

Teodor V. Marinov
Dept. of Computer Science
Johns Hopkins University
Baltimore, MD 21204
tmarino2@jhu.edu

Mehryar Mohri
Courant Institute and Google Research
New York, NY 10012
mohri@cims.nyu.edu

Abstract

The notion of *policy regret* in online learning is a well defined performance measure for the common scenario of adaptive adversaries, which more traditional quantities such as external regret do not take into account. We revisit the notion of policy regret and first show that there are online learning settings in which policy regret and external regret are incompatible: any sequence of play that achieves a favorable regret with respect to one definition must do poorly with respect to the other. We then focus on the game-theoretic setting where the adversary is a self-interested agent. In that setting, we show that external regret and policy regret are not in conflict and, in fact, that a wide class of algorithms can ensure a favorable regret with respect to both definitions, so long as the adversary is also using such an algorithm. We also show that the sequence of play of no-policy regret algorithms converges to a *policy equilibrium*, a new notion of equilibrium that we introduce. Relating this back to external regret, we show that coarse correlated equilibria, which no-external regret players converge to, are a strict subset of policy equilibria. Thus, in game-theoretic settings, every sequence of play with no external regret also admits no policy regret, but the converse does not hold.

1 Introduction

Learning in dynamically evolving environments can be described as a repeated game between a player, an online learning algorithm, and an adversary. At each round of the game, the player selects an action, e.g. invests in a specific stock, the adversary, which may be the stock market, chooses a utility function, and the player gains the utility value of its action. The player observes the utility value and uses it to update its strategy for subsequent rounds. The player's goal is to accumulate the largest possible utility over a finite number of rounds of play.¹

The standard measure of the performance of a player is its *regret*, that is the difference between the utility achieved by the best offline solution from some restricted class and the utility obtained by the online player, when utilities are revealed incrementally. Formally, we can model learning as the following problem. Consider an action set \mathcal{A} . The player selects an action a_t at round t , the adversary picks a utility function u_t , and the player gains the utility value $u_t(a_t)$. While in a full observation setting the player observes the entire utility function u_t , in a *bandit setting* the player only observes

¹Such games can be equivalently described in terms of minimizing losses rather than maximizing utilities. All our results can be equivalently expressed in terms of losses instead of utilities.

the utility value of its own action, $u_t(a_t)$. We use the shorthand $a_{1:t}$ to denote the player's sequence of actions (a_1, \dots, a_t) and denote by $\mathcal{U}_t = \{u: \mathcal{A}^t \rightarrow \mathbb{R}\}$ the family of utility functions u_t . The objective of the player is to maximize its expected cumulative utility over T rounds, i.e. maximize $\mathbb{E}[\sum_{t=1}^T f_t(a_{1:t})]$, where the expectation is over the player's (possible) internal randomization. Since this is clearly impossible to maximize without knowledge of the future, the algorithm instead seeks to achieve a performance comparable to that of the best fixed action in hindsight. Formally, *external regret* is defined as

$$R(T) = \mathbb{E} \left[\max_{a \in \mathcal{A}} \sum_{t=1}^T u_t(a_{1:t-1}, a) - \sum_{t=1}^T u_t(a_{1:t}) \right]. \quad (1)$$

A player is said to *admit no external regret* if the external regret is sublinear, that is $R(T) = o(T)$. In contrast to statistical learning, online learning algorithms do not need to make stochastic assumptions about data generation: strong regret bounds are possible even if the utility functions are adversarial.

There are two main adversarial settings in online learning: the *oblivious setting* where the adversary ignores the player's actions and where the utility functions can be thought of as determined before the game starts (for instance, in weather prediction); and the *adaptive setting* where the adversary can react to the player's actions, thus seeking to throw the player off track (e.g., competing with other agents in the stock market). More generally, we define an *m-memory bounded adversary* as one that at any time t selects a utility function based on the player's past m actions: $u_t(a'_1, \dots, a'_{t-m-1}, a_{t-m}, \dots, a_t) = u_t(a_1, \dots, a_{t-m-1}, a_{t-m}, \dots, a_t)$, for all a'_1, \dots, a'_{t-m-1} and all a_1, \dots, a_t . An oblivious adversary can therefore be equivalently viewed as a 0-memory bounded adversary and an adaptive adversary as an ∞ -memory bounded adversary. For an oblivious adversary, external regret in Equation 1 reduces to $R(T) = \mathbb{E}[\max_{a \in \mathcal{A}} \sum_{t=1}^T u_t(a) - u_t(a_t)]$, since the utility functions do not depend upon past actions. Thus, external regret is meaningful when the adversary is oblivious, but it does not admit any natural interpretation when the adversary is adaptive. The problem stems from the fact that in the definition of external regret, the benchmark is a function of the player's actions. Thus, if the adversary is adaptive, or even memory-bounded for some $m > 0$, then, external regret does not take into account how the adversary would react had the player selected some other action.

To resolve this critical issue, Arora et al. (2012b) introduced an alternative measure of performance called *policy regret* for which the benchmark does not depend on the player's actions. Policy regret is defined as follows

$$P(T) = \max_{a \in \mathcal{A}} \sum_{t=1}^T u_t(a, \dots, a) - \mathbb{E} \left[\sum_{t=1}^T u_t(a_{1:t}) \right]. \quad (2)$$

Arora et al. (2012b) further gave a reduction, using a mini-batch technique where the minibatch size is larger than the memory m of adversary, that turns any algorithm with a sublinear external regret against an oblivious adversary into an algorithm with a sublinear policy regret against an m -memory bounded adversary, albeit at the price of a somewhat worse regret bound, which is still sublinear.

In this paper, we revisit the problem of online learning against adaptive adversaries. Since Arora et al. (2012b) showed that there exists an adaptive adversary against which any online learning algorithm admits linear policy regret, even when the external regret may be sublinear, we ask if no policy regret implies no external regret. One could expect this to be the case since policy regret seems to be a *stronger* notion than external regret. However, our first main result (Theorem 3.2) shows that this in fact is *not* the case and that the two notions of regret are incompatible: there exist adversaries (or sequence of utilities) on which action sequences with sublinear external regret admit linear policy regret and action sequences with sublinear policy regret incur linear external regret.

We argue, however, that such sequences may not arise in practical settings, that is in settings where the adversary is a self-interested entity. In such settings, rather than considering a malicious opponent whose goal is to hurt the player by inflicting large regret, it seems more reasonable to consider an opponent whose goal is to maximize his own utility. In zero-sum games, maximizing one's utility comes at the expense of the other player's, but there is a subtle difference between an adversary who is seeking to maximize the player's regret and an adversary who is seeking to minimize the player's utility (or maximize his own utility). We show that in such strategic game settings there is indeed a strong relationship between external regret and policy regret. In particular, we show in Theorem 3.4 that a large class of *stable* online learning algorithms with sublinear external regret also benefit from sublinear policy regret.

Further, we consider a two-player game where each player is playing a no policy regret algorithm. It is known that no external regret play converges to a coarse correlated equilibrium (CCE) in such a game, but what happens when players are using no policy regret algorithms? We show in Theorem 4.8 that the average play in repeated games between no policy regret players converges to a *policy equilibrium*, a new notion of equilibrium that we introduce. Policy equilibria differ from more traditional notions of equilibria such as Nash or CCEs in a crucial way. Recall that a CCE is defined to be a recommended joint strategy for players in a game such that there is no incentive for any player to deviate unilaterally from the recommended strategy if other players do not deviate.

What happens if the other players react to one player’s deviation by deviating themselves? This type of reasoning is not captured by external regret, but is essentially what is captured by policy regret. Thus, our notion of policy equilibrium must take into account these counterfactuals, and so the definition is significantly more complex. But, by considering functions rather than just actions, we can define such equilibria and prove that they exactly characterize no policy regret play.

Finally, it becomes natural to determine the relationship between policy equilibria (which characterize no policy regret play) and CCEs (which characterize no external regret play). We show in Theorems 4.9 and 4.10 that the set of CCEs is a strict subset of policy equilibria. In other words, every CCE can be thought of as a policy regret equilibrium, but no policy regret play might not converge to a CCE.

2 Related work

The problem of minimizing policy regret in a fully adversarial setting was first studied by Merhav et al. (2002). Their work dealt specifically with the full observation setting and assumed that the utility (or loss) functions were m -memory bounded. They gave regret bounds in $O(T^{2/3})$. The follow-up work by Farias and Megiddo (2006) designed algorithms in a reactive bandit setting. However, their results were not in the form of regret bounds but rather introduced a new way to compare against acting according to a fixed expert strategy. Arora et al. (2012b) studied m -memory bounded adversaries both in the bandit and full information settings and provided extensions to more powerful competitor classes considered in swap regret and more general Φ -regret. Dekel et al. (2014) provided a lower bound in the bandit setting for switching cost adversaries, which also leads to a tight lower bound for policy regret in the order $\Omega(T^{2/3})$. Their results were later extended by Koren et al. (2017a) and Koren et al. (2017b). More recently, Heidari et al. (2016) considered the multi-armed bandit problem where each arm’s loss evolves with consecutive pulls. The process according to which the loss evolves was not assumed to be stochastic but it was not arbitrary either – in particular, the authors required either the losses to be concave, increasing and to satisfy a decreasing marginal returns property, or decreasing. The regret bounds given are in terms of the time required to distinguish the optimal arm from all others.

A large part of reinforcement learning is also aimed at studying sequential decision making problems. In particular, one can define a Markov Decision Process (MDP) by a set of states equipped with transition distributions, a set of actions and a set of reward or loss distributions associated with each state action pair. The transition and reward distributions are assumed unknown and the goal is to play according to a strategy that minimizes loss or maximizes reward. We refer the reader to (Sutton and Barto, 1998; Kakade et al., 2003; Szepesvári, 2010) for general results in RL. MDPs in the online setting with bandit feedback or arbitrary payoff processes have been studied by Even-Dar et al. (2009a); Yu et al. (2009); Neu et al. (2010) and Arora et al. (2012a).

The tight connection between no-regret algorithms and correlated equilibria was established and studied by Foster and Vohra (1997); Fudenberg and Levine (1999); Hart and Mas-Colell (2000); Blum and Mansour (2007). A general extension to games with compact, convex strategy sets was given by Stoltz and Lugosi (2007). No external regret dynamics were studied in the context of socially concave games (Even-Dar et al., 2009b). More recently, Hazan and Kale (2008) considered more general notions of regret and established an equivalence result between fixed-point computation, the existence of certain no-regret algorithms, and the convergence to the corresponding equilibria. In a follow-up work by Mohri and Yang (2014) and Mohri and Yang (2017), the authors considered a more powerful set of competitors and showed that the repeated play according to conditional swap-regret or transductive regret algorithms leads to a new set of equilibria.

3 Policy regret in reactive versus strategic environments

Often, distant actions in the past influence an adversary more than more recent ones. The definition of policy regret (2) models this influence decay by assuming that the adversary is m -memory bounded for some $m \in \mathbb{N}$. This assumption is somewhat stringent, however, since ideally we could model the current move of the adversary as a function of the entire past, even if actions taken further in the past have less significance. Thus, we extend the definition of Arora et al. (2012b) as follows.

Definition 3.1. The m -memory policy regret at time T of a sequence of actions $(a_t)_{t=1}^T$ with respect to a fixed action a in the action set \mathcal{A} and the sequence of utilities $(u_t)_{t=1}^T$, where $u_t : \mathcal{A}^t \rightarrow \mathbb{R}$ and $m \in \mathbb{N} \cup \{\infty\}$ is

$$P(T, a) = \sum_{t=1}^T u_t(a_1, \dots, a_{t-m}, a, \dots, a) - \sum_{t=1}^T u_t(a_1, \dots, a_t).$$

We say that the sequence $(a_t)_{t=1}^T$ has sublinear policy regret (or no policy regret) if $P(T, a) < o(T)$, for all actions $a \in \mathcal{A}$.

Let us emphasize that this definition is just an extension of the standard policy regret definition and that, when the utility functions are m -memory bounded, the two definitions exactly coincide.

While the motivation for policy regret suggests that this should be a stronger notion compared to external regret, we show that not only that these notions are incomparable in the general adversarial setting, but that they are also *incompatible* in a strong sense.

Theorem 3.2. *There exists a sequence of m -memory bounded utility functions $(u_t)_{t=1}^T$, where $u_t : \mathcal{A} \rightarrow \mathbb{R}$, such that for any constant $m \geq 2$ (independent of T), any action sequence with sublinear policy regret will have linear external regret and any action sequence with sublinear external regret will have linear policy regret.*

The proof of the above theorem constructs a sequence for which no reasonable play can attain sublinear external regret. In particular, the only way the learner can have sublinear external regret is if they choose to have very small utility. To achieve this, the utility functions chosen by the adversary are the following. At time t , if the player chose to play the same action as their past 2 actions then they get utility $\frac{1}{2}$. If the player's past two actions were equal but their current action is different, then they get utility 1, and if their past two actions differ then no matter what their current action is they receive utility 0. It is easy to see that the maximum utility play for this sequence (and the lowest 2-memory bounded policy regret strategy) is choosing the same action at every round. However, such an action sequence admits linear external regret. Moreover, every sublinear external regret strategy must then admit sublinear utility and thus linear policy regret.

As discussed in Section 1, in many realistic environments we can instead think of the adversary as a self-interested agent trying to maximize their own utility, rather than trying to maximize the regret of the player. This more strategic environment is better captured by the game theory setting, in particular a 2-player game where both players are trying to maximize their utility. Even though we have argued that external regret is not a good measure, our next result shows that minimizing policy regret in games can be done if both players choose their strategies according to certain no external regret algorithms. More generally, we adapt a classical notion of stability from the statistical machine learning setting and argue that if the players use no external regret algorithms that are *stable*, then the players will have no policy regret in expectation. To state the result formally we first need to introduce some notation.

Game definition: We consider a 2-player game \mathcal{G} , with players 1 and 2. The action set of player i is denoted by \mathcal{A}_i , which we think of as being embedded into $\mathbb{R}^{|\mathcal{A}_i|}$ in the obvious way where each action corresponds to a standard basis vector. The corresponding simplex is $\Delta_{\mathcal{A}_i}$. The action of player 1 at time t is a_t and of player 2 is b_t . The observed utility for player i at time t is $u_i(a_t, b_t)$ and this is a bi-linear form with corresponding matrix P_i . We assume that the utilities are bounded in $[0, 1]$.

Algorithm of the player: When discussing algorithms, we take the view of player 1. Specifically, at time t , player 1 plays according to an algorithm which can be described as $Alg_t : (\mathcal{A}_1 \times \mathcal{A}_2)^t \rightarrow \Delta_{\mathcal{A}_1}$. We distinguish between two settings: full information, in which the player observes the full utility

function at time t (i.e., $u_1(\cdot, b_t)$), and the bandit setting, in which the player only observes $u_1(a_t, b_t)$. In the full information setting, algorithms like multiplicative weight updates (MWU Arora et al. (2012c)) depend only on the past $t-1$ utility functions $(u_1(\cdot, b_\ell))_{\ell=1}^{t-1}$, and thus we can think of Alg_t as a function $f_t : \mathcal{A}_2^t \rightarrow \Delta \mathcal{A}_1$. In the bandit setting, though, the output at time t of the algorithm depends both on the previous $t-1$ actions $(a_\ell)_{\ell=1}^{t-1}$ and on the utility functions (i.e., the actions picked by the other player).

