

The complexity of approximate (coarse) correlated equilibrium for incomplete information games

Binghui Peng

Columbia University, United States

BP2601@COLUMBIA.EDU

Aviad Rubinstein

Stanford University, United States

AVIAD@CS.STANFORD.EDU

Editors: Shipra Agrawal and Aaron Roth

Abstract

We study the iteration complexity of decentralized learning of approximate correlated equilibria in incomplete information games.

On the negative side, we prove that in *extensive-form games*, assuming $\text{PPAD} \not\subseteq \text{TIME}(n^{\text{polylog}(n)})$, any polynomial-time learning algorithms must take at least $2^{\log_2^{1-o(1)}(|\mathcal{I}|)}$ iterations to converge to the set of ϵ -approximate correlated equilibrium, where $|\mathcal{I}|$ is the number of nodes in the game and $\epsilon > 0$ is an absolute constant. This nearly matches, up to the $o(1)$ term, the algorithms of (Peng and Rubinstein, 2024; Dagan et al., 2024) for learning ϵ -approximate correlated equilibrium, and resolves an open question of Anagnostides, Kalavasis, Sandholm, and Zampetakis (Anagnostides et al., 2024). Our lower bound holds even for the easier solution concept of ϵ -approximate *coarse* correlated equilibrium

On the positive side, we give uncoupled dynamics that reach ϵ -approximate correlated equilibria of a *Bayesian game* in polylogarithmic iterations, without any dependence of the number of types. This demonstrates a separation between Bayesian games and extensive-form games.

Keywords: Online learning, regret minimization, game theory, computational complexity

1. Introduction

Our work considers fundamental questions on modeling and designing agents who learn to interact in settings with incomplete information. From an online learning perspective, we want agents that achieve low *external regret*¹, ideally also satisfying the stronger notion of low *swap regret*. From a game theory perspective, we consider agents that learn to play in games of incomplete information; we ask whether they can efficiently converge to the set of approximate *coarse correlated* (CCE) equilibria, ideally also satisfying the stronger notion of approximate *correlated equilibria* (CE). The connection between online learning and equilibrium computation is well-known (e.g. see Blum and Mansour (2007)): If each agent runs an online learning algorithm with low external regret, then the empirical distribution converges to the set of approximate CCE; if it further satisfies low swap regret, then the distribution converges to the set of approximate CE.

Correlated equilibria (specifically, normal-form correlated equilibria, NFCE²) have long been conjectured intractable Von Stengel and Forges (2008), until recently Dagan et al. (2024); Peng

1. Definitions of external and swap regret, Bayesian and extensive-form games, and various notions of equilibria are deferred to Section 2.

2. Different variants of CCE and CE have been considered for games of incomplete information. In what is arguably the most natural and desired variant, normal-form CE (NFCE), the correlation can be implemented by a correlating device sending each player a signal ex-ante, i.e. one signal that is independent of their type (in Bayesian game) or node in the game tree (in an extensive-form game). Motivated primarily by the conjecture that NFCE are intractable, several other variants have been considered Von Stengel and Forges (2008); Fujii (2023).

and Rubinstein (2024) gave algorithms that, for both Bayesian games and more general extensive-form games, compute ϵ -NFCE in $(|\mathcal{I}| \log(n))^{O(1/\epsilon)}$ iterations of uncoupled dynamics, where $|\mathcal{I}|$ is the size of the game (number of types or game nodes) and n is the number of actions. Put in the language of online learning, agents can achieve ϵ -swap regret after $(|\mathcal{I}| \log(n))^{O(1/\epsilon)}$ iterations. While this is polynomial in the full description of the game, the dependence on $|\mathcal{I}|$ (which is often combinatorial in natural game parameters) is undesirable.

A natural question is whether the algorithms of Dagan et al. (2024); Peng and Rubinstein (2024) can be improved. Our results show that the answer depends on whether we have a Bayesian game or an extensive-form game:

A lower bound on (poly-time) external regret in extensive-form games Recent work of Anagnostides et al. (2024) showed that even with only three players, and even if one settles for the weaker desiderata of low *external* regret or approximate *coarse* correlated equilibrium, they cannot be achieved in extensive-form games by polynomial-time agents running for fewer than $2^{\log_2^{1/2-o(1)}(|\mathcal{I}|)}$ iterations, assuming the ‘‘ETH for PPAD’’ Babichenko et al. (2016). In this work we give an improved reduction that:

- Raises the lower bound on the number of iterations to $2^{\log_2^{1-o(1)}(|\mathcal{I}|)}$; notice that this is almost tight as $2^{\log_2^{1+o(1)}(|\mathcal{I}|)}$ is already sufficient by Dagan et al. (2024); Peng and Rubinstein (2024).
- Weakens the complexity assumption from $\text{PPAD} \not\subseteq \text{TIME}(2^{\tilde{\Omega}(n)})$ to $\text{PPAD} \not\subseteq \text{TIME}(n^{\text{polylog}(n)})$; this lower bound is almost tight since if $\text{PPAD} = \text{FP}$ players can simply compute and play a Nash equilibrium.

Theorem 1 *Let $\epsilon > 0$ be an absolute constant. Assuming $\text{PPAD} \not\subseteq \text{TIME}(n^{\text{polylog}(n)})$, there is no polynomial time learning dynamics that guarantee to converge to the set of ϵ -CCE of a three player extensive-form game in $2^{\log_2^{1-o(1)}(|\mathcal{I}|)}$ iterations.*

This resolves an open question of Anagnostides et al. (2024) on the iteration complexity of CCE in extensive-form games, up to the $o(1)$ term in the exponential.

An algorithm for swap regret minimization in Bayesian games Our second result shows that for Bayesian games, there is a much better learning dynamics: it runs in time polynomial in the game description, and achieves ϵ -swap regret (equivalently, ϵ -CE), using only $(\log(n)/\epsilon^2)^{O(1/\epsilon)}$ iterations; in particular there is no dependence whatsoever on the number of types!

Theorem 2 *Let $\epsilon > 0$ and n be the number of actions. There is an uncoupled dynamic that converges to the set of ϵ -CE of a Bayesian game in $(\log(n)/\epsilon^2)^{O(1/\epsilon)}$ iterations.*

This separates the complexity of computing ϵ -CE in Bayesian and extensive-form games.

1.1. A brief overview of our approach

1.1.1. LOWER BOUND FOR EXTENSIVE-FORM GAME

We first review the previous lower bound of Anagnostides et al. (2024), which is based on a framework of Foster et al. (2023). Then we give a high level overview of our new ideas.

Previous approach Recent work of [Foster et al. \(2023\)](#) gives a new framework for lower bounding the number of iterations for obtaining approximate CCE; [Foster et al. \(2023\)](#) originally used this idea to Markov games, and [Anagnostides et al. \(2024\)](#) very recently applied it to extensive form games. These reductions -as well as ours- are based on three key ideas:

Low rank CCE We say that a correlated distribution is *low rank* (called “sparse” in previous work³) if it can be written as a mixture of a small number of product distributions. Uncoupled learning dynamics where agents independently update their strategies based on the history of play induce correlated distributions whose rank is at most the number of iterations. The reductions prove hardness results for the problem of computing low-rank approximate CCE; as a corollary, computationally efficient uncoupled dynamics cannot approach approximate CCE / low external regret.

Repeated game At a high level, the reductions in [Foster et al. \(2023\)](#); [Anagnostides et al. \(2024\)](#) begin with a normal-form game where (approximate) Nash equilibrium is hard to find. The reductions construct an extensive form (or Markov) game which sequentially repeats the hard normal-form game H times. Even though it is easy to find an (approximate) CCE in the normal-form game, recall that it is crucial that each player only observes their signal. However if the CCE is low rank, the players can learn each other’s signals from their strategies in previous repetitions, breaking the correlation.

Kibitzer player The downside of sequentially repeating the normal-form games is that it brings spurious equilibria which don’t exist in the single-shot game; in particular, equilibria where players may be incentivized to play low-payoff actions with threats of retaliations in future repetitions of the game. To circumvent these artifacts of extensive-form games, we would like to re-normalize the payoff of each player in each repetition with respect to other players’ strategies, then players are incentivized to play myopically because any (dis)advantage they could obtain in future repetitions will be normalized away anyway. To correctly compute this normalization, [Borgs et al. \(2008\)](#) introduce an additional player, *the Kibitzer*: starting from the original game between Alice and Bob, in each iteration the Kibitzer competes with one player, say Alice; both Kibitzer and Alice are scored relative to each other’s strategy when measured against the same Bob strategies.

Our approach In this work, we improve over the reduction from [Anagnostides et al. \(2024\)](#) in both complexity assumption and lower bound on the number of iterations (in our reduction, both are nearly tight). The key to obtaining a better reduction is that we start from ϵ -Bayesian Nash equilibrium (ϵ -BNE⁴) for a one-shot Bayesian game instead of ϵ -NE of a normal-form complete information game. The advantage of starting the reduction from ϵ -BNE is that it is a much harder problem than ϵ -NE: While ϵ -NE can be found in quasipolynomial time [Lipton et al. \(2003\)](#), ϵ -BNE is PPAD-complete [Rubinsein \(2018\)](#).

The reduction from ϵ -BNE becomes more delicate because it is possible that even off-equilibrium players only have an ϵ -improving deviation for one type, while for a random/average type their

3. We believe that “low rank” is a better term to distinguish from distributions with sparse supports.

4. Specifically, we reduce from what we call (*every-type- ϵ -BNE*), where each player can, *at any type*, gain at most ϵ , by deviating from its current strategy. A formal definition is shown in Definition 4, followed by a detailed discussion.

strategy is approximately optimal⁵. Here we give another application of the Kibitzer player: we let Kibitzer find a type that admits an ϵ -improving deviation. Specifically, for each repetition of the Bayesian game, the Kibitzer first chooses a player, say Alice. Then Kibitzer chooses a type θ for Alice. Kibitzer’s goal is to outperform Alice, on the type it chooses and the action it plays. Nature (the “chance” player) then samples the type θ' of Bob, conditioning on type θ for Alice. Crucially, this way Kibitzer does not know θ' . Finally, Kibitzer and Alice choose actions for type θ , and Bob chooses an action for type θ' .

1.1.2. ALGORITHM FOR BAYESIAN GAMES

Our algorithm is based on the swap-to-external regret reduction of Peng and Rubinstein (2024); Dagan et al. (2024). The idea is to run multiple instance of low-external-regret algorithms, each instance restarts and updates in different frequency. Peng and Rubinstein (2024); Dagan et al. (2024) prove that as long as the low-external-regret algorithm guarantees at most ϵ -external regret in H iterations, the overall algorithm has ϵ -swap regret in $H^{O(1/\epsilon)}$ iterations. In a Bayesian game, if each player has K types and n actions, the entire strategy space has size n^K . By running MWU over the strategy space, it takes $H = O(\log(n^K)/\epsilon^2) = O(K \log(n)/\epsilon^2)$ iterations to guarantee ϵ -external regret, and therefore, a total number of $(K \log(n)/\epsilon^2)^{O(1/\epsilon)}$ iterations.

Recall the number of iterations in Theorem 2 is independent of the number of types K . Our key idea is, instead of running MWU over the entire strategy space, we run MWU over each type and the final strategy is taken as the product at each type⁶. The key observation is, the appearance probability of each type is fixed during no-regret learning, so after taking the product over other types, the external regret guarantee continues to hold (for every type). One can then apply the swap-to-external regret analysis of Peng and Rubinstein (2024); Dagan et al. (2024).

1.2. Related work

Decentralized learning in games There is a long line of work regard learning in games Freund and Schapire (1997, 1999); Foster and Vohra (1997, 1998, 1999); Hart and Mas-Colell (2000, 2001); Cesa-Bianchi and Lugosi (2003); Stoltz and Lugosi (2005); Blum and Mansour (2007); Stoltz and Lugosi (2007); Papadimitriou and Roughgarden (2008); Hart and Mansour (2010); Hart and Mas-Colell (2013); Syrgkanis et al. (2015); Foster et al. (2016); Chen and Peng (2020); Daskalakis et al. (2021a); Anagnostides et al. (2022a,b); Peng and Rubinstein (2023, 2024); Dagan et al. (2024).

Our work focuses on incomplete information games – Bayesian games and extensive-form games – a central topic in the literature Forges (1993); Hartline et al. (2015); Fujii (2023); Von Stengel and Forges (2008); Farina et al. (2022a); Zinkevich et al. (2007); Lanctot et al. (2009); Farina et al. (2019b,a, 2022b); Bai et al. (2022); Anagnostides et al. (2023); Peng and Rubinstein (2024); Dagan et al. (2024) with important real-world applications Bowling et al. (2015); Brown and Sandholm (2018, 2019); Brown et al. (2019). For CCE, there are decentralized learning algorithms (e.g. Farina et al. (2022b)) that converge to the set of ϵ -CCE of extensive-form games (resp. Bayesian game) in $O(|Z| \log(n))$ iterations (for constant $\epsilon > 0$). Our lower bound (Theorem 1) gives a near matching lower bound. For the stronger notion of NFCE, it was often conjectured intractable (e.g. see Von Stengel and Forges (2008)). The recent work Peng and Rubinstein (2024); Dagan et al.

5. This would not be an issue if we start the reduction from what we call (*ex-ante- ϵ*)-BNE, however this more relaxed solution concept is not known to be PPAD-hard

6. This is equivalent to running MWU over the entire strategy space, but assigning different learning rates to each type

(2024) refute this conjecture for constant $\epsilon > 0$ and give an algorithm that converges to an ϵ -NFCE in $(|\mathcal{I}| \log(n)/\epsilon^2)^{O(1/\epsilon)}$ iterations. Our work builds upon the framework of Peng and Rubinstein (2024); Dagan et al. (2024) but improves the number of iterations to $(\log(n)/\epsilon^2)^{O(1/\epsilon)}$ for Bayesian games, which are completely independent of the number of types.

Lower bounds for games The complexity class of PPAD (Polynomial Parity Arguments on Directed graphs) was first introduced by Papadimitriou Papadimitriou (1994) to capture a particular genre of total search functions. Since then, a variety of important problems were shown to be PPAD-hard Chen et al. (2009a); Vazirani and Yannakakis (2011); Chen et al. (2013); Othman et al. (2016); Rubinstein (2018, 2019); Filos-Ratsikas et al. (2021); Daskalakis et al. (2021b); Papadimitriou and Peng (2021); Deligkas et al. (2022); Chen et al. (2022); Chen and Peng (2023); Daskalakis et al. (2023); Jin et al. (2023), a well-known example is the complexity of Nash equilibrium in normal-form games Chen et al. (2009b); Daskalakis et al. (2009).

A recent line of work Foster et al. (2023); Anagnostides et al. (2024) studied the complexity of decentralized learning for coarse correlated equilibrium, by examining the computational complexity of finding a low rank coarse correlated equilibrium. Foster et al. (2023) initiated this line of work and showed it is PPAD-hard to find an ϵ -CCE in Markov game with polynomial rank, for inverse polynomially small ϵ ⁷. Anagnostides et al. (2024) studied the extensive-form game and proved that there is no polynomial time algorithm for finding rank- $2^{\log_2^{1/2-o(1)} |\mathcal{I}|}$ ϵ -CCE in extensive-form game for some constant $\epsilon > 0$, under the exponential time hypothesis (ETH) for PPAD. Here ETH for PPAD is a refinement of the $P \neq \text{PPAD}$ conjecture, which asserts that the PPAD-complete problem END-OF-LINE requires exponential time. This conjecture was motivated by the exponential lower bound in the black box model and it was used in the work of Rubinstein (2016) to rule out polynomial time algorithms for finding approximate Nash equilibrium in normal-form games. In our work, we use a substantial weaker conjecture, i.e., we only need to assume the complete problem in PPAD does not have quasipolynomial time algorithms.

2. Preliminary

Notation Throughout the paper, we write $[n] = \{1, 2, \dots, n\}$ and $[n_1 : n_2] = \{n_1, n_1 + 1, \dots, n_2\}$. For any finite set S , we use $\Delta(S)$ to denote probability distributions over S . For any distribution p, q , we use $\text{TV}(p; q)$ to denote the total variation distance between p and q .

2.1. Bayesian game

In an m -player Bayesian game, let Θ_i ($|\Theta_i| = K$) be the set of types and let A_i ($|A_i| = n$) be the set of actions for player i . Let $\Theta = \Theta_1 \times \dots \times \Theta_m$ be the set of type profiles of all players. The type profile $\theta \in \Theta$ is generated from a prior distribution $\rho \in \Delta(\Theta)$. Let $A = A_1 \times \dots \times A_m$ be the action profiles of all players. The utility function $u_i : \Theta \times A \rightarrow [0, 1]$ of player i maps a type profile and an action profile to a payoff value. A strategy $s_i : \Theta_i \rightarrow A_i$ determines the action $s_i(\theta_i) \in A_i$ to be selected when player i has private type $\theta_i \in \Theta_i$. Let $S_i = A_i^{\Theta_i}$ be the set of pure strategies of player i and $S = S_1 \times \dots \times S_m$ be the entire strategy space. We use $\Delta(S_i)$ to denote the set of mixed strategies, and $X_i := \prod_{\theta_i \in \Theta_i} \Delta(A_i)$ to denote the set of behaviour strategies.