But even in the bandit setting, we would like to think of the player's algorithm as a function $f_t : \mathcal{A}_2^t \rightarrow \Delta \mathcal{A}_1$. We cannot quite do this, however we *can* think of the player's algorithm as a *distribution* over such functions. So how do we remove the dependence on \mathcal{A}_1^t ? Intuitively, if we fix the sequence of actions played by player 2, we want to take the expectation of Alg_t over possible choices of the t actions played by player 1. In order to do this more formally, consider the distribution μ over $\mathcal{A}_1^{t-1} \times \mathcal{A}_2^{t-1}$ generated by simulating the play of the players for t rounds. Then let $\mu_{b_{0:t}}$ be the distribution obtained by conditioning μ on the actions of player 2 being $b_{0:t}$. Now we let $f_t(b_{0:t-1})$ be the distribution obtained by sampling $a_{0:t-1}$ from $\mu_{b_{0:t-1}}$ and using $Alg(a_{0:t-1}, b_{0:t-1})$. When taking expectations over f_t , the expectation is taken with respect to the above distribution. We also refer to the output $p_t = f_t(b_{0:t-1})$ as the strategy of the player at time t .

Now that we can refer to algorithms simply as functions (or distributions over functions), we introduce the notion of a stable algorithm.

Definition 3.3. Let $f_t : \mathcal{A}_2^t \rightarrow \Delta \mathcal{A}_1$ be a sample from Alg_t (as described above), mapping the past t actions in \mathcal{A}_2 to a distribution over the action set \mathcal{A}_1 . Let the distribution returned at time t be $p_t^1 = f_t(b_1, \dots, b_t)$. We call this algorithm *on average* $(m, S(T))$ *stable* with respect to the norm $\|\cdot\|$, if for any $b'_{t-m+1}, \dots, b'_t \in \mathcal{A}_2$ such that $\tilde{p}_t^1 = f_t(b_1, \dots, b_{t-m}, b'_{t-m+1}, \dots, b'_t) \in \Delta \mathcal{A}_1$, it holds that $\mathbb{E}[\sum_{t=1}^T \|p_t^1 - \tilde{p}_t^1\|] \leq S(T)$, where the expectation is taken with respect to the randomization in the algorithm.

Even though this definition of stability is given with respect to the game setting, it is not hard to see that it can be extended to the general online learning setting, and in fact this definition is similar in spirit to the one given in Saha et al. (2012). It turns out that most natural no external regret algorithms are stable. In particular we show, in the supplementary, that both Exp3 Auer et al. (2002) and MWU are on average $(m, m\sqrt{T})$ stable with respect to ℓ_1 norm for any $m < o(\sqrt{T})$. It is now possible to show that if each of the players are facing stable no external regret algorithms, they will also have bounded policy regret (so the incompatibility from Theorem 3.2 cannot occur in this case).

Theorem 3.4. Let $(a_t)_{t=1}^T$ and $(b_t)_{t=1}^T$ be the action sequences of player 1 and 2 and suppose that they are coming from no external regret algorithms modeled by functions f_t and g_t , with regrets $R_1(T)$ and $R_2(T)$ respectively. Assume that the algorithms are on average $(m, S(T))$ stable with respect to the ℓ_2 norm. Then

$$\begin{aligned}\mathbb{E}[P(T, a)] &\leq \|P_1\|S(T) + R_1(T) \\ \mathbb{E}[P(T, b)] &\leq \|P_2\|S(T) + R_2(T),\end{aligned}$$

where $u_t(a_{1:t})$ in the definition of $P(T, a)$ equals $u_1(a_t, g_t(a_{0:t-1}))$ and similarly in the definition of $P(T, b)$, equals $u_2(b_t, f_t(b_{0:t-1}))$. The above holds for any fixed actions $b \in \mathcal{A}_2$ and $a \in \mathcal{A}_1$. Here the matrix norm $\|\cdot\|$ is the spectral norm.

4 Policy equilibrium

Recall that unlike external regret, policy regret captures how other players in a game might react if a player decides to deviate from their strategy. The story is similar when considering different notions of equilibria. In particular Nash equilibria, Correlated equilibria and CCEs can be interpreted in the following way: if player i deviates from the equilibrium play, their utility will not increase no matter how they decide to switch, provided that *all other players continue to play according to the equilibrium*. This sentiment is a reflection of what no external and no swap regret algorithms guarantee. Equipped with the knowledge that no policy regret sequences are obtainable in the game setting under reasonable play from all parties, it is natural to reason how other players would react if player i deviated and what would be the cost of deviation when taking into account possible reactions.

Let us again consider the 2-player game setup through the view of player 1. The player believes their opponent might be m -memory bounded and decides to proceed by playing according to a no policy

regret algorithm. After many rounds of the game, player 1 has computed an empirical distribution of play $\widehat{\sigma}$ over $\mathcal{A} := \mathcal{A}_1 \times \mathcal{A}_2$. The player is familiar with the guarantees of the algorithm and knows that, if instead, they changed to playing any fixed action $a \in \mathcal{A}_1$, then the resulting empirical distribution of play $\widehat{\sigma}_a$, where player 2 has responded accordingly in a memory-bounded way, is such that $\mathbb{E}_{(a,b) \sim \widehat{\sigma}} [u_1(a, b)] \geq \mathbb{E}_{(a,b) \sim \widehat{\sigma}_a} [u_1(a, b)] - \epsilon$. This thought experiment suggests that if no policy regret play converges to an equilibrium, then the equilibrium is not only described by the deviations of player 1, but also through the change in player 2's behavior, which is encoded in the distribution $\widehat{\sigma}_a$. Thus, any equilibrium induced by no policy regret play, can be described by tuples of distributions $\{(\sigma, \sigma_a, \sigma_b) : (a, b) \in \mathcal{A}\}$, where σ_a is the distribution corresponding to player 1's deviation to the fixed action $a \in \mathcal{A}_1$ and σ_b captures player 2's deviation to the fixed action $b \in \mathcal{A}_2$. Clearly σ_a and σ_b are not arbitrary but we still need a formal way to describe how they arise.

For convenience, let's restrict the memory of player 2 to be 1. Thus, what player 1 believes is that at each round t of the game, they play an action a_t and player 2 plays a function $f_t : \mathcal{A}_1 \rightarrow \mathcal{A}_2$, mapping a_{t-1} to $b_t = f_t(a_{t-1})$. Finally, the observed utility is $u_1(a_t, f_t(a_{t-1}))$. The empirical distribution of play, $\widehat{\sigma}$, from the perspective of player 1, is formed from the observed play $(a_t, f_t(a_{t-1}))_{t=1}^T$. Moreover, the distribution, $\widehat{\sigma}_a$, that would have occurred if player 1 chose to play action a on every round is formed from the play $(a, f_t(a))_{t=1}^T$. In the view of the world of player 1, the actions taken by player 2 are actually functions rather than actions in \mathcal{A}_2 . This suggests that the equilibrium induced by a no-policy regret play, is a distribution over the functional space defined below.

Definition 4.1. Let $\mathcal{F}_1 := \{f : \mathcal{A}_1^{m_1} \rightarrow \mathcal{A}_2\}$ and $\mathcal{F}_2 := \{g : \mathcal{A}_1^{m_2} \rightarrow \mathcal{A}_2\}$ denote the *functional spaces of play* of players 1 and 2, respectively. Denote the product space by $\mathcal{F} := \mathcal{F}_1 \times \mathcal{F}_2$.

Note that when $m_1 = m_2 = 0$, \mathcal{F} is in a one-to-one correspondence with \mathcal{A} , i.e. when players believe their opponents are oblivious, we recover the action set studied in standard equilibria. For simplicity, for the remainder of the paper we assume that $m_1 = m_2 = 1$. However, all of the definitions and results that follow can be extended to the fully general setting of arbitrary m_1 and m_2 ; see the supplementary for details.

Let us now investigate how a distribution π over \mathcal{F} can give rise to a tuple of distributions $(\widehat{\sigma}, \widehat{\sigma}_a, \widehat{\sigma}_b)$. We begin by defining the utility of π such that it equals the utility of a distribution σ over \mathcal{A} i.e., we want $\mathbb{E}_{(f,g) \sim \pi} [u_1(f, g)] = \mathbb{E}_{(a,b) \sim \sigma} [u_1(a, b)]$. Since utilities are not defined for functions, we need an interpretation of $\mathbb{E}_{(f,g) \sim \pi} [u_1(f, g)]$ which makes sense. We notice that π induces a Markov chain with state space \mathcal{A} in the following way.

Definition 4.2. Let π be any distribution over \mathcal{F} . Then π induces a *Markov process* with transition probabilities $\mathbb{P}[(a_2, b_2)|(a_1, b_1)] = \sum_{(f,g) \in \mathcal{F}_1 \times \mathcal{F}_2 : f(b_1)=a_2, g(a_1)=b_2} \pi(f, g)$. We associate with this Markov process the transition matrix $M \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{A}|}$, with $M_{x_1, x_2} = \mathbb{P}[x_2|x_1]$ where $x_i = (a_i, b_i)$.

Since every Markov chain with a finite state space has a stationary distribution, we think of utility of π as the utility of a particular stationary distribution σ of M . How we choose σ among all stationary distributions is going to become clear later, but for now we can think about σ as the distribution which maximizes the utilities of both players. Next, we need to construct σ_a and σ_b , which capture the deviation in play, when player 1 switches to action a and player 2 switches to action b . The no-policy regret guarantee can be interpreted as $\mathbb{E}_{(f,g) \sim \pi} [u_1(f, g)] \geq \mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a))]$ i.e., if player 1 chose to switch to a fixed action (or equivalently, the constant function which maps everything to the action $a \in \mathcal{A}_1$), then their utility should not increase. Switching to a fixed action a , changes π to a new distribution π_a over \mathcal{F} . This turns out to be a product distribution which also induces a Markov chain.

Definition 4.3. Let π be any distribution over \mathcal{F} . Let δ_a be the distribution over \mathcal{F}_1 putting all mass on the constant function mapping all actions $b \in \mathcal{A}_2$ to the fixed action $a \in \mathcal{A}_1$. Let $\pi_{\mathcal{F}_2}$ be the marginal of π over \mathcal{F}_2 . The *distribution resulting from player 1 switching to playing a fixed action $a \in \mathcal{A}_1$* , is denoted as $\pi_a = \delta_a \times \pi_{\mathcal{F}_2}$. This distribution induces a Markov chain with transition probabilities $\mathbb{P}[(a_2, b_2)|(a_1, b_1)] = \sum_{(f,g) : g(a_1)=b_2} \pi(f, g)$ and the *transition matrix of this Markov process is denoted by M_a* . The distribution π_b and matrix M_b are defined similarly for player 2.

Since the no policy regret algorithms we work with do not directly induce distributions over the functional space \mathcal{F} but rather only distributions over the action space \mathcal{A} , we would like to state all of our utility inequalities in terms of distributions over \mathcal{A} . Thus, we would like to check if there is a stationary distribution σ_a of M_a such that $\mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a))] = \mathbb{E}_{(a,b) \sim \sigma_a} [u_1(a, b)]$. This is indeed the case as verified by the following theorem.

Theorem 4.4. Let π be a distribution over the product of function spaces $\mathcal{F}_1 \times \mathcal{F}_2$. There exists a stationary distribution σ_a of the Markov chain M_a for any fixed $a \in \mathcal{A}_1$ such that $\mathbb{E}_{(a,b) \sim \sigma_a} [u_1(a, b)] = \mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a))]$. Similarly, for every fixed action $b \in \mathcal{A}_2$, there exists a stationary distribution σ_b of M_b such that $\mathbb{E}_{(a,b) \sim \sigma_b} [u_2(a, b)] = \mathbb{E}_{(f,g) \sim \pi} [u_2(f(b), b)]$.

The proof of this theorem is constructive and can be found in the supplementary. With all of this notation we are ready to formally describe what no-policy regret play promises in the game setting in terms of an equilibrium.

Definition 4.5. A distribution π over $\mathcal{F}_1 \times \mathcal{F}_2$ is a *policy equilibrium* if for all fixed actions $a \in \mathcal{A}_1$ and $b \in \mathcal{A}_2$, which generate Markov chains M_a and M_b respectively, with stationary distributions σ_a and σ_b from Theorem 4.4, there exists a stationary distribution σ of the Markov chain M induced by π such that:

$$\begin{aligned} \mathbb{E}_{(a,b) \sim \sigma} [u_1(a, b)] &\geq \mathbb{E}_{(a,b) \sim \sigma_a} [u_1(a, b)] \\ \mathbb{E}_{(a,b) \sim \sigma} [u_2(a, b)] &\geq \mathbb{E}_{(a,b) \sim \sigma_b} [u_2(a, b)]. \end{aligned} \quad (3)$$

In other words, π is a policy equilibrium if there exists a stationary distribution σ of the Markov chain corresponding to π , such that, when actions are drawn according to σ , no player has incentive to change their action. For a simple example of a policy equilibrium see Section E in the supplementary.

4.1 Convergence to the set of policy equilibria

We have tried to formally capture the notion of equilibria in which player 1's deviation would lead to a reaction from player 2 and vice versa in Definition 4.5. This definition is inspired by the counterfactual guarantees of no policy regret play and we would like to check that if players' strategies yield sublinear policy regret then the play converges to a policy equilibrium. Since the definition of sublinear policy regret does not include a distribution over functional spaces but only works with empirical distributions of play, we would like to present our result in terms of distributions over the action space \mathcal{A} . Thus we begin by defining the set of all product distributions $\sigma \times \sigma_a \times \sigma_b$, induced by policy equilibria π as described in the previous subsection. Here σ_a and σ_b represent the deviation in strategy if player 1 changed to playing the fixed action $a \in \mathcal{A}_1$ and player 2 changed to playing the fixed action $b \in \mathcal{A}_2$ respectively as constructed in Theorem 4.4.

Definition 4.6. For a policy equilibrium π , let S_π be the set of all stationary distributions which satisfy the equilibrium inequalities (3), $S_\pi := \{\sigma \times \sigma_a \times \sigma_b : (a, b) \in \mathcal{A}\}$. Define $S = \bigcup_{\pi \in \Pi} S_\pi$, where Π is the set of all policy equilibria.

Our main result states that the sequence of empirical product distributions formed after T rounds of the game $\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b$ is going to converge to S . Here $\hat{\sigma}_a$ and $\hat{\sigma}_b$ denote the distributions of deviation in play, when player 1 switches to the fixed action $a \in \mathcal{A}_1$ and player 2 switches to the fixed action $b \in \mathcal{A}_2$ respectively. We now define these distributions formally.

Definition 4.7. Suppose player 1 is playing an algorithm with output at time t given by $f_t : \mathcal{A}_2^t \rightarrow \Delta \mathcal{A}_1$ i.e. $p_t^1 = f_t(b_{0:t-1})$. Similarly, suppose player 2 is playing an algorithm with output at time t given by $p_t^2 = g_t(a_{0:t-1})$. The empirical distribution at time T is $\hat{\sigma} := \frac{1}{T} \sum_{t=1}^T p_t$, where $p_t = p_t^1 \times p_t^2$ is the product distribution over \mathcal{A} at time t . Further let $(p_a^2)_t = g_t(a_{0:t-m}, a, \dots, a)$ denote the distribution at time t , provided that player 1 switched their strategy to the constant action $a \in \mathcal{A}_1$. Let δ_a denote the distribution over \mathcal{A}_1 which puts all the probability mass on action a . Let $(p_a)_t = \delta_a \times (p_a^2)_t$ be the product distribution over \mathcal{A} , corresponding to the change of play at time t . Denote by $\hat{\sigma}_a = \frac{1}{T} \sum_{t=1}^T (p_a)_t$ the empirical distribution corresponding to the change of play. The distribution $\hat{\sigma}_b$ is defined similarly.