7. Foster et al. (2023) also showed a lower bound for constant ϵ under ETH for PPAD

A normal-form correlated equilibrium (also known as strategic-form correlated equilibrium) is a joint distribution $\mu \in \Delta(S)$ over the strategy space, such that no player can gain from deviating from the recommend strategy. Formally

Definition 3 (Normal-form correlated equilibrium) *A distribution $\mu \in \Delta(S)$ is a normal-form correlated equilibrium if for every player $i \in [m]$ and all swap function $\phi : S_i \rightarrow S_i$,*

$$\mathbb{E}_{\theta \sim \rho} \mathbb{E}_{s \sim \mu} [u_i(\theta; s(\theta))] \geq \mathbb{E}_{\theta \sim \rho} \mathbb{E}_{s \sim \mu} [u_i(\theta; \phi(s_i)(\theta_i), s_{-i}(\theta_{-i}))].$$

A Bayesian Nash equilibrium (BNE) is the special case of NFCE where μ is a product distribution, i.e. each player samples their action independently. More generally, we consider *low-rank* distributions that interpolate between a product distribution (rank 1) and general correlated distribution (unbounded rank).

Two natural ways of extending NFCE and BNE to approximate equilibrium have been considered in the literature; the differences are subtle but important for our purpose. The first notion of approximate equilibrium, which we call *ex-ante-approximate*, considers the potential benefit for a player from deviating from the recommended strategy in expectation over the Bayesian prior over types. The second, stricter notion, which we call *every-type-approximate* considers the potential benefit for a player from deviating – maximized over her types and in expectation over all other players’ types. This distinction is similar (but orthogonal) to approximate Nash equilibrium and well-supported Nash equilibrium in normal-form games (e.g. [Daskalakis et al. \(2009\)](#)).

Definition 4 (Notions of approximate equilibria)

- We say that a distribution μ is an (ex-ante- ϵ)-NFCE (or (ex-ante- ϵ)-BNE when μ is a product distribution) if for every player $i \in [m]$ and every swap function $\phi : S_i \rightarrow S_i$,

$$\mathbb{E}_{\theta \sim \rho} \mathbb{E}_{s \sim \mu} [u_i(\theta; s(\theta))] \geq \mathbb{E}_{\theta \sim \rho} \mathbb{E}_{s \sim \mu} [u_i(\theta; \phi(s_i)(\theta_i), s_{-i}(\theta_{-i}))] - \epsilon.$$

- We say that a distribution μ is an (every-type- ϵ)-NFCE (or (every-type- ϵ)-BNE when μ is a product distribution) if for every player $i \in [m]$, every swap function $\phi : S_i \rightarrow S_i$, and for every type $k \in \Theta_i$,

$$\mathbb{E}_{\theta_{-i} \sim (\rho|\theta_i=k)} \mathbb{E}_{s \sim \mu} [u_i(\theta; s(\theta))] \geq \mathbb{E}_{\theta_{-i} \sim (\rho|\theta_i=k)} \mathbb{E}_{s \sim \mu} [u_i(\theta; \phi(s_i)(\theta_i), s_{-i}(\theta_{-i}))] - \epsilon.$$

Our algorithm for Bayesian games gives an every-type-approximate NFCE, which is another improvement over ([Daskalakis et al., 2023](#); [Peng and Rubinstein, 2024](#)) what compute ex-ante-approximate NFCE. Meanwhile, our hardness for extensive-form games holds also with respect to the more lenient notion of ex-ante approximation. A detailed discussions on different notions of correlated equilibrium in Bayesian games can be found at ([Fujii, 2023](#)).

2.2. Extensive-form game

The extensive-form game extends Bayesian game by incorporating both incomplete information and sequentiality. An m -player extensive-form game can be formulated as a directed game tree Γ . Let \mathcal{N} be all nodes of Γ , the nodes of the game tree are partitioned into decision nodes and chance nodes $\mathcal{N} = \mathcal{N}_1 \cup \dots \cup \mathcal{N}_m \cup \mathcal{N}_c$. Here \mathcal{N}_i ($i \in [m]$) is the set of nodes where player i takes

actions and \mathcal{N}_c are chance nodes. The function of a chance node is to assign an outcome of a chance event, and each outgoing edge represents one possible outcome of that chance event as well as the probability of the event. At a decision node, the edges represent actions and the successor states the result from the player taking those actions. Player i 's decision nodes \mathcal{N}_i are further partitioned into a family \mathcal{I}_i of information sets; for each information set $I \in \mathcal{I}_i$, Player i chooses an action from A_i . We consider *perfect recall* extensive-form games so that each player remembers the history action sequence leads to each information set. We use $|\mathcal{I}| = \sum_{i \in [m]} |\mathcal{I}_i|$ to denote the total number of information sets. Let $S_i^{\text{EFG}} = \prod_{I \in \mathcal{I}_i} A_i = (A_i)^{\mathcal{I}_i}$ be the set of pure strategies of player i , and the entire strategy space is $S^{\text{EFG}} = \prod_{i \in [m]} S_i^{\text{EFG}}$.

Let $\Delta(S_i^{\text{EFG}})$ be the set of *mixed strategies*. Note that this notion of mixed strategies is very general in that it allows a player to correlate their actions across information sets that cannot simultaneously occur in the same play of the game. We also consider a slightly more restricted notion of *behaviour strategies* that do not account for such correlations; formally the set of Player i 's behaviour strategies is given by $X_i^{\text{EFG}} = \prod_{I \in \mathcal{I}_i} \Delta(A_i)$. Given a strategy distribution $\mu \in \Delta(S^{\text{EFG}})$, we use $u_i(\mu) \in [0, 1]$ to denote the expected utility of player i under μ .

In a coarse correlated equilibrium, no player can gain by switching to another fixed strategy⁸.

Definition 5 (Coarse correlated equilibrium) *Let $\epsilon > 0$, a distribution $\mu \in \Delta(S^{\text{EFG}})$ is an ϵ -approximate coarse correlated equilibrium (ϵ -CCE) of an extensive-form game, if for all player $i \in [m]$ and strategy $s_i \in S_i^{\text{EFG}}$,*

$$u_i(\mu) \geq u_i(s_i, \mu_{-i}) - \epsilon.$$

Remark 6 (Mixed strategy, behaviour strategy and rank) *In the literature of Bayesian and extensive-form game, a mixed strategy places a distribution over the set of pure strategies, while a behaviour strategy places a distribution over the action set at each information set. The Kuhn's Theorem (Kuhn, 1953), a fundamental result of imperfect information game, states that the behaviour strategy is equivalent to the mixed strategy, in the sense that for any mixed strategy $p_i \in \Delta(S_i^{\text{EFG}})$, one can find a behaviour strategy $x_i \in \prod_{I \in \mathcal{I}_i} \Delta(A_i)$, such that the probability of reaching any terminal node is the same. This means that:*

- **CCE \Leftrightarrow behaviour strategies:** *When considering the rank of a coarse correlated equilibrium, it suffices to define the rank with respect to behaviour strategy – one can always replace a mixed strategy with a behaviour strategy and the utility remains the same⁹.*
- **CE \Leftrightarrow mixed strategies:** *Meanwhile, when considering decentralized learning of correlated equilibrium, we should still stick with mixed strategies. The reason is that, after replacing a mixed strategy with a behaviour strategy, even though the utility received by each agent is the same, the hypothetical utility after swapping the strategy is not the same! See Appendix A for an example.*

Formally, we define the rank of a joint distribution as follows.

Definition 7 (Rank) *Let $\mu \in \Delta(S^{\text{EFG}})$ be a joint distribution. When considering CE, we say μ is of rank T , if it can be written as $\mu = \frac{1}{T} \sum_{t=1}^T \otimes_{i \in [m]} p_i^{(t)}$ where $p_i^{(t)} \in \Delta(S_i)$ is a mixed strategy;*

8. This is also called normal-form coarse correlated equilibrium (NFCCE).

9. This notion of rank/sparsity is used in the previous work of (Foster et al., 2023; Anagnostides et al., 2024).

When considering CCE, we say μ is of rank T if it can be written as $\mu = \frac{1}{T} \sum_{t=1}^T \bigotimes_{i \in [m]} x_i^{(t)}$, where $x_i^{(t)} \in X_i^{\text{EFG}}$ is a behaviour strategy.