Suppose that f_t and g_t are no-policy regret algorithms, then our main result states that the sequence $(\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b)_T$ converges to the set S .

Theorem 4.8. If the algorithms played by player 1 in the form of f_t and player 2 in the form of g_t give sub-linear policy regret sequences, then the sequence of product distributions $(\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b)_{T=1}^\infty$ converges weakly to the set S .

In particular if both players are playing MWU or Exp3, we know that they will have sublinear policy regret. Not surprisingly, we can show something slightly stronger as well. Let $\tilde{\sigma}$, $\tilde{\sigma}_a$ and $\tilde{\sigma}_b$ denote the empirical distributions of observed play corresponding to $\hat{\sigma}$, $\hat{\sigma}_a$ and $\hat{\sigma}_b$, i.e. $\tilde{\sigma} = \frac{1}{T} \delta_t$, where δ_t

denotes the Dirac distribution, putting all weight on the played actions at time t . Then these empirical distributions also converge to S almost surely.

4.2 Sketch of proof of the main result

The proof of Theorem 4.8 has three main steps. The first step defines the natural empirical Markov chains \widehat{M} , \widehat{M}_a and \widehat{M}_b from the empirical play $(p_t)_{t=1}^\infty$ (see Definition B.2) and shows that the empirical distributions $\widehat{\sigma}$, $\widehat{\sigma}_a$ and $\widehat{\sigma}_b$ are stationary distributions of the respective Markov chains. The latter is done in Lemma B.3. The next step is to show that the empirical Markov chains converge to Markov chains M , M_a and M_b induced by some distribution π over \mathcal{F} . In particular, we construct an empirical distribution $\widehat{\pi}$ and distributions $\widehat{\pi}_a$ and $\widehat{\pi}_b$ corresponding to player's deviations (see Definition B.5), and show that these induce the Markov chains \widehat{M} , \widehat{M}_a and \widehat{M}_b respectively (Lemma B.7). The distribution π we want is now the limit of the sequence $(\widehat{\pi})_T$. The final step is to show that π is a policy equilibrium. The proof goes by contradiction. Assume π is not a policy equilibrium, this implies that no stationary distribution of M and corresponding stationary distributions of M_a and M_b can satisfy inequalities (3). Since the empirical distributions $\widehat{\sigma}$, $\widehat{\sigma}_a$ and $\widehat{\sigma}_b$ of the play satisfies inequalities (3) up to an $o(1)$ additive factor, we can show, in Theorem B.8, that in the limit, the policy equilibrium inequalities are exactly satisfied. Combined with the convergence of \widehat{M} , \widehat{M}_a and \widehat{M}_b to M , M_a and M_b , respectively, this implies that stationary distributions of M , M_a and M_b , satisfying (3), giving a contradiction.

We would like to emphasize that the convergence guarantee of Theorem 4.8 does not rely on there being a unique stationary distribution of the empirical Markov chains \widehat{M} , \widehat{M}_a and \widehat{M}_b or their respective limits M , M_a , M_b . Indeed, Theorem 4.8 shows that any limit point of $\{(\widehat{\sigma}, \widehat{\sigma}_a, \widehat{\sigma}_b)_T\}_{T=1}^\infty$ satisfies the conditions of Definition 4.5. The proof does not require that any of the respective Markov chains have a unique stationary distribution, but rather requires only that $\widehat{\sigma}$ has sublinear policy regret. We would also like to remark that $\{(\widehat{\sigma}, \widehat{\sigma}_a, \widehat{\sigma}_b)_T\}_{T=1}^\infty$ need not have a unique limit and our convergence result only guarantees that the sequence is going to the set S . This is standard when showing that any type of no regret play converges to an equilibrium, see for example Stoltz and Lugosi (2007).

4.3 Relation of policy equilibria to CCEs

So far we have defined a new class of equilibria and shown that they correspond to no policy regret play. Furthermore, we know that if both players in a 2-player game play stable no external regret algorithms, then their play also has sublinear policy regret. It is natural to ask if every CCE is also a policy equilibrium: if σ is a CCE, is there a corresponding policy equilibrium π which induces a Markov chain M for which σ is a stationary distribution satisfying (3)? We show that the answer to this question is positive:

Theorem 4.9. *For any CCE σ of a 2-player game \mathcal{G} , there exists a policy-equilibrium π which induces a Markov chain M with stationary distribution σ .*

To prove this, we show that for any CCE we can construct stable no-external regret algorithm which converge to it, and so since stable no-external regret algorithms always converge to policy equilibria (Theorem 3.4), this implies the CCE is also a policy equilibrium.

However, we show the converse is not true: policy equilibria can give rise to behavior which is not a CCE. Our proof appeals to a utility sequence which is similar in spirit to the one in Theorem 3.2, but is adapted to the game setting.

Theorem 4.10. *There exists a 2-player game \mathcal{G} and product distributions $\sigma \times \sigma_a \times \sigma_b \in S$ (where S is defined in Definition 4.6 as the possible distributions of play from policy equilibria), such that σ is not a CCE of \mathcal{G} .*

In Section E of the supplementary we give a simple example of a policy equilibrium which is not a CCE.

5 Discussion

In this work we gave a new twist on policy regret by examining it in the game setting, where we introduced the notion of policy equilibrium and showed that it captures the behavior of no policy

regret players. While our characterization is precise, we view this as only the first step towards truly understanding policy regret and its variants in the game setting. Many interesting open questions remain. Even with our current definitions, since we now have a broader class of equilibria to consider it is natural to go back to the extensive literature in algorithmic game theory on the price of anarchy and price of stability and reconsider it in the context of policy equilibria. For example Roughgarden (2015) showed that in “smooth games” the worst CCE is no worse than the worst Nash. Since policy equilibria contain all CCEs (Theorem 4.9), is the same true for policy equilibria?

Even more interesting questions remain if we change our definitions to be more general. For example, what happens with more than 2 players? With three or more players, definitions of “reaction” by necessity become more complicated. Or what happens when m is not a constant? No policy regret algorithms exist for superconstant m , but our notion of equilibrium requires m to be constant in order for the Markov chains to make sense. Finally, what if we compare against deviations that are more complicated than a single action, in the spirit of swap regret or Φ -regret?

From an online learning perspective, note that our notion of on average stable and the definition of m -memory boundedness are different notions of stability. Is there one unified definition of “stable” which would allow us to give no policy regret algorithms against stable adversaries even outside of the game setting?

Acknowledgments

This work was supported in part by NSF BIGDATA grant IIS-1546482, NSF BIGDATA grant IIS-1838139, NSF CCF-1535987, NSF IIS-1618662, NSF CCF-1464239, and NSF AITF CCF-1535887.

References

- Zeyuan Allen-Zhu and Yuanzhi Li. Lazysvd: Even faster svd decomposition yet without agonizing pain. In *Advances in Neural Information Processing Systems*, pages 974–982, 2016.
- Raman Arora, Ofer Dekel, and Ambuj Tewari. Deterministic MDPs with adversarial rewards and bandit feedback. In *Proceedings on Uncertainty in Artificial Intelligence (UAI)*, 2012a.
- Raman Arora, Ofer Dekel, and Ambuj Tewari. Online bandit learning against an adaptive adversary: from regret to policy regret. In *Proceedings of International Conference on Machine Learning (ICML)*, 2012b.
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012c.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, 2007.
- Ofer Dekel, Jian Ding, Tomer Koren, and Yuval Peres. Bandits with switching costs: $T^{2/3}$ regret. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 459–467. ACM, 2014.
- Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009a.
- Eyal Even-Dar, Yishay Mansour, and Uri Nadav. On the convergence of regret minimization dynamics in concave games. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 523–532. ACM, 2009b.
- Daniela Pucci De Farias and Nimrod Megiddo. Combining expert advice in reactive environments. *Journal of the ACM (JACM)*, 53(5):762–799, 2006.
- Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40, 1997.

- Drew Fudenberg and David K Levine. Conditional universal consistency. *Games and Economic Behavior*, 29(1-2):104–130, 1999.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Elad Hazan and Satyen Kale. Computational equivalence of fixed points and no regret algorithms, and convergence to equilibria. In *Advances in Neural Information Processing Systems*, pages 625–632, 2008.
- Hoda Heidari, Michael Kearns, and Aaron Roth. Tight policy regret bounds for improving and decaying bandits. In *IJCAI*, pages 1562–1570, 2016.
- Sham Machandranath Kakade et al. *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England, 2003.
- Tomer Koren, Roi Livni, and Yishay Mansour. Bandits with movement costs and adaptive pricing. In *Proceedings of the 2017 Conference on Learning Theory*, pages 1242–1268, 2017a.
- Tomer Koren, Roi Livni, and Yishay Mansour. Multi-armed bandits with metric movement costs. In *Advances in Neural Information Processing Systems*, pages 4119–4128, 2017b.
- Neri Merhav, Erik Ordentlich, Gadiel Seroussi, and Marcelo J Weinberger. On sequential strategies for loss functions with memory. *IEEE Transactions on Information Theory*, 48(7):1947–1958, 2002.
- Mehryar Mohri and Scott Yang. Conditional swap regret and conditional correlated equilibrium. In *Advances in Neural Information Processing Systems*, pages 1314–1322, 2014.
- Mehryar Mohri and Scott Yang. Online learning with transductive regret. In *Advances in Neural Information Processing Systems*, pages 5220–5230, 2017.
- Gergely Neu, Andras Antos, András György, and Csaba Szepesvári. Online markov decision processes under bandit feedback. In *Advances in Neural Information Processing Systems*, pages 1804–1812, 2010.
- Tim Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)*, 62(5):32, 2015.
- Ankan Saha, Prateek Jain, and Ambuj Tewari. The interplay between stability and regret in online learning. *arXiv preprint arXiv:1211.6158*, 2012.
- Gilles Stoltz and Gábor Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 59(1):187–208, 2007.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- Csaba Szepesvári. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.
- Jia Yuan Yu, Shie Mannor, and Nahum Shimkin. Markov decision processes with arbitrary reward processes. *Mathematics of Operations Research*, 34(3):737–757, 2009.

Supplementary Material to “Policy Regret in Repeated Games”

A Proofs of results from Section 3

Theorem A.1. *Let the space of actions be $\{0, 1\}$ and define the m -memory bounded utility functions as follows:*

$$f_t(a_{t-m+1}, \dots, a_t) = \begin{cases} 1 & a_{t-m+i} = a_{t-m+i+1} = 1 \text{ for } i \in \{1, \dots, m-2\} \wedge a_{t-1} \neq a_t \\ \frac{1}{2} & a_{t-m+i} = a_{t-m+i+1} = 1 \text{ for } i \in \{1, \dots, m-1\} \\ 0 & \text{otherwise} \end{cases}.$$

Assuming $m \geq 3$ is a fixed constant (independent of T) any sequence with sublinear policy regret will have linear external regret and every sequence with sublinear external regret will have linear policy regret.

Proof of Theorem 3.2. Let $(a_t)_{t=1}^T$ be a sequence with sublinear policy regret. Then this sequence has utility at least $\frac{T}{2} - o(T)$ and so there are at most $o(T)$ terms in the sequence which are not equal to 1. Let the subsequence consisting of all $a_t = 0$ be indexed by \mathcal{I} . Define $\tilde{\mathcal{I}} = \{t, t+1, \dots, t+m-1 : t \in \mathcal{I}\}$ and consider the subsequence of functions $(f_t)_{t \in \tilde{\mathcal{I}}}$. This is precisely the sequence of functions which have utility $\frac{1}{2}$ with respect to the sequence of play $(a_t)_{t=1}^T$. Notice that the length of this sequence is at least $\frac{T}{2} - o(T) = T - o(T)$. The utility of this sequence is $\sum_{t \in \tilde{\mathcal{I}}} f_t(a_{t-m+1}, \dots, a_t) = \sum_{t \in \tilde{\mathcal{I}}} f_t(1, \dots, 1) = \frac{T - o(T)}{2}$, however, this subsequence has linear external regret, since $\sum_{t \in \tilde{\mathcal{I}}} f_t(a_{t-m+1}, \dots, a_{t-1}, 0) = \sum_{t \in \tilde{\mathcal{I}}} f_t(1, \dots, 1, 0) = T - o(T)$. Thus the external regret of $(a_t)_{t=1}^T$ is

$$\begin{aligned} \sum_{t=1}^T [f_t(a_{t-m+1}, \dots, 0) - f_t(a_{t-m+1}, \dots, a_t)] &= \sum_{t \notin \tilde{\mathcal{I}}} [f_t(a_{t-m+1}, \dots, 0) - f_t(a_{t-m+1}, \dots, a_t)] \\ &\quad + \sum_{t \in \tilde{\mathcal{I}}} [f_t(a_{t-m+1}, \dots, 0) - f_t(a_{t-m+1}, \dots, a_t)] \\ &\geq \frac{T - o(T)}{2} + \sum_{t \in \tilde{\mathcal{I}}} [f_t(a_{t-m+1}, \dots, 0) - f_t(a_{t-m+1}, \dots, a_t)] \\ &\geq \frac{T}{2} - o(T), \end{aligned}$$

where the last inequality follows from the fact that the cardinality of $\tilde{\mathcal{I}}$ is at most $o(T)$ and thus $\sum_{t \in \tilde{\mathcal{I}}} [f_t(a_{t-m+1}, \dots, 0) - f_t(a_{t-m+1}, \dots, a_t)] \geq -o(T)$.

Assume that $(a_t)_{t=1}^T$ has sublinear external regret. From the above argument, it follows that the utility of the sequence is at most $o(T)$ (otherwise if the sequence has utility $\omega(T)$, we can repeat the previous argument and get a contradiction with the fact the sequence is no-external regret). This implies that the policy regret of the sequence is $\sum_{t=1}^T [f_t(1, 1, \dots, 1) - f_t(a_{t-m+1}, \dots, a_t)] = \frac{T}{2} - o(T)$. \square

Theorem A.2. *Let $(a_t)_{t=1}^T$ and $(b_t)_{t=1}^T$ be the action sequences of player 1 and 2 respectively and suppose that they are coming from no-regret algorithms with regrets $R_1(T)$ and $R_2(T)$ respectively. Assume that the algorithms are on average $(m, S(T))$ stable with respect to the ℓ_2 norm. Then*

$$\mathbb{E} \left[u_2(f_t(b_{0:t-m+1}, b, \dots, b), b) - \sum_{t=1}^T u_2(f_t(b_{0:t-1}), b_t) \right] \leq \|P_2\| S(T) + R_2(T),$$

for any fixed action $b \in \mathcal{A}_2$, where P_2 is the utility matrix of player 2. A similar inequality holds for any fixed action $a \in \mathcal{A}_1$ and the utility of player 1.

Proof of Theorem 3.4.

$$\begin{aligned}
& \mathbb{E} \left[u_2(f_t(b_0, \dots, b_{t-m+1}, b, \dots, b), b) - \sum_{t=1}^T u_2(f_t(b_0, \dots, b_{t-2}, b_{t-1}), b_t) \right] \\
&= \mathbb{E} \left[u_2(f_t(b_0, \dots, b_{t-m+1}, b, \dots, b), b) - \sum_{t=1}^T u_2(f_t(b_0, \dots, b_{t-2}, b_{t-1}), b) \right] \\
&+ \mathbb{E} \left[u_2(f_t(b_0, \dots, b_{t-1}), b) - \sum_{t=1}^T u_2(f_t(b_0, \dots, b_{t-2}, b_{t-1}), b_t) \right] \\
&\leq \sum_{t=1}^T \|b^\top P_2\|_2 \|f_t(b_0, \dots, b_{t-m+1}, b, \dots, b) - f_t(b_0, \dots, b_{t-2}, b_{t-1})\|_2 + R_2(T) \\
&\leq \|b\|_1 \|P_2\| S(T) + R_2(T) = \|P_2\| S(T) + R_2(T)
\end{aligned}$$

where the first inequality holds by Cauchy-Schwartz, the second inequality holds by using the m -stability of the algorithm, together with the inequality between l_1 and l_2 norms. \square

The next theorems show that MWU and Exp3 are stable algorithms.