Note we define the rank of a distribution as the number of *uniform mixture* of product distributions. This is wlog – if the distribution can be written as k non-uniform mixture of production distributions, then one can approximate it with $O(k)$ uniform mixture.

2.3. Useful tools

No-regret learning We make use of the classic algorithm of Multiplicative Weights Update (MWU) for regret minimization (Littlestone and Warmuth, 1994).

Algorithm 1 MWU

- 1: **Input parameters** T (number of rounds), n (number of actions), B (bound on payoff)
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: Compute $p_t \in \Delta_n$ over experts such that $p_t(i) \propto \exp(\eta \sum_{\tau=1}^{t-1} r_\tau(i))$ for $i \in [n]$
 - 4: Play p_t and observes $r_t \in [0, B]^n$
 - 5: **end for**
-

Lemma 8 (Arora et al. (2012)) *Let $n, T \geq 1$ and the reward $r_t \in [0, B]^n$ ($t \in [T]$). If one takes $\eta = \sqrt{\log(n)/T}/B$, then the MWU algorithm guarantees an external regret of at most*

$$\max_{i^* \in [n]} \sum_{t \in [T]} r_t(i^*) - \sum_{t \in [T]} \langle p_t, r_t \rangle \leq \frac{\log(n)}{\eta} + \eta T B^2 \leq 2B \sqrt{T \log(n)}.$$

Online density estimation We make use of the online density estimation algorithm (Vovk, 1990). In the task of online density estimation, there is a set of outcomes O and a set of possible contexts C . The task proceeds in H rounds, and at each round $h \in [H]$,

- Nature first reveals a context $c_h \in C$
- The learner predicts a distribution $p_h \in \Delta(O)$ over outcomes based on the context c_h ; and
- Nature chooses an outcome $o_h \in O$

The learner has access to a set of *experts* $\{p^{(i)}\}_{i \in [n]}$, where each expert is a fixed function from context to distribution of outcome. That is, for any context $c \in C$, $p_c^{(i)} \in \Delta(O)$ is a distribution of outcome, and $p_c^{(i)}(o)$ is the probability of outcome $o \in O$. Vovk’s algorithm (see Algorithm 2) places a distribution over experts at each round, by running MWU over the logarithmic loss function.

Lemma 9 (Vovk (1990)) *Suppose the distribution of outcomes is realizable under the set of expert $[n]$, i.e., there exists $i^* \in [n]$ such that $o_h \sim p_{c_h}^{(i^*)}$ for each $h \in [H]$. Then Vovk’s aggregating rule (Algorithm 2) satisfies*

$$\frac{1}{H} \sum_{h=1}^H \mathbb{E}[\text{TV}(p_h; p_{c_h}^{(i^*)})] \leq \sqrt{\frac{\log(n)}{H}}.$$

Algorithm 2 Vovk's aggregating algorithm

- 1: **for** $h = 1, 2, \dots, H$ **do**
- 2: Observe the context c_h
- 3: Compute a distribution $q_h \in \Delta_n$ over experts $[n]$

$$q_h(i) = \frac{\exp\left(-\sum_{\nu=1}^{h-1} \log\left(\frac{1}{p_{c_\nu}^{(i)}(o_\nu)}\right)\right)}{\sum_{i' \in [n]} \exp\left(-\sum_{\nu=1}^{h-1} \log\left(\frac{1}{p_{c_\nu}^{(i')}(o_\nu)}\right)\right)} \quad \forall i \in [n]$$

- 4: Aggregate predictions from experts $p_h = \mathbb{E}_{i \sim q_h}[p_{c_h}^{(i)}] \in \Delta(O)$ and observe the outcome $o_h \in O$
 - 5: **end for**
-

The complexity of Bayesian Nash equilibrium We make use of the following hardness result of BNE.

Theorem 10 (Rubinstein (2018)) *It is PPAD-hard to find an (every-type- ϵ)-BNE of a two-player, n -type, 2-action Bayesian game for some obsolete constant $\epsilon > 0$.*

3. Lower bound for extensive-form game

Theorem 1 is a direct corollary of the following hardness result on computing low rank CCE.

Theorem 11 *Let $\epsilon > 0$ be an absolute constant. Assuming $\text{PPAD} \not\subseteq \text{TIME}(n^{\text{polylog}(n)})$, there is no polynomial time algorithm that can find a rank- $2^{\log_2^{1-o(1)}(|Z|)}$ ϵ -CCE of a three player extensive-form game.*

We reduce (every-type- ϵ)-BNE to low rank CCE of extensive-form game. Given a two-player Bayesian game \mathcal{G} , with type space $\Theta = \Theta_1 \times \Theta_2$, action space $A = A_1 \times A_2$ and prior distribution ρ , we construct a three-player extensive-form game Γ .

Gadget We first construct a three-player extensive-form game gadget that incorporates the Bayesian game. Among the first two players, we use $i \in [2]$ to represent a player and we use i' to denote the other player than i , i.e., $i' = 3 - i \in [2]$. We use the special symbol $K = 3$ to represent the third player Kibitzer.

The extensive-form game proceeds as follows. Kibitzer first picks a player $i \in [2]$, a type $\theta_i \in \Theta_i$ as well as an action $\tilde{a}_i \in A_i$, that is, the action set A_K is defined as

$$A_K = \{(i, \theta_i, \tilde{a}_i), i \in [2], \theta_i \in \Theta_i, \tilde{a}_i \in A_i\}.$$

Nature then samples $\theta_{i'}$ from ρ conditioning on θ_i . Then player i (resp. i') observes its private type θ_i (resp. $\theta_{i'}$) and plays action a_i (resp. $a_{i'}$). We use $O = A_1 \times A_2 \times A_K \times \Theta_1 \times \Theta_2$ to denote the outcome space of the gadget.

When Kibitzer plays $a_K = (1, \theta_1, \tilde{a}_1)$, then the utility of player $j \in [3]$ equals

$$u_j(a_1, a_2, a_K, \theta_1, \theta_2) = \begin{cases} u_1(\theta_1, \theta_2; a_1, a_2) - u_1(\theta_1, \theta_2; \tilde{a}_1, a_2) & j = 1 \\ 0 & j = 2 \\ u_1(\theta_1, \theta_2; \tilde{a}_1, a_2) - u_1(\theta_1, \theta_2; a_1, a_2) & j = K \end{cases}. \quad (1)$$

When Kibitzer plays $a_K = (2, \theta_2, \tilde{a}_2)$, the utility of player $j \in [3]$ is defined similarly

$$u_j(a_1, a_2, a_K, \theta_1, \theta_2) = \begin{cases} 0 & j = 1 \\ u_2(\theta_1, \theta_2; a_1, a_2) - u_2(\theta_1, \theta_2; a_1, \tilde{a}_2) & j = 2 \\ u_2(\theta_1, \theta_2; a_1, \tilde{a}_2) - u_2(\theta_1, \theta_2; a_1, a_2) & j = K \end{cases}. \quad (2)$$

Final construction Let T be the rank of the CCE. The final extensive-form game Γ is constructed by repeating the gadget for $H \geq \frac{\log(T)}{\epsilon^2}$ times. For each repetition $h \in [H]$, let N_h be the collection of gadgets at the h -th repetition, corresponding to different paths of play in earlier repetition. For any gadget $n_h \in N_h$, let $o_h = (a_{1,h}, a_{2,h}, a_{K,h}, \theta_{1,h}, \theta_{2,h}) \in O$ be the outcome of the h -th repetition. The utility of each of the three players is determined by the average over the gadget utilities defined Eq. (1)(2), averaged across the H repetitions; we denote them as $(u_{j,n_h})_{j \in [3]}$.

The proof of Theorem 11 can be found at Appendix B.

4. Efficient algorithm for Bayesian game

Our goal is to prove

Theorem 12 (Formal statement of Theorem 2) *Let $\epsilon > 0$, for any m -player, n -action, K -type Bayesian game, there exists uncoupled dynamics that converge to the set of (every-type- ϵ)-NFCE in $(\log(n)/\epsilon^2)^{O(1/\epsilon)}$ iterations. The computation cost at each iteration is $O(mnK \log(mnK/\epsilon)/\epsilon^3)$.*

We have each player run the no-regret learning algorithm (shown in Algorithm 3) for a sequence of T days. Algorithm 3 is an adaptation of the multi-scale MWU (Peng and Rubinstein, 2024) for Bayesian game. Given the regret parameter $\epsilon > 0$, it sets $H = \log(n)/\epsilon^2$, $L = 1/\epsilon$ and runs for $T = H^L$ days. It maintains L threads and plays uniformly over them. The L threads restart and update in different frequency. The key innovation of Algorithm 3 is that it runs MWU for each type $k \in [K]$ independently (see Procedure $\text{MWU}_{\ell,k}$) and the strategy of the ℓ -th thread is the product strategy of each type (see Line 7). Concretely, for each type $k \in [K]$, $\text{MWU}_{\ell,k}$ runs MWU and the reward vector is constructed as the expected reward conditioning on type k (see Line 17).