Theorem A.3. *MWU is an on average $(m, m\sqrt{T})$ stable algorithm with respect to ℓ_1 , for any $m < o(\sqrt{T})$.*

Proof of Theorem A.3. We think of MWU as Exponentiated Gradient (EG) where the loss vector has i -th entry equal to the negative utility if the player decided to play action i . Let the observed loss vector at time j be \widehat{l}_j and the output distribution be p_j , then the update of EG can be written as $p_{j+1} = \arg \min_{p \in \mathcal{C}} \langle \widehat{l}_j, p \rangle + \frac{1}{\eta} D(p, p_j)$, where \mathcal{C} is the simplex of the set of possible actions and D is the KL-divergence. Using Lemma 3 in Saha et al. (2012), with $f(p) = \langle \widehat{l}_j, p \rangle + \frac{1}{\eta} D(p, p_j)$ and the fact that the KL-divergence is 1-strongly convex over the simplex \mathcal{C} with respect to the ℓ_1 norm, we have:

$$\begin{aligned}
\frac{1}{2\eta} \|p_j - p_{j+1}\|_1^2 &\leq f(p_j) - f(p_{j+1}) = \langle p_j - p_{j+1}, \widehat{l}_j \rangle - \frac{1}{\eta} D(p_{j+1}, p_j) \\
&\leq \|p_j - p_{j+1}\|_1 \|\widehat{l}_j\|_\infty \leq \|p_j - p_{j+1}\|_1,
\end{aligned}$$

where the second inequality follows from Hölder's inequality and the fact that $D(p_{j+1}, p_j) \geq 0$. Thus with step size $\eta \sim \sqrt{\frac{1}{T}}$, we have $\|p_j - p_{j+1}\|_1 \leq \frac{1}{2\sqrt{T}}$. Using triangle inequality, we can get $\|p_{j-1} - p_{j+1}\|_1 \leq \|p_{j-1} - p_j\|_1 + \|p_j - p_{j+1}\|_1 \leq \frac{2}{2\sqrt{T}}$ and induction shows that $\|p_{j-m+1} - p_{j+1}\|_1 \leq \frac{m}{2\sqrt{T}}$. Suppose for the last m iterations, a fixed loss function l_a was played instead and the resulting output of the algorithm becomes \tilde{p}_{j+1} . Then using the same argument as above we have $\|p_{j-m+1} - \tilde{p}_{j+1}\|_1 \leq \frac{m}{2\sqrt{T}}$ and thus $\|p_{j+1} - \tilde{p}_{j+1}\|_1 \leq \frac{m}{\sqrt{T}}$. Summing over all T rounds concludes the proof. \square

Theorem A.4. *Exp3 is an on average $(m, m\sqrt{T})$ stable algorithm with respect to ℓ_1 , for any $m < o(\sqrt{T})$.*

Proof of Theorem A.4. The update at time t , conditioning on the the draw being i , is given by

$$p_{t+1}^i = \frac{w_t^i \exp\left(\frac{\gamma}{kp_t^i} u_t^i\right)}{w_t^i \exp\left(\frac{\gamma}{kp_t^i} u_t^i\right) + \sum_{j \neq i} w_t^j} \text{ and for } j \neq i, p_{t+1}^j = \frac{w_t^j}{w_t^i \exp\left(\frac{\gamma}{kp_t^i} u_t^i\right) + \sum_{j \neq i} w_t^j}, \text{ where } u_t \text{ is the utility}$$

vector at time t , w_t is the weight vector at time t i.e. $w_{t+1}^i = w_t^i \exp\left(\frac{\gamma}{kp_t^i} u_t^i\right)$, $w_{t+1}^j = w_t^j$, and k is the number of actions. We have the following bound:

$$\begin{aligned}
|p_{t+1}^i - p_t^i| &= \left| \frac{w_t^i \exp\left(\frac{\gamma}{kp_t^i} u_t^i\right)}{w_t^i \exp\left(\frac{\gamma}{kp_t^i} u_t^i\right) + \sum_{j \neq i} w_t^j} - \frac{w_t^i}{\sum_j w_t^j} \right| \leq \left| \frac{w_t^i (\exp\left(\frac{\gamma}{kp_t^i} u_t^i\right) - 1)}{\sum_j w_t^j} \right| \\
&= p_t^i (\exp\left(\frac{\gamma}{kp_t^i} u_t^i\right) - 1) \leq p_t^i 2 \frac{\gamma}{kp_t^i} u_t^i \leq 2 \frac{\gamma}{k},
\end{aligned}$$

where the first inequality uses the fact that $p_{t+1}^i \geq p_t^i$ and the second inequality uses the choice of γ together with $\exp(x) \leq 2x + 1$ for $x \in [0, 1]$. Similarly for $j \neq i$, we have:

$$\begin{aligned} |p_{t+1}^j - p_t^j| &= \left| \frac{w_t^j}{w_t^i \exp\left(\frac{\gamma}{kp_t^i} u_t^i\right) + \sum_{j \neq i} w_t^j} - \frac{w_t^j}{\sum_j w_t^j} \right| \leq \left| \frac{w_t^j}{\sum_j w_t^j} \left(1 - \frac{1}{\exp\left(\frac{\gamma u_t^i}{kp_t^i}\right)}\right) \right| \\ &= p_t^j \left(1 - \exp\left(-\frac{\gamma u_t^i}{kp_t^i}\right)\right) \leq p_t^j \frac{\gamma u_t^i}{kp_t^i} \leq \frac{p_t^j}{kp_t^i} \gamma, \end{aligned}$$

where we have used $p_{t+1}^j \leq p_t^j$ and $\exp(-x) \geq 1 - x$ for all x . We can now proceed to bound $\mathbb{E}_{i_t} [\|p_{t+1} - p_t\|_1 | i_{1:t-1}]$, where i_t is the random variable denoting the draw at time t :

$$\begin{aligned} \mathbb{E}_{i_t} [\|p_{t+1} - p_t\|_1 | i_{1:t-1}] &= \sum_i p_t^i \|p_{t+1} - p_t\|_1 = \sum_i p_t^i \left(|p_{t+1}^i - p_t^i| + \sum_{j \neq i} |p_{t+1}^j - p_t^j| \right) \\ &\leq 2\gamma + \sum_{i,j} p_t^i \frac{p_t^j}{kp_t^i} \gamma = 3\gamma. \end{aligned}$$

Setting $\gamma \sim \frac{1}{\sqrt{T}}$ finishes the proof. \square

Combining the above results together, we can show that both players will also have no-policy regret for any $m < o(\sqrt{T})$.

Corollary A.5. *Let $(a_t)_{t=1}^T$ and $(b_t)_{t=1}^T$, be the action sequences of players 1 and 2 respectively and suppose that the sequences are coming from MWU. Then for any fixed m , it holds:*

$$\mathbb{E} \left[u_2(f_t(b_{0:t-m+1}, b, \dots, b), b) - \sum_{t=1}^T u_2(f_t(b_{0:t-1}), b_t) \right] \leq O(\|P_2\| m \sqrt{T}),$$

for any fixed action $b \in \mathcal{A}_2$, where f_t are the functions corresponding to the MWU algorithm used by player 1.

Proof. From Theorem A.3 it follows that MWU is on average $(m, m\sqrt{T})$ stable and has regret at most $O(\sqrt{T})$. \square

B Proofs of results from Section 4

Theorem B.1. *Let π be a distribution over the product of function spaces $\mathcal{F}_1 \times \mathcal{F}_2$. There exists a stationary distribution $\tilde{\sigma}$ of the Markov chain M_a for any fixed $a \in \mathcal{A}_1$ and $(f, g) \sim \pi$ such that $\mathbb{E}_{(a,b) \sim \tilde{\sigma}} [u_1(a, b)] = \mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a))]$*

Proof of Theorem 4.4. Note that, by definition $(M_a)_{(\tilde{a}, \tilde{b}), (\tilde{a}, \tilde{b})} = 0$ if $\tilde{a} \neq a$ and $M_{(\tilde{a}, \tilde{b}), (\tilde{a}, \tilde{b})} = \sum_{f,g: g(\tilde{a})=\tilde{b}} \pi(f, g)$ if $\tilde{a} = a$. Consider the distribution $\tilde{\sigma}$ over \mathcal{A} , where $\tilde{\sigma}_{(\tilde{a}, \tilde{b})} = 0$ if $\tilde{a} \neq a$ and $\tilde{\sigma}_{(\tilde{a}, \tilde{b})} = \sum_{f,g: g(a)=\tilde{b}} \pi(f, g)$ if $\tilde{a} = a$. We now show that $\tilde{\sigma}$ is a stationary distribution of M_a :

$$\begin{aligned} (\tilde{\sigma}^\top M_a)_{(a, \tilde{b})} &= \sum_{(\tilde{a}, \tilde{b})} \tilde{\sigma}_{(\tilde{a}, \tilde{b})} (M_a)_{(\tilde{a}, \tilde{b}), (a, \tilde{b})} = \sum_{\tilde{b}} \tilde{\sigma}_{(a, \tilde{b})} (M_a)_{(a, \tilde{b}), (a, \tilde{b})} \\ &= \sum_{\tilde{b}} \left(\sum_{f,g: g(a)=\tilde{b}} \pi(f, g) \right) \left(\sum_{f,g: g(a)=\tilde{b}} \pi(f, g) \right) \\ &= \left(\sum_{f,g: g(a)=\tilde{b}} \pi(f, g) \right) \left(\sum_{\tilde{b}} \sum_{f,g: g(a)=\tilde{b}} \pi(f, g) \right) \\ &= \sum_{f,g: g(a)=\tilde{b}} \pi(f, g) = \tilde{\sigma}_{(a, \tilde{b})}. \end{aligned}$$

Finally, notice that:

$$\mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a))] = \sum_{b \in \mathcal{A}_2} u_1(a, b) \mathbb{P}[g(a) = b] = \sum_{b \in \mathcal{A}_2} u_1(a, b) \sum_{f, g: g(a)=b} \pi(f, g).$$

□

B.1 Sketch of proof of the main result

The proof of Theorem 4.8 has three main steps. The first step defines the natural empirical Markov chains \widehat{M} , \widehat{M}_a and \widehat{M}_b from the empirical play $(p_t)_{t=1}^T$ and shows that the empirical distributions $\widehat{\sigma}$, $\widehat{\sigma}_a$ and $\widehat{\sigma}_b$ are stationary distributions of the respective Markov chains. The next step is to show that the empirical Markov chains converge to Markov chains M , M_a and M_b induced by some distribution π over \mathcal{F} . The final step is to show that π is a policy equilibrium.

We begin with the definition of the empirical Markov chains.

Definition B.2. Let the empirical Markov chain be \widehat{M} , with $\widehat{M}_{i,j} = \frac{\frac{1}{T} \sum_{t=1}^T p_t(x_i) p_t(x_j)}{\frac{1}{T} \sum_{t=1}^T p_t(x_i)}$ if $\frac{1}{T} \sum_{t=1}^T p_t(x_i) \neq 0$ and 0 otherwise, where p_t is defined in 4.7. For any fixed $a \in \mathcal{A}_1$, let the empirical Markov chain corresponding to the deviation in play of player 1 be \widehat{M}_a , with $(\widehat{M}_a)_{i,j} = \frac{\frac{1}{T} \sum_{t=1}^T (p_a)_t(x_i) (p_a)_t(x_j)}{\frac{1}{T} \sum_{t=1}^T (p_a)_t(x_i)}$, if $\frac{1}{T} \sum_{t=1}^T (p_a)_t(x_i) \neq 0$ and 0 otherwise, where $(p_a)_t$ is defined in 4.7. The Markov chain \widehat{M}_b is defined similarly for any $b \in \mathcal{A}_2$.

The intuition behind constructing these Markov chains is as follows – if we were only provided with the observed empirical play $(x_t)_{t=1}^T = (a_t, b_t)_{t=1}^T$ and someone told us that the x_t 's were coming from a Markov chain, we could try to build an estimator of the Markov chain by approximating each of the transition probabilities. In particular the estimator of transition from state i to state j is given by $\tilde{M}_{i,j} = \frac{\frac{1}{T} \sum_{t=1}^T \delta_{t-1}(x_i) \delta_t(x_j)}{\frac{1}{T} \sum_{t=1}^T \delta_t(x_i)}$, where $\delta_t(x_i) = 1$ if x_i occurred at time t and 0 otherwise. When the players are playing according to a no-regret algorithm i.e. at time t , x_t is sampled from p_t , it is possible to show that $\tilde{M}_{i,j}$ concentrates to $\widehat{M}_{i,j}$ (see section B.4). Not only does \widehat{M} arise naturally, but it turns out that the empirical distribution $\widehat{\sigma}$ defined in 4.7 is also a stationary distribution of \widehat{M} .

Lemma B.3. *The distribution of play $\widehat{\sigma} = \frac{1}{T} \sum_{t=1}^T p_t$ is a stationary distribution of \widehat{M} . Similarly the distributions $\widehat{\sigma}$, $\widehat{\sigma}_a$, $\widehat{\sigma}_a$, $\widehat{\sigma}_b$ and $\widehat{\sigma}_b$ are stationary distributions of the Markov chains \widehat{M} , \widehat{M}_a , \widehat{M}_a , \widehat{M}_b and \widehat{M}_b respectively.*

Proof. We show the result for $\widehat{\sigma}$ and \widehat{M} . The rest of the results can then be derived in the same way.

$$(\widehat{\sigma}^\top \widehat{M})_j = \sum_{i=1}^{|\mathcal{A}|} \left(\frac{1}{T} \sum_{t=1}^T p_t(x_i) \right) \frac{\frac{1}{T} \sum_{t=1}^T p_t(x_i) p_t(x_j)}{\frac{1}{T} \sum_{t=1}^T p_t(x_i)} = \frac{1}{T} \sum_{t=1}^T p_t(x_j) \sum_{i=1}^{|\mathcal{A}|} p_t(x_i) = \frac{1}{T} \sum_{t=1}^T p_t(x_j),$$

where the first equality holds because the i -th entry of the vector $\widehat{\sigma}$ is exactly $\frac{1}{T} \sum_{t=1}^T p_t(x_i)$ and the (i, j) -th entry of \widehat{M} by definition is $\frac{\frac{1}{T} \sum_{t=1}^T p_t(x_i) p_t(x_j)}{\frac{1}{T} \sum_{t=1}^T p_t(x_i)}$, and the last equality holds because p_t is a distribution over actions so $\sum_{i=1}^{|\mathcal{A}|} p_t(x_i) = 1$. □

Suppose that both players are playing MWU for T rounds. Then Lemma B.3 together with Theorems 3.4 and the stability of MWU imply that $\mathbb{E}_{(a,b) \sim \widehat{\sigma}} [u_1(a, b)] \geq \mathbb{E}_{(a,b) \sim \widehat{\sigma}_a} [u_1(a, b)] - O(m/\sqrt{T})$. A similar inequality holds for player 2 and $\widehat{\sigma}_b$. As $T \rightarrow \infty$, the inequality above becomes similar to (3). This will play a crucial role in the proof of our convergence result, which shows that $\widehat{\sigma}$, $\widehat{\sigma}_a$ and $\widehat{\sigma}_b$ converge to the set of policy equilibria. We would also like to guarantee that the empirical distributions of observed play $\tilde{\sigma}$, $\tilde{\sigma}_a$ and $\tilde{\sigma}_b$ also converge to this set. To show this second result, we are going to proof that $\tilde{\sigma}$ approaches $\widehat{\sigma}$ almost surely as T goes to infinity.

Lemma B.4. *Let $\widehat{\sigma} = \frac{1}{T} \sum_{t=1}^T p_t$ be the empirical distribution after T rounds of the game and let $\tilde{\sigma} = \frac{1}{T} \sum_{t=1}^T \delta_t$ be the empirical distribution of observed play. Then $\limsup_{T \rightarrow \infty} \|\tilde{\sigma} - \widehat{\sigma}\|_1 = 0$ almost surely. Similarly, for the distributions corresponding to deviation in play we have $\limsup_{T \rightarrow \infty} \|\tilde{\sigma}_a - \widehat{\sigma}_a\|_1 = 0$ and $\limsup_{T \rightarrow \infty} \|\tilde{\sigma}_b - \widehat{\sigma}_b\|_1 = 0$ almost surely.*

Proof. The proof follows the same reasoning as the proof of lemma B.19. \square

Our next step is to show that the empirical Markov chains \widehat{M} converge to a Markov chain M induced by some distribution π over the functional space \mathcal{F} . We do so by constructing a sequence of empirical distributions $\widehat{\pi}$ over \mathcal{F} , based on the players' strategies, which induce \widehat{M} . We can then consider every convergent subsequence of $(\widehat{\pi})_{T_\ell=1}^\infty$ with limit point π and argue that the corresponding sequence $(\widehat{M})_{T_\ell=1}^\infty$ of Markov chains converges to the Markov chain induced by π .

Definition B.5. Let $\widehat{\pi}$ be the distribution over \mathcal{F} , such that the probability to sample any fixed $f : \mathcal{A}_2 \rightarrow \mathcal{A}_1$ and $g : \mathcal{A}_1 \rightarrow \mathcal{A}_2$ is $\widehat{\pi}(f, g) = \prod_{i \in [\mathcal{A}]} \frac{\sum_t p_t(x_i) p_t(y_i)}{\sum_{t=1} p_t(x_i)}$, where $x_i = (a_i, b_i)$ and $y_i = (f(b_i), g(a_i))$. Similarly, let $\widehat{\pi}_a$ and $\widehat{\pi}_b$ be the distributions over \mathcal{F} constructed as above but by using the empirical distribution of deviated play induced by player 1 deviating to action $a \in \mathcal{A}_1$ and player 2 deviating to action $b \in \mathcal{A}_2$.

The next lemma checks that $\widehat{\pi}$ is really a probability distribution.

Lemma B.6. *The functionals $\widehat{\pi}$, $\widehat{\pi}_a$ and $\widehat{\pi}_b$ are all probability distributions.*

Proof. Consider the space of all transition events for a fixed (a, b) pair i.e. $\mathcal{S}_{(a,b)} = \{((a', b') \times (a, b)) : (a', b') \in \mathcal{A}\}$. There is an inherent probability measure on this set, given by $\mathbb{P}[(a', b') \times (a, b)] = \frac{\sum_t p_t(a, b) p_t(a', b')}{\sum_t p_t(a, b)}$. It is easy to see that this is a probability measure, since the measure of the whole set is exactly

$$\sum_{(a', b') \in \mathcal{A}} \frac{\sum_t p_t(a, b) p_t(a', b')}{\sum_t p_t(a, b)} = \frac{\sum_t p_t(a, b) \sum_{(a', b') \in \mathcal{A}} p_t(a', b')}{\sum_t p_t(a, b)} = 1.$$

The set of all \mathcal{F} can exactly be thought of as $\times_{(a,b) \in \mathcal{A}} \mathcal{S}_{(a,b)}$ and the function $\widehat{\pi}$ defined in B.5 is precisely the product measure on that set. Similar arguments show that $\widehat{\pi}_a$ and $\widehat{\pi}_b$ are probability distributions. \square

The proof of the above lemma reveals something interesting about the construction of $\widehat{\pi}$. Fix the actions $(a, b) \in \mathcal{A}$. Then the probability to sample a function pair (f, g) which map (a, b) to (a', b') i.e. $a' = f(a)$ and $b' = g(b)$ is precisely equal to the entry $\widehat{M}_{(a,b),(a',b')}$ of the empirical Markov chain. Since every function pair $(f, g) \in \mathcal{F}$ is determined by the way \mathcal{A} is mapped, and we have already have a probability distribution for a fixed mapping (a, b) to (a', b') , we can just extend this to $\widehat{\pi}$ by taking the product distribution over all pairs $(a, b) \in \mathcal{A}$. This construction gives us exactly that \widehat{M} is induced by $\widehat{\pi}$.

Lemma B.7. *Let \widehat{M} , \widehat{M}_a and \widehat{M}_b be the empirical Markov chains defined in B.2, then the induced Markov chain from $\widehat{\pi}$ is exactly \widehat{M} and the induced Markov chains from $\widehat{\pi}_a$ and $\widehat{\pi}_b$ are exactly \widehat{M}_a and \widehat{M}_b .*

Proof. Consider $\widehat{M}_{(a,b),(a',b')} = \frac{\sum_{t=1}^T p_t(a', b') p_t(a, b)}{\sum_{t=1}^T p_t(a, b)}$. The transition probability induced by $\widehat{\pi}$ is exactly

$$\begin{aligned} \mathbb{P}[(a', b')|(a, b)] &= \sum_{(f,g):(f(b),g(a))=(a',b')} \widehat{\pi}(f, g) = \sum_{(f,g):(f(b),g(a))=(a',b')} \prod_{i \in [\mathcal{A}]} \frac{\sum_t p_t(x_i) p_t(y_i)}{\sum_{t=1} p_t(x_i)} \\ &= \sum_{(f,g):(f(b),g(a))=(a',b')} \frac{\sum_t p_t(a, b) p_t(a', b')}{\sum_{t=1} p_t(a, b)} \prod_{i \in [\mathcal{A}], (x_i, y_i) \neq ((a, b), (a', b'))} \frac{\sum_t p_t(x_i) p_t(y_i)}{\sum_{t=1} p_t(x_i)} \\ &= \frac{\sum_t p_t(a, b) p_t(a', b')}{\sum_{t=1} p_t(a, b)} \sum_{(f,g):(f(b),g(a))=(a',b')} \prod_{i \in [\mathcal{A}], (x_i, y_i) \neq ((a, b), (a', b'))} \frac{\sum_t p_t(x_i) p_t(y_i)}{\sum_{t=1} p_t(x_i)} \\ &= \frac{\sum_t p_t(a, b) p_t(a', b')}{\sum_{t=1} p_t(a, b)}, \end{aligned}$$

where the last equality holds, because for fixed (f, g) with $x_i = (a_i, b_i)$ and $y_i = (f(b_i), g(a_i))$, the product $\prod_{i \in [\mathcal{A}], (x_i, y_i) \neq ((a, b), (a', b'))} \frac{\sum_t p_t(x_i) p_t(y_i)}{\sum_{t=1} p_t(x_i)}$ is exactly the conditional probability $\widehat{\pi}((f, g)|(f(b), g(a)) = (a', b'))$. The result for $\widehat{\pi}_a$ and $\widehat{\pi}_b$ is shown similarly. \square

The last step of the proof is to show that any limit point π of $(\hat{\pi})_T$ is necessarily a policy equilibrium. This is done through an argument by contradiction. In particular we assume that a limit point π is not a policy equilibrium. The limit point π induces a Markov chain M , which we can show is the limit point of the corresponding subsequence of $(\hat{M})_T$ by using lemma B.7. Since π is not a policy equilibrium, no stationary distribution of M can satisfy the inequalities (3). We can now show that the subsequence of $(\hat{\sigma})_T$ which are stationary distributions of the corresponding \hat{M} 's, converges to a stationary distribution of M . This, however, is a contradiction because of the next theorem.

Theorem B.8. *Let P be the set of all product distributions $\sigma \times \sigma_a \times \sigma_b$ which satisfy the inequalities in 3:*

$$\begin{aligned}\mathbb{E}_{(a,b) \sim \sigma} [u_1(a, b)] &\geq \mathbb{E}_{(a,b) \sim \sigma_a} [u_1(a, b)] \\ \mathbb{E}_{(a,b) \sim \sigma} [u_2(a, b)] &\geq \mathbb{E}_{(a,b) \sim \sigma_b} [u_2(a, b)].\end{aligned}$$

Let $\hat{\sigma}^T$ be the empirical distribution of play after T rounds and let $\hat{\sigma}_a^T$ be the empirical distribution when player 1 switches to playing action a and define $\hat{\sigma}_b^T$ similarly for player 2. Then the product distribution $\hat{\sigma}^T \times \hat{\sigma}_a^T \times \hat{\sigma}_b^T$ converges to weakly to the set P .

Proof. Theorem B.8 follows from the fact that convergence in the Prokhorov metric implies weak convergence. First notice that by Prokhorov's theorem $\mathcal{P}(\mathcal{A})$ is a compact metric space with the Prokhorov metric. Thus by Tychonoff's Theorem the product space $\mathcal{P}(\mathcal{A})^3$ is compact in the maximum metric. Suppose for a contradiction that the sequence $(\hat{\sigma}^T \times \hat{\sigma}_a^T \times \hat{\sigma}_b^T)_T$ does not converge to the set S . This implies that there exists some subsequence $(\hat{\sigma}^k \times \hat{\sigma}_a^k \times \hat{\sigma}_b^k)_k$, converging to some $\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b \notin S$. If $\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b \notin S$, then either $\mathbb{E}_{(a,b) \sim \sigma} [u_1(a, b)] < \mathbb{E}_{(a,b) \sim \sigma_a} [u_1(a, b)]$ or $\mathbb{E}_{(a,b) \sim \sigma} [u_2(a, b)] < \mathbb{E}_{(a,b) \sim \sigma_b} [u_2(a, b)]$. WLOG suppose the first inequality holds. From our assumption, the continuity of u_1 and the definition of the maximum metric we have $\lim_{k \rightarrow \infty} \mathbb{E}_{(a,b) \sim \hat{\sigma}^k} [u_1(a, b)] = \mathbb{E}_{(a,b) \sim \hat{\sigma}} [u_1(a, b)]$ and $\lim_{k \rightarrow \infty} \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^k} [u_1(a, b)] = \mathbb{E}_{(a,b) \sim \hat{\sigma}_a} [u_1(a, b)]$. Notice that by the fact $\hat{\sigma}_a^k$ is the average empirical distribution if player 1 changed its play to the fixed action $a \in \mathcal{A}_1$ and $\hat{\sigma}^k$ being the average empirical distribution it holds that $\mathbb{E}_{(a,b) \sim \hat{\sigma}^k} [u_1(a, b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^k} [u_1(a, b)] \geq -o(1)$ and thus $\lim_{k \rightarrow \infty} [\mathbb{E}_{(a,b) \sim \hat{\sigma}^k} [u_1(a, b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^k} [u_1(a, b)]] \geq 0$. The above implies:

$$\begin{aligned}0 &\leq \lim_{k \rightarrow \infty} [\mathbb{E}_{(a,b) \sim \hat{\sigma}^k} [u_1(a, b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^k} [u_1(a, b)]] \\ &= \lim_{k \rightarrow \infty} \mathbb{E}_{(a,b) \sim \hat{\sigma}^k} [u_1(a, b)] - \lim_{k \rightarrow \infty} \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^k} [u_1(a, b)] \\ &= \mathbb{E}_{(a,b) \sim \hat{\sigma}} [u_1(a, b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}_a} [u_1(a, b)] < 0,\end{aligned}$$

which is a contradiction. Since $\mathcal{A} \times \mathcal{A} \times \mathcal{A}$ is separable then convergence in the Prokhorov metric in $\mathcal{P}(\mathcal{A} \times \mathcal{A} \times \mathcal{A})$ is equivalent to weak convergence. Again we can argue by contradiction – if we assume that $(\hat{\sigma}^T \times \hat{\sigma}_a^T \times \hat{\sigma}_b^T)_T$ doesn't converge to the set S in the Prokhorov metric, then there exists some subsequence $(\hat{\sigma}^k \times \hat{\sigma}_a^k \times \hat{\sigma}_b^k)_k$ which converges to some $\mu \in \mathcal{P}(\mathcal{A}^3)$ such that $\mu \notin S$. First we argue that μ must be a product measure i.e. $\mu = (\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b)$. Let $(\hat{\sigma}^j \times \hat{\sigma}_a^j \times \hat{\sigma}_b^j)_j$ be a convergence subsequence of $(\hat{\sigma}^k \times \hat{\sigma}_a^k \times \hat{\sigma}_b^k)_k$ in $\mathcal{P}(\mathcal{A})^3$, with limit $(\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b)$, then each of $\hat{\sigma}^j$, $\hat{\sigma}_a^j$ and $\hat{\sigma}_b^j$ converge weakly to $\hat{\sigma}$, $\hat{\sigma}_a$ and $\hat{\sigma}_b$ respectively and thus $(\hat{\sigma}^j \times \hat{\sigma}_a^j \times \hat{\sigma}_b^j)_j$ converges weakly to $\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b$ and thus it converges in the Prokhorov metric of $\mathcal{P}(\mathcal{A}^3)$. This implies that $(\hat{\sigma}^k \times \hat{\sigma}_a^k \times \hat{\sigma}_b^k)_k$ also converges weakly to $\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b$ and so μ is a product measure. Again since $\mu \notin S$, assume WLOG $\mathbb{E}_{(a,b) \sim \hat{\sigma}} [u_1(a, b)] < \mathbb{E}_{(a,b) \sim \hat{\sigma}_a} [u_1(a, b)]$. Define $f : \mathcal{A}^3 \rightarrow \mathbb{R}$, $f(a, b, c, d, e, f) = u_1(a, b) - u_1(c, d)$. f is continuous and from the no-policy regret of the pair $\hat{\sigma}^k, \hat{\sigma}_a^k$ we have:

$$0 \leq \lim_{k \rightarrow \infty} \mathbb{E}_{(a,b,c,d,e,f) \sim (\hat{\sigma}^k \times \hat{\sigma}_a^k \times \hat{\sigma}_b^k)_k} [f(a, b, c, d, e, f)] = \mathbb{E}_{(a,b,c,d,e,f) \sim \mu} [f(a, b, c, d, e, f)] < 0,$$

which is again a contradiction. \square

For completeness we restate the main result below and give its proof in full.

Theorem B.9. *If the algorithms played by player 1 in the form of f_t and player 2 in the form of g_t give sub-linear policy regret sequences, then sequence of product distributions $(\hat{\sigma}^T \times \hat{\sigma}_a^T \times \hat{\sigma}_b^T)_{T=1}^\infty$ converges weakly to the set S .*

Proof of Theorem 4.8. We consider the sequence of empirical distributions $\hat{\pi}^T$ defined in B.5, over the functional space $\mathcal{F}_1 \times \mathcal{F}_2$ and show that this sequence must converge to the set of all policy equilibria Π in the Prokhorov metric on $\mathcal{P}(\mathcal{F}_1 \times \mathcal{F}_2)$. First, notice that since the functions $f : \mathcal{A}_2 \rightarrow \mathcal{A}_1$ are from finite sets of actions to finite sets of actions, we can consider the set \mathcal{F}_1 as a subset of a finite dimensional vector space, with the underlying field of real numbers and the metric induced by the l_1 norm. Similarly, we can also equip \mathcal{F}_2 with the l_1 norm. Since both \mathcal{F}_1 and \mathcal{F}_2 are closed sets with respect to this metric and they are clearly bounded, they are compact. Thus the set $\mathcal{F}_1 \times \mathcal{F}_2$ is a compact set with the underlying metric d being the maximum metric. By Prokhorov's theorem we know that $\mathcal{P}(\mathcal{F}_1 \times \mathcal{F}_2)$ is a compact metric space with the Prokhorov metric. Suppose that the sequence $(\hat{\pi}^T)_T$ does not converge to Π . This implies that there is some convergent subsequence $(\hat{\pi}^t)_t$ with a limit π outside of Π . Let M be the Markov chain induced by π and let \hat{M}^T be the Markov chain induced by $\hat{\pi}^T$.