We bound the external regret for each thread $\ell \in [L]$ and each type $k \in [K]$ in Lemma 13, and bound the swap regret in Lemma 14. The proof can be found in Appendix C.

Lemma 13 (External regret) *For each thread $\ell \in [L]$, restart $\beta \in [T/H^\ell]$, type $k \in [K]$ and action $j \in [n]$, we have*

$$\sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} r_{t,k}(j) - \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} \sum_{s_i \in S_i} q_{t,\ell}(s_i) r_{t,k}(s_i(k)) \leq 2\epsilon H^\ell.$$

Lemma 14 (External to swap regret reduction) *For any type $k \in [K]$ and any swap function $\phi : S_i \rightarrow S_i$, we have*

$$\sum_{t \in [T]} \sum_{s_i \in S_i} p_t(s_i) r_{t,k}(\phi(s_i)(k)) - p_t(s_i) r_{t,k}(s_i(k)) \leq 3\epsilon T.$$

Algorithm 3 Multi-scale MWU for Bayesian game

▷ No-regret learning algorithm for player $i \in [m]$, the opponents play $\mu_t^{(-i)} \in \Delta(S_{-i})$ at day $t \in [T]$

- 1: **Parameters** $H = \log(n)/\epsilon^2$, $L = 1/\epsilon$, $T = H^L$
 - 2: Let $q_{t,\ell} \in (\Delta(A_i))^{\Theta_i}$ be the strategy of `THREAD $_{\ell}$` ($\ell \in [L]$) at day $t \in [T]$
 - 3: Play the uniform mixture $p_t = \frac{1}{L} \sum_{\ell \in [L]} q_{t,\ell} \in \Delta(S_i)$
 - 4:
 - 5: **procedure** `THREAD $_{\ell}$` ▷ Thread $\ell \in [L]$
 - 6: Let $w_{t,\ell,k} \in \Delta(A_i)$ be the strategy of `MWU $_{\ell,k}$` ($k \in [K]$) at day $t \in [T]$
 - 7: Play the product distribution $q_{t,\ell} = w_{t,\ell,1} \times \cdots \times w_{t,\ell,K} \in (\Delta(A_i))^{\Theta_i}$
 - 8: **end procedure**
 - 9:
 - 10: **procedure** `MWU $_{\ell,k}$` ▷ Thread $\ell \in [L]$, type $k \in [K]$
 - 11: **for** $\beta = 1, 2, \dots, T/H^{\ell}$ **do** ▷ β -th restart, restart every H^{ℓ} days
 - 12: Initiate MWU with parameters $H, n, H^{\ell-1}$
 - 13: **for** $h = 1, 2, \dots, H$ **do** ▷ Round h
 - 14: Let $z_{\beta,h,\ell,k} \in \Delta(A_i)$ be the strategy of MWU at the h -th round
 - 15: Play $w_{t,\ell,k} = z_{\beta,h,\ell,k}$ at day $t \in [(\beta-1)H^{\ell} + (h-1)H^{\ell-1} + 1 : (\beta-1)H^{\ell} + hH^{\ell-1}]$
 - 16: Update MWU with the aggregated rewards of the past $H^{\ell-1}$ days

$$\left\{ \sum_{t=(\beta-1)H^{\ell} + (h-1)H^{\ell-1} + 1}^{(\beta-1)H^{\ell} + hH^{\ell-1}} r_{t,k}(j) \right\}_{j \in [n]} \in [0, H^{\ell-1}]^n$$
 - 17: where

$$r_{t,k}(j) = \mathbb{E}_{s_{-i} \sim \mu_t^{(-i)}} \mathbb{E}_{\theta_{-i} \sim (\rho|\theta_i=k)} [u_i(\theta; j, s_{-i}(\theta_{-i}))] \quad \forall j \in [n]$$
 - 18: **end for**
 - 19: **end for**
 - 20: **end procedure**
-

Proof [Proof of Theorem 12] We make each player $i \in [m]$ run Algorithm 3 for $T = (\log(n)/\epsilon^2)^{1/\epsilon}$ rounds. Let $p_t^{(i)} \in \Delta(S_i)$ be the strategy distribution of player i at round $t \in [T]$, then the joint distribution is $\mu_t = \bigotimes_{i \in [m]} p_t^{(i)}$. We take $\mu_t^{(-i)} = \bigotimes_{i' \neq i} p_t^{(i')}$ for any $i \in [m]$. We prove the empirical distribution $\mu = \frac{1}{T} \sum_{t \in [T]} \mu_t$ is a (every-type- 3ϵ)-NFCE.

For each player $i \in [m]$, type $k \in [K]$, swap function $\phi : S_i \rightarrow S_i$, we have

$$\begin{aligned}
 & \mathbb{E}_{\theta \sim (\rho|\theta_i=k)} \mathbb{E}_{s \sim \mu} [u_i(\theta; \phi(s_i)(\theta_i), s_{-i}(\theta_{-i})) - u_i(\theta; s(\theta))] \\
 &= \frac{1}{T} \mathbb{E}_{\theta \sim (\rho|\theta_i=k)} \sum_{t \in [T]} \sum_{s_i \in S_i} p_t^{(i)}(s_i) \mathbb{E}_{s_{-i} \sim \mu_t^{(-i)}} [u_i(\theta; \phi(s_i)(k), s_{-i}(\theta_{-i})) - u_i(\theta; s_i(k), s_{-i}(\theta_{-i}))] \\
 &= \frac{1}{T} \sum_{t \in [T]} \sum_{s_i \in S_i} p_t^{(i)}(s_i) \underbrace{\mathbb{E}_{s_{-i} \sim \mu_t^{(-i)}} \mathbb{E}_{\theta \sim (\rho|\theta_i=k)} [u_i(\theta; \phi(s_i)(k), s_{-i}(\theta_{-i})) - u_i(\theta; s_i(k), s_{-i}(\theta_{-i}))]}_{r_{t,k}^{(i)}(\phi(s_i)(k)) - r_{t,k}^{(i)}(s_i(k))} \\
 &= \frac{1}{T} \sum_{t \in [T]} \sum_{s_i \in S_i} p_t^{(i)}(s_i) (r_{t,k}^{(i)}(\phi(s_i)(k)) - r_{t,k}^{(i)}(s_i(k))) \leq 3\epsilon.
 \end{aligned}$$

The first step follows from the construction of $\mu = \frac{1}{T} \sum_{t \in [T]} \mu_t = \frac{1}{T} \sum_{t \in [T]} \bigotimes_{i \in [m]} p_t^{(i)}$, the third step follows from the definition of reward $r_{t,k}^{(i)}$ for player i (Line 17 of Algorithm 3), the last step follows from Lemma 14. Hence, we have proved μ is a (every-type- 3ϵ)-NFCE.

For computation efficiency, a player needs to compute $O(nK)$ reward entries per iteration (Line 17 of Algorithm 3). Each of those entries can be approximated (with high probability, up to $\pm\epsilon$ error) by random sampling $O(\log(mTnKL)/\epsilon^2) = O(\log(mnK/\epsilon)/\epsilon^3)$ entries from the joint strategy distribution μ_t and the prior distribution over Θ – we assume one can sample from the conditional distribution $(\rho|\theta_i = k)$ in $O(m)$ time. The total runtime per player equals $O(mnK \log(mnK/\epsilon)/\epsilon^3)$. ■

Acknowledgement

The research of B.P. is supported by NSF CCF-1703925, IIS-1838154, CCF-2106429, CCF-2107187, CCF-1763970, AF2212233, COLL2134095, COLL2212745. The research of A.R. is supported by NSF CCF-1954927, and a David and Lucile Packard Fellowship.