First we show that $\lim_{t \rightarrow \infty} \|\hat{M}^t - M\|_1 = 0$. Recall that $M_{(a,b),(a',b')} = \sum_{(f,g): f(b)=a', g(a)=b'} \pi(f, g)$ and that by lemma B.7 $\hat{M}_{(a,b),(a',b')}^t = \sum_{(f,g): f(b)=a', g(a)=b'} \hat{\pi}^t(f, g)$. Notice that f, g are continuous functions on \mathcal{F}_1 and \mathcal{F}_2 , since the topology induced by the l_1 metric on both sets is exactly the discrete topology and every function from a topological space equipped with the discrete topology is continuous. Since convergence in the Prokhorov metric implies weak convergence, we have that for any fixed f, g , $\lim_{t \rightarrow \infty} \hat{\pi}^t(f, g) = \pi(f, g)$. Since the sum $\sum_{(f,g): f(b)=a', g(a)=b'} \hat{\pi}^t(f, g)$ is finite this implies that $\lim_{t \rightarrow \infty} \sum_{(f,g): f(b)=a', g(a)=b'} \hat{\pi}^t(f, g) = \sum_{(f,g): f(b)=a', g(a)=b'} \pi(f, g)$ and so $\lim_{t \rightarrow \infty} \|\hat{M}^t - M\|_1 = 0$.

Next we show that any convergent subsequence $(\hat{\sigma}^k)_k$ of $(\hat{\sigma}^t)_t$ in the Prokhorov metric, converges to a stationary distribution σ of M . First notice that $(\hat{\sigma}^k)_k$ exists, since $\mathcal{P}(\mathcal{A})$ is compact. Next, suppose σ is the limit of $(\hat{\sigma}^k)_k$ in the Prokhorov metric. This implies that $\lim_{k \rightarrow \infty} \hat{\sigma}^k(a, b) = \sigma(a, b)$, in particular if we consider $\mathcal{A} \subset \mathbb{R}^{|\mathcal{A}|}$ and $\hat{\sigma}^k, \sigma \in \mathbb{R}^{|\mathcal{A}|}$ as vectors, then the above implies that $\lim_{k \rightarrow \infty} \|\sigma - \hat{\sigma}^k\|_1 = 0$. Next we construct the following sequence $(\sigma_{k_n})_{k_n}$ of stationary distributions of M – choose k_n large enough, so that $\|M - \hat{M}^{k_n}\| \leq \frac{1}{n}$. Such a k_n exists, because $(\hat{M}^k)_k$ is a subsequence of $(\hat{M}^t)_t$ which converges to M . By lemma B.3, there exists a stationary distribution σ_{k_n} of M such that $\|\hat{\sigma}_{k_n} - \sigma_{k_n}\|_1 \leq \frac{c}{n}$, for some constant c . We show that σ_{k_n} converges to σ . Fix some $\epsilon > 0$, we find an N , such that for any $n \geq N$ we have $\|\sigma_{k_n} - \sigma\|_1 < \epsilon$. Notice that $\|\sigma_{k_n} - \sigma\|_1 \leq \|\sigma_{k_n} - \hat{\sigma}_{k_n}\|_1 + \|\hat{\sigma}_{k_n} - \sigma\|_1$. Since $\|\sigma_{k_n} - \hat{\sigma}_{k_n}\|_1 \leq \frac{c}{n}$ and by convergence, we know that for $\frac{\epsilon}{2}$, there exists N' such that for any $n \geq N'$, $\|\hat{\sigma}_{k_n} - \sigma\|_1 < \frac{\epsilon}{2}$, we can set $N = \max(\frac{2}{c\epsilon}, N')$. Suppose, for a contradiction, that σ is not a stationary distribution of M . Then there exists some ϵ such that $\|\sigma^\top M - \sigma\|_2 > \epsilon$. This implies:

$$\epsilon < \|\sigma^\top M - \sigma\|_2 \leq \|\sigma^\top M - \sigma_{k_n}^\top M\|_2 + \|\sigma_{k_n}^\top M - \sigma\|_2 < 2\|\sigma_{k_n} - \sigma\|_2,$$

where the last inequality holds by the fact $\sigma - \sigma_{k_n}$ is not a stationary distribution of M and thus M can only shrink the difference as a stochastic matrix. The inequality $2\|\sigma_{k_n} - \sigma\|_2 > \epsilon$ is a contradiction since we know that σ_{k_n} converges to σ and thus σ is a stationary distribution of M . Since strong convergence, implies weak convergence, which in hand implies convergence in the Prokhorov metric for separable metric spaces, we have shown that every convergent subsequence of $(\hat{\sigma}_t)_t$ converges to a stationary distribution of M in the Prokhorov metric.

Next, we show that $(\hat{\pi}_a^t)_t$ converges to π_a . By assumption $(\hat{\pi}^t)_t$ converges weakly to π . Since we are in a finite dimensional space, we also have strong convergence. In particular, for any $g \in \mathcal{F}_2$, we have

$$\lim_{t \rightarrow \infty} \sum_{f \in \mathcal{F}_1} \hat{\pi}^t(f, g) = \sum_{f \in \mathcal{F}_1} \lim_{t \rightarrow \infty} \hat{\pi}^t(f, g) = \sum_{f \in \mathcal{F}_1} \pi(f, g)$$

and so the sequence of marginal distribution also converges to the respective marginal of π . Since $\hat{\pi}_a^t$ is exactly the product distribution of the dirac distribution over \mathcal{F}_1 putting all weight on the constant function mapping everything to the fixed action a and the marginal of $\hat{\pi}^t$ over \mathcal{F}_2 , by the convergence of marginals we conclude that $(\hat{\pi}_a^t)_t$ converges to π_a in the strong sense and thus in the Prokhorov metric. In the same way we can show that $(\hat{\pi}_b^t)_t$ converges to π_b .

With a similar argument as for $(\hat{\sigma}^t)_t$ we show that every convergent subsequence of $(\hat{\sigma}_a^t)_t$ converges to a stationary distribution of M_a and any convergent subsequence of $(\hat{\sigma}_b^t)_t$ converges to a stationary

distribution of M_b . Because of the construction in theorem B.1 and the convergence of $\widehat{\pi}_a$ to π_a , we can guarantee that $(\widehat{\sigma}_a^t)_t$ converges precisely to σ_a :

$$\begin{aligned}\sigma_a(a, b) &= \sum_{f, g: g(a)=b} \pi_a(f, g) = \sum_{f, g: g(a)=b} \lim_{t \rightarrow \infty} \widehat{\pi}_a^t(f, g) \\ &= \lim_{t \rightarrow \infty} \sum_{f, g: g(a)=b} \widehat{\pi}_a^t(f, g) = \lim_{t \rightarrow \infty} \widehat{\sigma}_a^t(a, b).\end{aligned}$$

Similarly $\widehat{\sigma}_b^t$ converges to $\widehat{\sigma}_b$. However, we assumed that π is not a policy equilibrium and thus no stationary distributions of M , M_a and M_b can satisfy the policy equilibrium inequalities. We now arrive at a contradiction since by theorem B.8 and the above, we have that any limit point of $(\widehat{\sigma}^t)_t$ and the corresponding distributions for fixed actions a and b are stationary distributions of M , M_a and M_b , respectively, which satisfy the policy equilibrium inequalities. \square

Not surprisingly, we are able to show that the empirical distributions of observed play, also converge to the set of S almost surely.

Corollary B.10. *The sequence of product distributions $(\tilde{\sigma} \times \tilde{\sigma}_a \times \tilde{\sigma}_b)_{T=1}^\infty$ converges weakly to the set S almost surely.*

Proof. We are going to show that for any convergent subsequence $(\widehat{\sigma}^t)_t$ of $(\widehat{\sigma})_T$ with limit point σ , the corresponding subsequence $(\tilde{\sigma}^t)_t$ converges to σ a.s.. Let E be the event that $\lim_{t \rightarrow \infty} \|\sigma - \tilde{\sigma}^t\|_1 = 0$. Consider the following event $\lim_{t \rightarrow \infty} \|\sigma - \widehat{\sigma}^t\|_1 + \lim_{t \rightarrow \infty} \|\widehat{\sigma}^t - \tilde{\sigma}^t\|_1 = 0$ denoted by E' . Notice that every time E' occurs E also occurs because

$$\|\sigma - \tilde{\sigma}^t\|_1 \leq \|\sigma - \widehat{\sigma}^t\|_1 + \|\widehat{\sigma}^t - \tilde{\sigma}^t\|_1,$$

and taking limits, preserves the non-strict inequality. By Theorem B.19, E' occurs with probability 1 and thus E occurs w.p. 1 as well. \square

B.2 Proof that CCEs are a subset of policy equilibria

Lemma B.11. *For any fixed game \mathcal{G} and a CCE σ of \mathcal{G} , there exist on average $(m, R(T))$ stable no-external regret algorithms, which if the players follow, the empirical distribution of play, converges to σ . Here $m < o(T)$ and $R(T) \leq O(\sqrt{T})$.*

Proof. Fix the time horizon to be T . We assume that at round t of the game, player i has access to an oracle which provides $\widehat{\sigma}_t$, the empirical distribution of play. The algorithm of player i is the following – split the time horizon into mini-batches of size \sqrt{T} . For the first mini-batch the player plays according to σ . At the end of the mini-batch, the oracle provides the player with $\widehat{\sigma}_t$ and the player checks if $\|\widehat{\sigma}_t - \sigma\|_1 \leq \frac{|\mathcal{A}|}{T^{1/6}}$. If the condition is satisfied, then the algorithm continues playing according to σ . At the j -th epoch, the player checks if $\|\widehat{\sigma}_t - \sigma\|_1 \leq \frac{|\mathcal{A}|}{j^{1/6} T^{1/6}}$. If at any of the epochs the inequality turns out to be false, then the player switches to playing Exp3 at the end of the next epoch for the rest of the game i.e. if at the end of the j -th epoch, the inequality does not hold, then at the end of the $j+1$ -st epoch, the player switches to playing Exp3. First we show that this is a no-external regret algorithm. Suppose until epoch j the player has not switched to Exp3 and then at epoch j the player switches. Let u_t denote the utility function which the player observed at time t . Let $(a_t)_{t=1}^T$ be the action sequence of the player. For any fixed action $a \in \mathcal{A}_i$ we have:

$$\begin{aligned}\mathbb{E} \left[\sum_{t=1}^T u_t(a) - \sum_{t=1}^T u_t(a_t) \right] &= \mathbb{E} \left[\sum_{t=1}^{(j-1)\sqrt{T}} u_t(a) - \sum_{t=1}^{(j-1)\sqrt{T}} u_t(a_t) \right] + \sum_{t=(j-1)\sqrt{T}+1}^{(j+1)\sqrt{T}} [u_t(a) - u(a_t)] \\ &\quad + \mathbb{E} \left[\sum_{t=(j+1)\sqrt{T}+1}^T u_t(a) - u(a_t) \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^{(j-1)\sqrt{T}} u_t(a) - \sum_{t=1}^{(j-1)\sqrt{T}} u_t(a_t) \right] + 2\sqrt{T} + c\sqrt{|\mathcal{A}_i| \log(|\mathcal{A}_i|) T}\end{aligned}$$

where in the first inequality we have used the fact that the utilities are in $[0, 1]$ and the bound on the regret of Exp3 where c is some constant. We now bound the term $\mathbb{E} \left[\sum_{t=1}^{(j-1)\sqrt{T}} u_t(a) - \sum_{t=1}^{(j-1)\sqrt{T}} u_t(a_t) \right]$ via two applications of lemma B.13.

$$\begin{aligned}
& \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^{(j-1)\sqrt{T}} u_t(a) - \sum_{t=1}^{(j-1)\sqrt{T}} u_t(a_t) \right] = \mathbb{E}_{(a', b') \sim \hat{\sigma}_t} [u_i(a, b')] - \mathbb{E}_{(a', b') \sim \hat{\sigma}_t} [u_i(a', b')] \\
& = \mathbb{E}_{(a', b') \sim \hat{\sigma}_t} [u_i(a, b')] - \mathbb{E}_{(a', b') \sim \sigma} [u_i(a, b')] + \mathbb{E}_{(a', b') \sim \sigma} [u_i(a, b')] - \mathbb{E}_{(a', b') \sim \sigma} [u_i(a', b')] \\
& + \mathbb{E}_{(a', b') \sim \sigma} [u_i(a', b')] - \mathbb{E}_{(a', b') \sim \hat{\sigma}_t} [u_i(a', b')] \\
& \leq \|P_i\| \frac{(j-1)|\mathcal{A}|}{jT^{1/6}} + \langle P_i, \mathbb{E}_{(a', b') \sim \sigma} [b' a'^T] \rangle - \langle P_i, \mathbb{E}_{(a', b') \sim \hat{\sigma}_t} [b' a'^T] \rangle \\
& \leq 2\|P_i\| \frac{|\mathcal{A}|}{T^{1/6}},
\end{aligned}$$

where the first inequality follows from lemma B.13 and the fact σ is a CCE. The second inequality again follows from lemma B.13. The above implies $\mathbb{E} \left[\sum_{t=1}^{(j-1)\sqrt{T}} u_t(a) - \sum_{t=1}^{(j-1)\sqrt{T}} u_t(a_t) \right] \leq \frac{\sqrt{T}|\mathcal{A}|^2}{T^{1/6}}$. This follows fact that the spectral norm of the utility matrix for player i can not exceed $|\mathcal{A}|$, which again follows from the boundedness of utilities.

Next, we show that if all players, play according to the algorithm, then with high probability, the empirical distribution of play converges to σ . For this, we analyze the probability of the event that $\|\hat{\sigma}_t - \sigma\|_1 > \frac{|\mathcal{A}|}{T^{1/6}}$ after the first epoch. By Azuma's inequality, we have that $\mathbb{P} \left[|\hat{\sigma}_t(x) - \sigma(x)| > \frac{\epsilon}{|\mathcal{A}|} \right] \leq 2 \exp(-\epsilon^2 t / (2|\mathcal{A}|))$, for any fixed action $x \in \mathcal{A}$. By a union bound over all possible actions, we get that $\mathbb{P} [\|\hat{\sigma}_1 - \sigma\|_1 > \epsilon] \leq 2|\mathcal{A}| \exp(-\epsilon^2 t / (2|\mathcal{A}|))$. Plugging in $\epsilon = \frac{|\mathcal{A}|}{T^{1/6}}$, we get that the probability the algorithm fails after the first epoch is bounded by $2|\mathcal{A}| \exp(-T^{1/6}|\mathcal{A}|/2)$. Now the probability that the algorithm fails after epoch j is bounded by

$$(1 - 2|\mathcal{A}| \exp(-T^{1/6}|\mathcal{A}|/2))^{j-1} 2|\mathcal{A}| \exp(-j^2 T^{1/6}|\mathcal{A}|/2) < 2|\mathcal{A}| \exp(-T^{1/6}|\mathcal{A}|/2).$$

A very pessimistic union bound now gives us that the probability the algorithm fails to converge to σ , provided that all players use it, is upper bounded by $2\sqrt{T}|\mathcal{A}| \exp(-T^{1/6}|\mathcal{A}|/2)$.