References

- Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, and Tuomas Sandholm. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 736–749, 2022a.
- Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Tuomas Sandholm. Uncoupled learning dynamics with $O(\log T)$ swap regret in multiplayer games. *Advances in Neural Information Processing Systems*, 35:3292–3304, 2022b.
- Ioannis Anagnostides, Gabriele Farina, and Tuomas Sandholm. Near-optimal Φ -regret learning in extensive-form games. In *International Conference on Machine Learning*, pages 814–839. PMLR, 2023.
- Ioannis Anagnostides, Alkis Kalavasis, Tuomas Sandholm, and Manolis Zampetakis. On the complexity of computing sparse equilibria and lower bounds for no-regret learning in games. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2024.
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of computing*, 8(1):121–164, 2012.
- Yakov Babichenko, Christos H. Papadimitriou, and Aviad Rubinfeld. Can almost everybody be almost happy? In Madhu Sudan, editor, *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 1–9. ACM, 2016. doi: 10.1145/2840728.2840731. URL <https://doi.org/10.1145/2840728.2840731>.
- Yu Bai, Chi Jin, Song Mei, Ziang Song, and Tiancheng Yu. Efficient phi-regret minimization in extensive-form games via online mirror descent. *Advances in Neural Information Processing Systems*, 35:22313–22325, 2022.
- Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8(6), 2007.
- Christian Borgs, Jennifer Chayes, Nicole Immorlica, Adam Tauman Kalai, Vahab Mirrokni, and Christos Papadimitriou. The myth of the folk theorem. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 365–372, 2008.
- Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015.
- Noam Brown and Tuomas Sandholm. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.
- Noam Brown and Tuomas Sandholm. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890, 2019.

- Noam Brown, Adam Lerer, Sam Gross, and Tuomas Sandholm. Deep counterfactual regret minimization. In *International conference on machine learning*, pages 793–802. PMLR, 2019.
- Nicolo Cesa-Bianchi and Gábor Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51:239–261, 2003.
- Thomas Chen, Xi Chen, Binghui Peng, and Mihalis Yannakakis. Computational hardness of the hylland-zeckhauser scheme. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2253–2268. SIAM, 2022.
- Xi Chen and Binghui Peng. Hedging in games: Faster convergence of external and swap regrets. *Advances in Neural Information Processing Systems*, 33:18990–18999, 2020.
- Xi Chen and Binghui Peng. Complexity of equilibria in first-price auctions under general tie-breaking rules. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 698–709, 2023.
- Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. Settling the complexity of arrow-debreu equilibria in markets with additively separable utilities. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 273–282. IEEE, 2009a.
- Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009b.
- Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 181–190, 2013.
- Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. From external to swap regret 2.0: An efficient reduction and oblivious adversary for large action spaces. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, 2024.
- Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. Near-optimal no-regret learning in general games. *Advances in Neural Information Processing Systems*, 34:27604–27616, 2021a.
- Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In *Proceedings of the 53th Annual ACM SIGACT Symposium on Theory of Computing*, 2021b.
- Constantinos Daskalakis, Noah Golowich, and Kaiqing Zhang. The complexity of markov equilibrium in stochastic games. In *The Thirty Sixth Annual Conference on Learning Theory*. PMLR, 2023.
- Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. Pure-circuit: Strong inapproximability for ppad. In *2022 IEEE 63th Annual Symposium on Foundations of Computer Science (FOCS)*, 2022.

- Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Optimistic regret minimization for extensive-form games via dilated distance-generating functions. *Advances in neural information processing systems*, 32, 2019a.
- Gabriele Farina, Chun Kai Ling, Fei Fang, and Tuomas Sandholm. Efficient regret minimization algorithm for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems*, 32, 2019b.
- Gabriele Farina, Andrea Celli, Alberto Marchesi, and Nicola Gatti. Simple uncoupled no-regret learning dynamics for extensive-form correlated equilibrium. *Journal of the ACM*, 69(6):1–41, 2022a.
- Gabriele Farina, Chung-Wei Lee, Haipeng Luo, and Christian Kroer. Kernelized multiplicative weights for 0/1-polyhedral games: Bridging the gap between learning in extensive-form and normal-form games. In *International Conference on Machine Learning*, pages 6337–6357. PMLR, 2022b.
- Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. On the complexity of equilibrium computation in first-price auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 454–476, 2021.
- Françoise Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and decision*, 35:277–310, 1993.
- Dean P Foster and Rakesh Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, 29(1-2):7–35, 1999.
- Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40, 1997.
- Dean P Foster and Rakesh V Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.
- Dylan J Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. *Advances in Neural Information Processing Systems*, 29, 2016.
- Dylan J Foster, Noah Golowich, and Sham M Kakade. Hardness of independent learning and sparse equilibrium computation in markov games. In *International Conference on Machine Learning*. PMLR, 2023.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- Kaito Fujii. Bayes correlated equilibria and no-regret dynamics. *arXiv preprint arXiv:2304.05005*, 2023.
- Sergiu Hart and Yishay Mansour. How long to equilibrium? the communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69(1):107–126, 2010.

- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Sergiu Hart and Andreu Mas-Colell. A reinforcement procedure leading to correlated equilibrium. In *Economics Essays: A Festschrift for Werner Hildenbrand*, pages 181–200. Springer, 2001.
- Sergiu Hart and Andreu Mas-Colell. *Simple adaptive strategies: from regret-matching to uncoupled dynamics*, volume 4. World Scientific, 2013.
- Jason Hartline, Vasilis Syrgkanis, and Eva Tardos. No-regret learning in bayesian games. *Advances in Neural Information Processing Systems*, 28, 2015.
- Yujia Jin, Vidya Muthukumar, and Aaron Sidford. The complexity of infinite-horizon general-sum stochastic games. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.
- HW Kuhn. Extensive games and the problem of information. *Contributions to the Theory of Games*, page 193, 1953.
- Marc Lanctot, Kevin Waugh, Martin Zinkevich, and Michael Bowling. Monte carlo sampling for regret minimization in extensive games. *Advances in neural information processing systems*, 22, 2009.
- Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 36–41, 2003.
- Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.
- Abraham Othman, Christos Papadimitriou, and Aviad Rubinstein. The complexity of fairness through equilibrium. *ACM Transactions on Economics and Computation (TEAC)*, 4(4):1–19, 2016.
- Christos Papadimitriou and Binghui Peng. Public goods games in directed networks. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 745–762, 2021.
- Christos H Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994.
- Christos H Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):1–29, 2008.
- Binghui Peng and Aviad Rubinstein. Near optimal memory-regret tradeoff for online learning. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1171–1194. IEEE, 2023.
- Binghui Peng and Aviad Rubinstein. Fast swap regret minimization and applications to approximate correlated equilibria. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, 2024.

- Aviad Rubinstein. Settling the complexity of computing approximate two-player nash equilibria. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 258–265. IEEE, 2016.
- Aviad Rubinstein. Inapproximability of nash equilibrium. *SIAM Journal on Computing*, 47(3): 917–959, 2018.
- Aviad Rubinstein. *Hardness of Approximation Between P and NP*. Morgan & Claypool, 2019.
- Gilles Stoltz and Gábor Lugosi. Internal regret in on-line portfolio selection. *Machine Learning*, 59:125–159, 2005.
- Gilles Stoltz and Gábor Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 59(1):187–208, 2007.
- Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. *Advances in Neural Information Processing Systems*, 28, 2015.
- Vijay V Vazirani and Mihalis Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)*, 58(3):1–25, 2011.
- Bernhard Von Stengel and Françoise Forges. Extensive-form correlated equilibrium: Definition and computational complexity. *Mathematics of Operations Research*, 33(4):1002–1022, 2008.
- Vladimir Vovk. Aggregating strategies. In *Proceedings of 3rd Annu. Workshop on Comput. Learning Theory*, pages 371–383, 1990.
- Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20, 2007.

Appendix A. On the rank of CE

We give a simple example of Bayesian game (thus also an example of extensive-form game), where behaviourizing a rank-2 CE is no longer a CE. There are two players, Alice and Bob, in the game. Bob has 2 actions, one type and his utility is always 0. Alice has n types and 4 actions. For each type, Alice has utility

$$u_A(0, 0) = 0, u_A(1, 0) = 0, u_A(2, 0) = 1, u_B(3, 0) = -2$$

and

$$u_A(0, 1) = 0, u_A(1, 1) = 0, u_A(2, 1) = -2, u_A(3, 1) = 1.$$

Consider the following rank-2 CE $\mu = \frac{1}{2}p_A^{(1)} \otimes p_B^{(1)} + \frac{1}{2}p_A^{(2)} \otimes p_B^{(2)}$, where $p_A^{(1)} = \frac{1}{3}(0, 0, \dots, 0) + \frac{2}{3}(1, 1, \dots, 1)$, $p_A^{(2)} = \frac{2}{3}(0, 0, \dots, 0) + \frac{1}{3}(1, 1, \dots, 1)$; $p_B^{(1)} = 0$, $p_B^{(2)} = 1$. That is, in the first product distribution, Alice plays all 0 w.p. $1/3$ and plays all 1 w.p. $2/3$, Bob always plays 0; in the second product distribution, Bob plays all 0 w.p. $2/3$ and plays all 1 w.p. $1/3$, Bob always plays 1. It is not hard to verify that μ is indeed a CE and Alice receives 0 utility.

Meanwhile, the only way to behaviourize μ , is to replace $p_A^{(1)}$ with $x_A^{(1)} = (\frac{1}{3}, \frac{2}{3}) \otimes \dots \otimes (\frac{1}{3}, \frac{2}{3})$ and $x_B^{(1)} = (\frac{2}{3}, \frac{1}{3}) \otimes \dots \otimes (\frac{2}{3}, \frac{1}{3})$. Let $\mu' = \frac{1}{2}x_A^{(1)} \otimes p_B^{(1)} + \frac{1}{2}x_A^{(2)} \otimes p_B^{(2)}$. Consider the swap function $\phi_A = \{0, 1, 2, 3\}^n \rightarrow \{0, 1, 2, 3\}^n$

$$\phi_A(s) = \begin{cases} (2, \dots, 2) & s \in \{0, 1\}^n, |s| \in (n/3 - \sqrt{n} \log(n), n/3 + \sqrt{n} \log(n)) \\ (3, \dots, 3) & s \in \{0, 1\}^n, |s| \in (2n/3 - \sqrt{n} \log(n), 2n/3 + \sqrt{n} \log(n)) \\ s & \text{otherwise} \end{cases}$$

We claim that applying the swap function ϕ_A , the utility of Alice becomes $1 - o(1)$ (under μ'). First, we know that with probability $1 - o(1)$, for any $(s_A, s_B) \sim \mu'$, $|s_A| \in (n/3 - \sqrt{n} \log(n), n/3 + \sqrt{n} \log(n))$ or $|s_A| \in (2n/3 - \sqrt{n} \log(n), 2n/3 + \sqrt{n} \log(n))$. In the former case, Bob plays 1 w.p. $1 - o(1)$ so switching to action $(3, 3, \dots, 3)$ has utility 1; in the latter case, Bob plays 0 w.p. $1 - o(1)$ so switching to action $(2, 2, \dots, 2)$ has utility 1. This implies μ' is not a CE.

Appendix B. Missing proof from Section 3

We complete the proof of Theorem 11.

Notation We first introduce additional notations. A gadget $n_h \in N_h$ is uniquely determined by its history $\sigma(n_h) = (n_1, o_1, \dots, n_{h-1}, o_{h-1})$, where $n_\tau \in N_\tau$ ($\tau \in [h-1]$) is the gadget at the root path to n_h and o_τ is the outcome at n_τ . We use S_j^{EFG} to denote the strategy space of player $j \in [3]$ in the extensive-form game Γ , and S_i to denote the strategy space of player $i \in [2]$ in the Bayesian game \mathcal{G} . For player $i \in [2]$, we have $S_i^{\text{EFG}} = \prod_{h \in [H]} \prod_{n_h \in N_h} S_i$ and $X_i^{\text{EFG}} = \prod_{h \in [H]} \prod_{n_h \in N_h} X_i$, i.e., the pure/behaviour strategy space is the product of pure/behaviour strategy space at each gadget. Given a behaviour strategy $x_i \in X_i^{\text{EFG}}$, for any gadget $n_h \in N_h$, we write $x_{i, n_h} \in X_i$ to denote the behaviour strategy of x_i on n_h , i.e., x_{i, n_h} is the marginal distribution of x_i on n_h . For player K , we have $S_K^{\text{EFG}} = \prod_{h \in [H]} \prod_{n_h \in N_h} A_K$, $X_K^{\text{EFG}} = \prod_{h \in [H]} \prod_{n_h \in N_h} \Delta(A_K)$, and we define $x_{K, n_h} \in \Delta(A_K)$ similarly.

Given a rank- T ϵ -CCE μ of the extensive-form game Γ , we write

$$\mu = \sum_{t \in [T]} x^{(t)} = \sum_{t \in [T]} x_1^{(t)} \otimes x_2^{(t)} \otimes x_K^{(t)},$$

where $x_j^{(t)} \in X_j^{\text{EFG}}$ ($j \in [3]$) is a behaviour strategy of player j .

B.1. Analysis of players 1,2

We first prove that at any approximate CCE, the utilities of the first two players are not far below zero. To this end, for player $i \in [2]$, we consider the deviation strategy $x_i^\dagger \in S_i^{\text{EFG}}$, whose construction is shown at Algorithm 4. For any product distribution $t \in [T]$, player $i \in [2]$, let $x_{-i}^{(t)} = x_{i'}^{(t)} \times x_K^{(t)}$. Let $O_{-i} = A_{i'} \times A_K \times \Theta_1 \times \Theta_2$ be the outcome space of player K, player i' and the chance player. For any gadget $n_h \in N_h$, let $x_{-i, n_h}^{(t)} \in \Delta(O_{-i})$ be the distribution of outcome (except player i) at gadget n_h under strategy $x_{-i}^{(t)}$.

Algorithm 4 Deviation strategy of player 1, 2 ▷ Analysis only

- 1: **Input:** rank- T ϵ -CCE $\mu = \frac{1}{T} \sum_{t \in [T]} x_1^{(t)} \otimes x_2^{(t)} \otimes x_K^{(t)}$, $i \in [2]$
- 2: **for** $h = 1, 2, \dots, H$ **do** ▷ repetition $h \in [H]$
- 3: **for** $n_h \in N_h$ **do** ▷ gadget n_h
- 4: $\sigma(n_h) = (n_1, o_1, \dots, n_{h-1}, o_{h-1})$ ▷ history of n_h
- 5: Compute a distribution $q_{i, n_h} \in \Delta(T)$ over $[T]$ ▷ Vovk's rule

$$q_{-i, n_h}(t) = \frac{\exp\left(-\sum_{\nu=1}^{h-1} \log\left(\frac{1}{x_{-i, n_\nu}^{(t)}(o_{-i, \nu})}\right)\right)}{\sum_{t' \in [T]} \exp\left(-\sum_{\nu=1}^{h-1} \log\left(\frac{1}{x_{-i, n_\nu}^{(t')}(o_{-i, \nu})}\right)\right)} \quad t \in [T]$$

- 6: Compute $p_{-i, n_h} = \mathbb{E}_{t \sim q_{-i, n_h}} [x_{-i, n_h}^{(t)}] \in \Delta(O_{-i})$ ▷ predicted outcome
- 7: Define the deviation strategy $x_{i, n_h}^\dagger \in S_i$ at gadget n_h

$$x_{i, n_h}^\dagger = \arg \max_{s_i \in S_i} \mathbb{E}_{(a_{i'}, a_K, \theta_1, \theta_2) \sim p_{-i, n_h}} [u_{i, n_h}(s_i(\theta_i), a_{i'}, a_K, \theta_1, \theta_2)]$$

- 8: **end for**
 - 9: **end for**
-

First, we prove

Lemma 15 For player $i \in [2]$, product distribution $t \in [T]$, we have

$$\frac{1}{H} \sum_{h \in [H]} \mathbb{E}_{x_i^\dagger \times x_{-i}^{(t)}} \left[\text{TV}(p_{-i, n_h}; x_{-i, n_h}^{(t)}) \right] \leq \sqrt{\frac{\log(T)}{H}}.$$

Here n_1, n_2, \dots, n_H is the (random) sequence of gadgets visited under strategy $x_i^\dagger \times x_{-i}^{(t)}$.

Proof Consider the following online density estimation procedure. The outcome space is O_{-i} , the set of experts are T strategies $\{x_{-i,n_h}^{(t)}\}_{t \in [T]}$, the context comes from the set of gadgets $\cup_{h \in [H]} N_h$. At each round $h \in [H]$,

- Nature provides the context n_h
- The algorithm predicts the distribution $p_{-i,n_h} \in \Delta(O_{-i})$ over outcome space O_{-i} ;
- Nature reveals the outcome $o_{-i,h} = (a_{i'}, a_{\mathcal{K}}, \theta_1, \theta_2) \sim x_{-i,n_h}^{(t)}$.
- The next context $n_{h+1} = n_h | (o_{-i,h}, x_i^\dagger(\theta_i))$

By Lemma 9, we have

$$\frac{1}{H} \sum_{h=1}^H \mathbb{E}_{x_i^\dagger \times x_{-i}^{(t)}} \left[\text{TV}(p_{-i,n_h}; x_{-i,n_h}^{(t)}) \right] \leq \sqrt{\frac{\log(T)}{H}}.$$

■

Now we can bound the utility of the first two players at ϵ -CCE.