Using a doubling trick, we can achieve the same results for any T . As in theorem B.19, we can now get $\limsup_{T \rightarrow \infty} \|\hat{\sigma}_T - \sigma\|_1 = 0$ almost surely and thus the empirical distribution of play converges to σ a.s.. To check the stability of the algorithm we consider three possibilities – the player is playing according to σ and is not in an epoch after which they switch to Exp3, the player is playing according to σ and is in an epoch after which they switch to Exp3 and the player has switched to playing Exp3. We take the view of player 1. Consider the first case and suppose the current epoch of the algorithm is j . In particular consider time step $j\sqrt{T} + i$ at which the play was according to $\sigma_{j\sqrt{T}+i} = \sigma$ and suppose player 2 in the last $m < o(\sqrt{T})$ switched from playing distributions $(p_t)_{t=j\sqrt{T}+i-m}^{j\sqrt{T}}$ to $(\tilde{p}_t)_{t=j\sqrt{T}+i-m}^{j\sqrt{T}}$ so that player 1 plays $\tilde{\sigma}_{j\sqrt{T}+i}$. Since $m < o(\sqrt{T})$, the only change in the current algorithm can occur at the $j-1$ -st epoch and is in the form of the algorithm decides to switch to Exp3. Since the switch can only occur after the j -th epoch is done, this implies that at time $j\sqrt{T} + i$ again the play is according to the distribution σ and we have $\|\sigma_{j\sqrt{T}+i} - \tilde{\sigma}_{j\sqrt{T}+i}\|_1 = 0$. Consider the second case in which at the end of epoch j the algorithm switches to playing Exp3. If the past m actions of player 2 change, this can only change the decision at the end of $j-1$ -st epoch of player 1 to switch to Exp3. However, as in the previous case, the player at the j -th epoch is still playing according to σ and again the algorithm is stable at each iteration during the epoch. Finally, in the third case we have stability from the fact Exp3 is stable and once we enter an epoch during which Exp3 is played, changing the past m -actions of player 2 can not change the decision of player 1 to switch to Exp3, since $m < o(\sqrt{T})$ and the epochs are of size \sqrt{T} .

□

As a consequence of the proof we see that, the algorithm is on average stable for any $m < o(\sqrt{T})$.

Table 1: Utility for player 1

player 1 \ player 2	c_1	c_2
a	3/4	0
b	1	0

Theorems 4.9 and 4.10 are combined in the the following theorem.

Theorem B.12. *For any CCE σ of a 2-player game \mathcal{G} , there exists a policy-equilibrium π , which induces a Markov chain M , with stationary distribution σ . Further, there exists a 2-player game \mathcal{G} and product distributions $\sigma \times \sigma_a \times \sigma_b \in S$, where S is defined in 4.6, such that σ is not a CCE of \mathcal{G} .*

Proof. First note that lemma B.11 guarantees that for any CCE σ there exists an m -stable on average no-external regret algorithm, which generates an empirical distribution of play $\tilde{\sigma}$ converging to σ almost surely. From the m -stability of the algorithm and theorem A.2, we know that if both players play according to $\tilde{\sigma}$ then the empirical sequence of play also has no-policy regret and thus by theorem B.9, we know that $\tilde{\sigma} \times \tilde{\sigma}_a \times \tilde{\sigma}_b$ converges to the set of policy-regret equilibria.

To show the second part of the theorem, we consider a game in which the pay-off matrix for player 2 is just some constant multiple of $\mathbf{1}_{\mathcal{A}_1} \mathbf{1}_{\mathcal{A}_2}^\top$, where $\mathbf{1}_{\mathcal{A}_i} \in \mathbb{R}^{|\mathcal{A}_i|}$ is the vector with each entry equal to 1. This implies that whatever sequence player 2 chooses to play, they will have no-policy regret and no-external regret, since the observed utility in expectation is going to be equal for all actions $b \in \mathcal{A}_2$. Let $\mathcal{A}_1 = \{a, b\}$. The utility matrix of player 1 is given in 1. The strategy of player 2 is the following:

$$f_t(a_{t-1}, a_t) = \begin{cases} 1 & a_{t-1} = a, a_t = b \\ \frac{3}{4} & a_{t-1} = a_t = a \\ 0 & \text{otherwise} \end{cases}.$$

Similarly to the proof of theorem A.1, we can show that any no-policy regret strategy for player 1, incurs linear external regret. For any time step t , let $\hat{\sigma}_t^1$ be the dirac distribution for player 1, putting all probability mass on action a and let $\hat{\sigma}_t^2$ be the dirac distribution putting all probability mass on the action played by player 2 at time t . Let $\hat{\sigma} = \frac{1}{T} \sum_{t=1}^T \hat{\sigma}_t^1 (\hat{\sigma}_t^2)^\top$ be the empirical distribution after T rounds. Then by theorem B.9 $\hat{\sigma} \times \hat{\sigma}_a \times \hat{\sigma}_b$ converges to the set of policy equilibria. However, by construction, if player 1, plays according to $\hat{\sigma}$, then they will have linear external regret and thus the $\hat{\sigma}$ can not be converging to the set of CCEs. \square

B.3 Auxiliary results.

Lemma B.13. *Let σ and σ' be two distributions supported on a finite set and let f be a utility/loss function uniformly bounded by 1. If $\|\sigma - \sigma'\|_1 \leq \epsilon$ then $|\mathbb{E}_{a \sim \sigma} [f(a)] - \mathbb{E}_{a \sim \sigma'} [f(a)]| \leq \epsilon$.*

Proof.

$$\begin{aligned} |\mathbb{E}_{s \sim \sigma} [f(s)] - \mathbb{E}_{s \sim \sigma'} [f(s)]| &= \left| \sum_{s \in S} \sigma(s) f(s) - \sum_{s \in S} \sigma'(s) f(s) \right| \\ &= \left| \sum_{s \in S} f(s) (\sigma(s) - \sigma'(s)) \right| \leq \sum_{s \in S} |\sigma(s) - \sigma'(s)| = \|\sigma - \sigma'\|_1 \leq \epsilon. \end{aligned}$$

\square

Lemma B.14. *Let $M \in \mathbb{R}^{d \times d}$ and $\hat{M} \in \mathbb{R}^{d \times d}$ be two row-stochastic matrices, such that $\|M - \hat{M}\| \leq \epsilon$, then for any stationary distribution $\hat{\sigma}$ of \hat{M} , there exists a stationary distribution σ of M , such that $\|\hat{\sigma} - \sigma\|_1 \leq \frac{4d^2\epsilon}{\delta}$.*

Proof. Let $U \in \mathbb{R}^{d \times k}$ be the left singular vectors corresponding to the singular value 1 of M and let $\hat{U} \in \mathbb{R}^{d \times l}$ be the left singular vectors corresponding to the singular value 1 of \hat{M} . First notice that

$$\|MM^\top - \hat{M}\hat{M}^\top\| \leq \|MM^\top - M\hat{M}^\top\| + \|M\hat{M}^\top - \hat{M}\hat{M}^\top\| \leq (\|M\| + \|\hat{M}\|)\|M - \hat{M}\| \leq 2\epsilon$$

Denote the eigen-gap of M by δ , then by Wedin's theorem (see for example lemma B.3 in Allen-Zhu and Li (2016)) theorem we have

$$\|\widehat{U}^\top U^\perp\| \leq \frac{\|MM^\top - \widehat{M}\widehat{M}^\top\|}{\delta} \leq \frac{2\epsilon}{\delta}.$$

WLOG assume $\widehat{\sigma} = \frac{(\widehat{U})_i}{\|(\widehat{U})_i\|_1}$. This implies that $\|\widehat{\sigma}^\top U^\perp\|_2 \leq \frac{2d\epsilon}{\delta}$ and thus:

$$\|UU^\top \widehat{\sigma} - \widehat{\sigma}\|_2 = \|(I - UU^\top)\widehat{\sigma}\|_2 = \|U^\perp (U^\perp)^\top \widehat{\sigma}\|_2 \leq \|\widehat{\sigma}^\top U^\perp\|_2 \leq \frac{2d\epsilon}{\delta}.$$

Let $\sigma_i = \frac{U_i}{\|U_i\|_1}$ be the stationary distribution of M , corresponding to the i -th left singular vector and let $\alpha_i = (U^\top \widehat{\sigma})_i \|U_i\|_1 \geq 0$. Then we have $\|\sum_i \alpha_i \sigma_i - \widehat{\sigma}\|_1 \leq \frac{2d^2\epsilon}{\delta}$, where the inequality follows from the derivation above and the inequality between l_1 and l_2 norms. Let $\sigma = \frac{\sum_i \alpha_i \sigma_i}{\|\sum_i \alpha_i \sigma_i\|_1}$. This is a stationary distribution of M , since

$$\sigma^\top M = \frac{1}{\|\sum_i \alpha_i \sigma_i\|_1} \sum_i \alpha_i \sigma_i^\top M = \frac{\sum_i \alpha_i \sigma_i}{\|\sum_i \alpha_i \sigma_i\|_1} = \sigma.$$

Notice that by reverse triangle inequality we have

$$\left| \left\| \sum_i \alpha_i \sigma_i \right\|_1 - \|\widehat{\sigma}\|_1 \right| \leq \frac{2d^2\epsilon}{\delta},$$

or equivalently

$$\left| \left\| \sum_i \alpha_i \sigma_i \right\|_1 - 1 \right| \leq \frac{2d^2\epsilon}{\delta}.$$

Thus we have:

$$\|\sigma - \widehat{\sigma}\|_1 \leq \left\| \sum_i \alpha_i \sigma_i - \widehat{\sigma} \right\|_1 + \left\| \sigma - \sum_i \alpha_i \sigma_i \right\|_1 \leq \frac{2d^2\epsilon}{\delta} + \|\sigma\|_1 |1 - \left\| \sum_i \alpha_i \sigma_i \right\|_1| \leq \frac{4d^2\epsilon}{\delta}.$$

□

Corollary B.15. *Let the empirical distribution of observed play be $\tilde{\sigma}^T = \frac{1}{T} \sum_{t=1}^T \delta_t$, the empirical distribution of play if player 1 deviated to playing fixed action $a \in \mathcal{A}_1$ be $\tilde{\sigma}_a^T$ and the empirical distribution of play if player 2 to action $b \in \mathcal{A}_2$ be $\tilde{\sigma}_b^T$. The sequence $(\tilde{\sigma}^T, \tilde{\sigma}_a^T, \tilde{\sigma}_b^T)_T$ converges to the set P almost surely.*

Proof. Lemma B.13, together with theorem B.19 imply that

$$\limsup_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)]| = 0$$

almost surely i.e.

$$\mathbb{P} \left[\limsup_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)]| = 0 \right] = 1.$$

Since

$$\limsup_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)]| \geq \liminf_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)]| \geq 0,$$

this implies that

$$\mathbb{P} \left[\liminf_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)]| = 0 \right] = 1.$$

On the other hand this implies

$$\mathbb{P} \left[\liminf_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)]| = 0 \cap \limsup_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)]| = 0 \right] = 1.$$

and so

$$\mathbb{P} \left[\lim_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}^T} [u_1(a,b)]| = 0 \right] = 1.$$

In a similar way we can get $\lim_{T \rightarrow \infty} |\mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^T} [u_1(a,b)]| = 0$ a.s. The above imply that $\lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}^T} [u_1(a,b)] = 0$ a.s. and $\lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^T} [u_1(a,b)] = 0$ a.s. and thus:

$$\begin{aligned} 0 &= \lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}^T} [u_1(a,b)] + \lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}_a^T} [u_1(a,b)] \\ &= \lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] + \lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}^T} [u_1(a,b)] \end{aligned}$$

a.s.. Since $\mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}^T} [u_1(a,b)] < o(1)$, this implies that

$$\begin{aligned} 0 &\leq - \lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \hat{\sigma}^T} [u_1(a,b)] \\ &= \lim_{T \rightarrow \infty} \mathbb{E}_{(a,b) \sim \tilde{\sigma}^T} [u_1(a,b)] - \mathbb{E}_{(a,b) \sim \tilde{\sigma}_a^T} [u_1(a,b)] \end{aligned}$$

a.s.. Now we can proceed as in the proof of theorem B.8. \square

B.4 Concentration of \tilde{M}

Lemma B.16. *With probability at least $1 - |\mathcal{A}|6 \exp\left(-\frac{T\epsilon^2}{4}\right)$ it holds that $|\frac{1}{T} \sum_{t=1}^T p_{t-1}(x_i)p_t(x_j) - \frac{1}{T} \sum_{t=1}^T \delta_{t-1}(x_i)\delta_t(x_j)| < \epsilon$ and $|\frac{1}{T} \sum_{t=1}^T p_t(x_i) - \frac{1}{T} \sum_{t=1}^T \delta_t(x_i)| < \epsilon$, simultaneously for all i .*

Proof. We consider the random variable $Z_t = \delta_t(x_i) - p_t(x_i)$, notice that $\mathbb{E}[Z_t | p_1, \dots, p_{t-1}] = 0$ so that $\{Z_t\}_t$ is a bounded martingale sequence with $|Z_t| < 1$ and thus by Azuma's inequality we have $\mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^T Z_t\right| \geq \epsilon\right] < 2 \exp\left(-\frac{T\epsilon^2}{2}\right)$ which shows that $\frac{1}{T} \sum_{t=1}^T \delta_t(x_i)$ concentrates around $\frac{1}{T} \sum_{t=1}^T p_t(x_i)$. Let $R_t = \delta_{t-1}(x_i)\delta_t(x_j) - p_{t-1}(x_i)p_t(x_j)$ and consider the filtration $\{\mathcal{F}_t\}_t$, where $\mathcal{F}_1 = \emptyset$, $\mathcal{F}_t = \Sigma(\delta_1, \dots, \delta_t)$ is the sigma algebra generated by the random variables δ_1 to δ_t . Then $|R_t| \leq 1$ and $\mathbb{E}[R_t | \mathcal{F}_1, \dots, \mathcal{F}_{t-2}] = 0$, so $\{R_t\}_t$ is also a bounded martingale difference and thus $\mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^T R_{2t}\right| \geq \frac{\epsilon}{2}\right] < 2 \exp\left(-\frac{T\epsilon^2}{4}\right)$. A similar argument allows us to bound the sum of the R_{2t+1} 's and a union bound gives us $\mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^T R_t\right| \geq \epsilon\right] < 4 \exp\left(-\frac{T\epsilon^2}{4}\right)$. A union bound over all i finishes the proof. \square

Definition B.17. Define the perturbed distribution of player i at time t to be $\tilde{p}_t^i = (1 - \sqrt{|\mathcal{A}|\tilde{\epsilon}})p_t^i + \frac{1}{|\mathcal{A}|} \sqrt{|\mathcal{A}|\tilde{\epsilon}}$.

Lemma B.18. *The difference of expected utilities from playing according to $(\tilde{p}_t^i)_{t=1}^T$ instead of $(p_t^i)_{t=1}^T$ is at most $2T\sqrt{|\mathcal{A}|\tilde{\epsilon}}$*

Proof. From lemma B.13 at each time step the difference of expected utility is bounded by $\sqrt{|\mathcal{A}|\tilde{\epsilon}}$ in absolute value. \square

Theorem B.19. *If at time t player i plays according to \tilde{p}_t^i , where $\tilde{\epsilon} = \frac{T^{-1/4}}{|\mathcal{A}|}$ and $p_t = \tilde{p}_t^1(\tilde{p}_t^2)^\top$, then the regret for playing according to \tilde{p}_t^i is at most $O(T^{7/8})$. Further $\limsup_{T \rightarrow \infty} \|\tilde{M} - \hat{M}\|_2 = 0$, almost surely. Additionally if $\tilde{\sigma} = \frac{1}{T} \sum_{t=1}^T \delta_t$ is the stationary distribution of \tilde{M} corresponding to the observed play and $\hat{\sigma} = \frac{1}{T} \sum_{t=1}^T p_t$ is the stationary distribution of \hat{M} corresponding to the averaged empirical distribution, then $\limsup_{T \rightarrow \infty} \|\tilde{\sigma} - \hat{\sigma}\|_1 = 0$ almost surely.*

Proof. Set $\tilde{\epsilon} = \frac{T^{-1/4}}{|\mathcal{A}|}$. The regret bound of the no-external regret algorithms now becomes $O(T^{7/8})$. We can, however, now guarantee that $\sum_{t=1}^T p_t(x_i) \geq \frac{T^{-1/4}}{|\mathcal{A}|}$ and thus, combining this with the high probability bound we obtain that with probability at least $1 - C \exp\left(-\frac{T\epsilon^2}{4}\right)$ it holds that

$|\tilde{M}_{i,j} - \widehat{M}_{i,j}| < 2\epsilon \frac{T^{1/4}}{|\mathcal{A}|}$. To see this, let $x = \frac{1}{T} \sum_{t=1}^T p_{t-1}(x_i)p_t(x_j)$, $\widehat{x} = \frac{1}{T} \sum_{t=1}^T \delta_{t-1}(x_i)\delta_t(x_j)$, $y = \frac{1}{T} \sum_{t=1}^T p_t(x_i)$, $\widehat{y} = \frac{1}{T} \sum_{t=1}^T \delta_t(x_i)$. Then

$$|\tilde{M}_{i,j} - \widehat{M}_{i,j}| = \left| \frac{x}{y} - \frac{\widehat{x}}{\widehat{y}} \right| \leq \frac{|x - \widehat{x}|}{|y|} + \frac{|\widehat{x}|}{|1/y - 1/\widehat{y}|} \leq \frac{\epsilon}{\epsilon} + \frac{|\widehat{x}||y - \widehat{y}|}{|y\widehat{y}|} \leq 2\frac{\epsilon}{\epsilon},$$

where the last inequality holds because $\widehat{x} \leq \widehat{y}$. Setting $\epsilon = T^{-1/3}$ and a union bound we arrive at $\mathbb{P}[\|\tilde{M} - \widehat{M}\|_2 > T^{-1/12}] < C \exp(-\frac{T^{1/3}}{4})$. By Borel-Cantelli lemma we have $\limsup_{T \rightarrow \infty} \|\tilde{M} - \widehat{M}\|_2 = 0$ almost surely. From lemma B.16 and a union bound we know that with $\mathbb{P}[\|\tilde{\sigma} - \widehat{\sigma}\|_1 > \epsilon] < 2|\mathcal{A}| \exp(-\frac{T\epsilon^2}{4})$. Setting $\epsilon = T^{-1/3}$ and again using Borel-Cantelli's lemma we see that $\limsup_{T \rightarrow \infty} \|\tilde{\sigma} - \widehat{\sigma}\|_1 = 0$. \square

C Extending the framework to arbitrary memories

Definition C.1. Let player 1 have memory m_1 and player 2 have memory m_2 . Let the function spaces be $\mathcal{F}_1 = \{f : \mathcal{A}_2^{m_1} \rightarrow \mathcal{A}_1\}$ and $\mathcal{F}_2 = \{f : \mathcal{A}_1^{m_2} \rightarrow \mathcal{A}_2\}$. Let π be a distribution over $\mathcal{F}_1 \times \mathcal{F}_2$. Define the m -memory bounded Markov process, where $m = \max(m_1, m_2)$ to be $\mathbb{P}[(a_t, b_t)|(a_{t-1}, b_{t-1}), \dots, (a_{t-m}, b_{t-m})] = \sum_{(f,g) \in \mathcal{F}_1 \times \mathcal{F}_2 : f(b_{t-1}, \dots, b_{t-m}) = a_t, g(a_{t-1}, \dots, a_{t-m}) = b_t} \pi(f, g)$. We associate with this Markov process the matrix $M \in \mathbb{R}^{|\mathcal{A}|^m \times |\mathcal{A}|^m}$, with $M_{(x_{t-1}, \dots, x_{t-m}), (x_t, x_{t-1}, \dots, x_{t-m+1})} = \mathbb{P}[x_t | x_{t-1}, \dots, x_{t-m}]$ and $M_{i,j} = 0$ for all other entries.

Definition C.2. The utility of π is defined through the stationary distribution γ of M . In particular γ is a distribution over $|\mathcal{A}|^m$ with entries indexed by (x_t, \dots, x_{t-m+1}) . Let σ be the marginal distribution of x_t , then $u_i(\pi) = \sup_{\sigma} \mathbb{E}_{(a,b) \sim \sigma} [u_i(a, b)]$

Definition C.3. The empirical m -memory bounded Markov process is \widehat{M} with $\mathbb{P}[x_1 | x_2, \dots, x_{m+1}] = \frac{\sum_t \prod_{i=0}^{m-1} p_{t-i}(x_{i+1}) \times \prod_{i=0}^{m-1} p_{t-i}(x_{i+2})}{\sum_t \prod_{i=0}^{m-1} p_{t-i}(x_{i+2})}$

Theorem C.4. The distribution $\widehat{\sigma}$ over $|\mathcal{A}|^m$ with entries indexed as

$$\widehat{\sigma}((a_1, b_1), \dots, (a_m, b_m)) = \frac{1}{T} \sum_t \prod_{i=0}^{m-1} p_{t-i}((a_{i+1}, b_{i+1}))$$

, is a stationary distribution of \widehat{M} .

Proof. Consider $(\widehat{\sigma}^\top \widehat{M})_{x_1, \dots, x_m}$, we show it is equal to $\widehat{\sigma}(x_1, \dots, x_m)$:

$$\begin{aligned} (\widehat{\sigma}^\top \widehat{M})_{x_1, \dots, x_m} &= \sum_{x_2, \dots, x_{m+1}} \widehat{M}_{(x_2, \dots, x_{m+1}), (x_1, \dots, x_m)} \widehat{\sigma}(x_2, \dots, x_{m+1}) \\ &= \sum_{x_2, \dots, x_{m+1}} \mathbb{P}[x_1 | x_2, \dots, x_{m+1}] \widehat{\sigma}(x_2, \dots, x_{m+1}) \\ &= \frac{1}{T} \sum_{x_2, \dots, x_{m+1}} \frac{\sum_t \prod_{i=0}^{m-1} p_{t-i}(x_{i+1}) \times \prod_{i=0}^{m-1} p_{t-i}(x_{i+2})}{\sum_t \prod_{i=0}^{m-1} p_{t-i}(x_{i+2})} \sum_t \prod_{i=0}^{m-1} p_{t-i}(x_{m-i+2}) \\ &= \frac{1}{T} \sum_{x_2, \dots, x_{m+1}} \sum_t \prod_{i=0}^{m-1} p_{t-i}(x_{i+1}) \times \prod_{i=0}^{m-1} p_{t-i}(x_{i+2}) \\ &= \frac{1}{T} \sum_t \prod_{i=0}^{m-1} p_{t-i}(x_{i+1}) = \widehat{\sigma}(x_1, \dots, x_m). \end{aligned}$$

\square

We also need a theorem which states that the utility $\mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a, \cdot, a))]$ is the expectation of the stationary distribution of the Markov process coming from the play $(a, g(a, \cdot, a))$ according to the marginal of g .

Definition C.5. Let M_1^a be the m_2 -memory bounded Markov process which comes from player 1 playing a fixed action $a \in \mathcal{A}_1$ and player 2 playing $g \in \mathcal{F}_2$ according to the marginal of π , i.e. $(M_1^a)_{(x_1, \dots, x_{m_2}), (x_1, \dots, x_{m_2+1})} = \mathbb{P}[x_{m_2+1} | x_1, \dots, x_{m_2}] = \sum_{(f,g): g(a_1, \dots, a_{m_2}) = b_{m_2+1}} \pi(f, g)$ if $a_{m_2+1} = a$ or 0 otherwise (here $x_i = (a_i, b_i)$). Similarly let M_2^b be the m_1 -memory bounded Markov process which arises when player 2 switches to playing the constant action.

Theorem C.6. Consider M_1^a and let $m = m_2$. Let $\bar{\sigma}$ be the following distribution over \mathcal{A}^m – $\bar{\sigma}(x_1, \dots, x_m) = \sum_{(f,g): g(a, \dots, a) = b} \pi(f, g)$ if $x_1 = \dots = x_m = (a, b)$ and 0 otherwise. Then $\bar{\sigma}$ is a stationary distribution of M_1^a and its marginal distribution $\gamma(a, b) = \sum_{(f,g): g(a, \dots, a) = b} \pi(f, g)$ is such that $\mathbb{E}_{(a,b) \sim \gamma} [u_1(a, b)] = \mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a, \dots, a))]$.

Proof.

$$\begin{aligned} (\bar{\sigma}^\top M_1^a)_{x_2, \dots, x_{m+1}} &= \sum_{x_1} \bar{\sigma}(x_1, \dots, x_m) (M_1^a)_{(x_1, \dots, x_m), (x_2, \dots, x_{m+1})} \\ &= \sum_{x_1} \bar{\sigma}(x_1, \dots, x_m) \mathbb{P}[x_{m+1} | x_1, \dots, x_m] \\ &= \sum_{b_1 \in \mathcal{A}_2} \sigma((a, b_1), \dots, (a, b_1)) \sum_{f, g(a, \dots, a) = b_{m+1}} \pi(f, g) \\ &= \sum_{b_1 \in \mathcal{A}_2} \left[\sum_{f, g(a, \dots, a) = b_1} \pi(f, g) \right] \sum_{f, g(a, \dots, a) = b_{m+1}} \pi(f, g) \\ &= \sum_{f, g(a, \dots, a) = b_{m+1}} \pi(f, g) = \sigma((a, b_{m+1}), \dots, (a, b_{m+1})) = \sigma(x_2, \dots, x_{m+1}) \end{aligned}$$

□

The rest of the proofs extend in similar ways to the case with general memory. There is the question of not requiring $\bar{\sigma}$ to be a stationary distribution but to be any distribution such that the marginal with respect to the first coordinate of $\bar{\sigma}^\top M_1^a$, satisfies the expectation equality $\mathbb{E}_{(a,b) \sim \gamma} [u_1(a, b)] = \mathbb{E}_{(f,g) \sim \pi} [u_1(a, g(a, \dots, a))]$.

D Extending the framework to arbitrary number of players

Definition D.1. Consider an n player game where player i has memory m_i . Define the set of functions $\mathcal{F}_i = \{f : \mathcal{A}_{-i}^{m_i} \rightarrow \mathcal{A}_i\}$. We consider a distribution π over $\mathcal{F} = \times_{i=1}^n \mathcal{F}_i$. Let $m = \max(m_1, \dots, m_n)$. Let $x = (a^1, \dots, a^n) \in \mathcal{A}$ and let $x^{-i} = (a^1, \dots, a^{i-1}, a^{i+1}, \dots, a^n) \in \mathcal{A}_{-i}$. Define the m -memory Markov process M such that $\mathbb{P}[x_{m+1} | x_1, \dots, x_m] = \sum_{(f_1(x_m^{-1}, \dots, x_{m-m_1+1}^{-1}) = a_{m+1}^1, \dots, f_n(x_m^{-n}, \dots, x_{m-m_n+1}^{-n}) = a_{m+1}^n)} \pi(f_1, \dots, f_n)$.

All other definitions follow the same form. The problem with this most general setting is – how do we construct no-policy regret algorithms and what does it even mean to have no-policy regret? The utility function of player i at time $t - u_i(\cdot, x_t^{-i})$ is no longer interpretable as an m -memory bounded function, since x_t^{-i} depends on all other players' memories. One possible solution is to look at this utility as an m -memory bounded function, where m is the maximum memory among all other players.

E Simple example of a policy equilibrium

We now present a simple 2-player game with strategies of the players which lead to a policy equilibrium, which in fact is not a CCE. Further these strategies give the asymptotically maximum utility for both row and column players over repeated play of the game. The idea behind the construction is very similar to the one showing incompatibility of policy regret and external regret. The utility matrix for the game is given in Table 2. Since the column player has the same payoff for all his actions they will always have no policy and no external regret. The strategy the column player chooses is to always play the function $f : \mathcal{A}_1 \rightarrow \mathcal{A}_2$:

$$f(x) = \begin{cases} c & x = a \\ d & x = b. \end{cases}$$

Table 2: Utility matrix

Player 1 \ Player 2	c	d
a	(3/4, 1)	(0, 1)
b	(1, 1)	(0, 1)

In the view of the row player, this strategy corresponds to playing against an adversary which plays the familiar utility functions:

$$u_t(a_{t-1}, a_t) = \begin{cases} 1 & a_{t-1} = a, a_t = b \\ \frac{3}{4} & a_{t-1} = a_t = a \\ 0 & \text{otherwise.} \end{cases}$$

We have already observed that on these utilities, the row player can have either no policy regret or no external regret but not both. What is more the utility of no policy regret play is higher than the utility of any of the no external regret strategies. This already implies that the row player is better off playing according to the no policy regret strategy which consists of always playing the fixed function $g : \mathcal{A}_2 \rightarrow \mathcal{A}_1$ given by $g(x) = a$. Below we present the policy equilibrium $\pi \in \Delta \mathcal{F}$, corresponding Markov chain $M \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{A}|}$ and its stationary distribution $\sigma \in \Delta \mathcal{A}$ satisfying the no policy regret requirement.

$$\pi(\tilde{f}, \tilde{g}) = \delta_{(f, g)}, M = \begin{matrix} & \begin{matrix} (a, c) & (a, d) & (b, c) & (b, d) \end{matrix} \\ \begin{matrix} (a, c) \\ (a, d) \\ (b, c) \\ (b, d) \end{matrix} & \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \end{matrix}, \sigma(x, y) = \delta_{(a, c)}.$$

Suppose the row player was playing any no policy regret strategy, for example one coming from a no policy regret algorithm, as a response to the observed utilities $u_t(\cdot, \cdot)$. Since the only sublinear policy regret play for these utilities is to only deviate from playing a a sublinear number of times we see that the empirical distribution of play for the row player converges to the dirac distribution δ_a . Together with the strategy of the column player, this implies the column player chooses the action d only a sublinear number of times and thus their empirical distribution of play converges to δ_c . It now follows that the empirical distribution of play converges to $\delta_a \times \delta_c = \delta_{(a, c)} \in \Delta \mathcal{A}$. We can similarly verify that the empirical Markov chain will converge to M and the empirical functional distribution $\hat{\pi}$ converges to π . Theorem B.9 guarantees that because both players incur only sublinear regret π is a policy equilibrium. It should also be intuitively clear why this is the case without the theorem – suppose that the row player switches to playing the fixed action b . The resulting functional distribution, Markov chain and stationary distributions become:

$$\pi_b(\tilde{f}, \tilde{g}) = \delta_{(f, g=b)}, M_b = \begin{matrix} & \begin{matrix} (a, c) & (a, d) & (b, c) & (b, d) \end{matrix} \\ \begin{matrix} (a, c) \\ (a, d) \\ (b, c) \\ (b, d) \end{matrix} & \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}, \sigma_b(x, y) = \delta_{(b, d)}.$$

The resulting utility for the row player is now 0, compared to the utility gained from playing according to π , which is 3/4.