Lemma 16 (Equilibrium utility) *The utility at ϵ -CCE of player $i \in [2]$ satisfies $u_i(\mu) \geq -5\epsilon$.*

Proof For any player $i \in [2]$, production distribution $t \in [T]$, gadget $n_h \in N_h$, define

$$s_{i,n_h}^{(t)} = \arg \max_{s_i \in S_i} \mathbb{E}_{(a_{i'}, a_{\mathcal{K}}, \theta_1, \theta_2) \sim x_{-i,n_h}^{(t)}} u_{i,n_h}(s_i(\theta_i), a_{i'}, a_{\mathcal{K}}, \theta_1, \theta_2).$$

That is $s_{i,n_h}^{(t)} \in S_i$ is the best strategy of player i , given that the other players use strategies $x_{i',n_h}^{(t)} \times x_{\mathcal{K},n_h}^{(t)}$. We claim that

$$u_{i,n_h}(s_{i,n_h}^{(t)}, x_{-i,n_h}^{(t)}) \geq 0. \quad (3)$$

This is because player i can always best-respond to $x_{i',n_h}^{(t)}$ in the Bayesian game, and the strategy $x_{\mathcal{K},n_h}^{(t)}$ of Kibitzer is independent of $x_{i'}^{(t)}$ so it cannot find a better action.

Now we compute its utility under deviation x_i^\dagger . For any $t \in [T]$, gadget $n_h \in N_h$, we have

$$\begin{aligned} & u_{i,n_h}(x_{i,n_h}^\dagger, x_{-i,n_h}^{(t)}) \\ &= u_{i,n_h}(x_{i,n_h}^\dagger, x_{-i,n_h}^{(t)}) - u_{i,n_h}(x_{i,n_h}^\dagger, p_{-i,n_h}) + u_{i,n_h}(x_{i,n_h}^\dagger, p_{-i,n_h}) - u_{i,n_h}(s_{i,n_h}^{(t)}, p_{-i,n_h}) \\ & \quad + u_{i,n_h}(s_{i,n_h}^{(t)}, p_{-i,n_h}) - u_{i,n_h}(s_{i,n_h}^{(t)}, x_{-i,n_h}^{(t)}) + u_{i,n_h}(s_{i,n_h}^{(t)}, x_{-i,n_h}^{(t)}) \\ & \geq -\frac{2}{H} \text{TV}(p_{-i,n_h}; x_{-i,n_h}^{(t)}) + 0 - \frac{2}{H} \text{TV}(p_{-i,n_h}; x_{-i,n_h}^{(t)}) + u_{i,n_h}(s_{i,n_h}^{(t)}, x_{-i,n_h}^{(t)}) \\ & \geq -\frac{4}{H} \text{TV}(p_{-i,n_h}; x_{-i,n_h}^{(t)}). \end{aligned} \quad (4)$$

The second step follows from (i) the value of u_{i,n_h} is between $[-\frac{1}{H}, \frac{1}{H}]$ and (ii) x_{i,n_h}^\dagger is the best response under p_{-i,n_h} . The third step follows from Eq. (3).

Taking an expectation, we have

$$\begin{aligned}
u_i(x_i^\dagger, x_{-i}^{(t)}) &= \sum_{h=1}^H \mathbb{E}_{x_i^\dagger \times x_{-i}^{(t)}} \left[u_{i,n_h}(x_{i,n_h}^\dagger, x_{-i,n_h}^{(t)}) \right] \\
&\geq -\frac{4}{H} \sum_{h=1}^H \mathbb{E}_{x_i^\dagger \times x_{-i}^{(t)}} \left[\text{TV}(p_{-i,n_h}; x_{-i,n_h}^{(t)}) \right] \\
&\geq -4\sqrt{\frac{\log(T)}{H}} \geq -4\epsilon.
\end{aligned}$$

Here the second step follows from Eq. (4), the third step follows from Lemma 15.

Taking an expectation over $t \sim [T]$, we have

$$u_i(x_i^\dagger, \mu_{-i}) = \frac{1}{T} \sum_{t \in [T]} u_i(x_i^\dagger, x_{-i}^{(t)}) \geq -4\epsilon.$$

Note μ is an ϵ -CCE, we have $u_i(\mu) \geq -5\epsilon$ ■

A direct corollary of Lemma 16 is that the utility of player K at ϵ -CCE is not much above 0.

Lemma 17 *The utility at ϵ -CCE of player K satisfies $u_K(\mu) \leq 10\epsilon$.*

Proof The utility of three players sums to 0. By Lemma 16, we have that

$$u_K(\mu) = -u_1(\mu) - u_2(\mu) \leq 10\epsilon. \quad \blacksquare$$

B.2. Reduction

The reduction is formally presented at Algorithm 5, using some of the following notation. Let $O_{-K} = A_1 \times A_2 \times (\Theta_1 \cup \Theta_2)$ be the outcome space of player 1, 2 and the chance player. For any $t \in [T]$, let $x_{-K}^{(t)} = x_1^{(t)} \times x_2^{(t)}$. For any gadget $n_h \in N_h$ and action $a_K = (i, \theta_i, \tilde{a}_i) \in A_K$, let $x_{-K,n_h,a_K}^{(t)} \in \Delta(O_{-K})$ be the distribution of outcomes (except player K) at gadgets n_h , given that player K plays a_K and players 1, 2 use (mixed) strategy profile $x_{1,n_1}^{(t)} \times x_{2,n_2}^{(t)}$.

For any node $n_h \in N_h$, $a_K = (i, \theta_i, \tilde{a}_i) \in A_K$, define $p_{-K,n_h,a_K} \in \Delta(\Theta_{i'} \times A_1 \times A_2)$ such that

$$p_{-K,n_h,a_K}(\theta_{i'}, a_i, a_{i'}) = \Pr_{\hat{\theta}_{i'} \sim \rho|\theta_i} [\hat{\theta}_{i'} = \theta_{i'}] \cdot \Pr_{s_i \sim p_{i,n_h}} [a_i = s_i(\theta_i)] \cdot \Pr_{s_{i'} \sim p_{i',n_h}} [a_{i'} = s_{i'}(\theta_{i'})]$$

That is, it first samples $\theta_{i'} \sim \rho|\theta_i$, then it independently samples strategies $s_i, s_{i'}$ from p_{i,n_h}, p_{i',n_h} (respectively), and finally takes $a_i = s_i(\theta_i), a_{i'} = s_{i'}(\theta_{i'})$. For $i \in [2], \theta_i \in \Theta_i$, we further define $p_{i,n_h,\theta_i} \in \Delta(A_i)$ such that $p_{i,n_h,\theta_i}(a_i) = \mathbb{E}_{s_i \sim p_{i,n_h}} [s_i(\theta_i)]$.

Our goal is to prove, if Algorithm 5 doesn't find an approximate BNE, then Kibitzer has a good deviation strategy x_K^\dagger . First, we prove Kibitzer approximately learns the outcome distribution.

Algorithm 5 Reduction

- 1: **Input:** rank- T ϵ -CCE $\mu = \frac{1}{T} \sum_{t \in [T]} x_1^{(t)} \otimes x_2^{(t)} \otimes x_K^{(t)}$
 - 2: **for** $h = 1, 2, \dots, H$ **do** \triangleright repetition h
 - 3: **for** $n_h \in N_h$ **do** \triangleright gadget n_h
 - 4: $\sigma(n_h) = (n_1, o_1, \dots, n_{h-1}, o_{h-1})$ \triangleright history of n_h
 - 5: Compute a distribution $q_{-K, n_h} \in \Delta(T)$ over $[T]$
- $$q_{-K, n_h}(t) = \frac{\exp\left(-\sum_{\nu=1}^{h-1} \log\left(\frac{1}{x_{-K, n_\nu, a_{K, \nu}}^{(t)}(o_{-K, \nu})}\right)\right)}{\sum_{t' \in [T]} \exp\left(-\sum_{\nu=1}^{h-1} \log\left(\frac{1}{x_{-K, n_\nu, a_{K, \nu}}^{(t')} (o_{-K, \nu})}\right)\right)} \quad \forall t \in [T]$$
- 6: Compute $p_{i, n_h} = \mathbb{E}_{t \sim q_{-K, n_h}} [x_{i, n_h}^{(t)}] \in \Delta(S_i)$ for $i = 1, 2$
 - 7: **if** (p_{1, n_h}, p_{2, n_h}) is an (every-type- 16ϵ)-BNE **then**
 - 8: **return** (p_{1, n_h}, p_{2, n_h})
 - 9: **else**
 - 10: Choose $a_{K, n_h}^\dagger = (i, \theta_i, \tilde{a}_i) \in A_K$ such that
$$\mathbb{E}_{\theta_{i'} \sim \rho_{\theta_i}} \mathbb{E}_{s_i \sim p_{i, n_h}} \mathbb{E}_{s_{i'} \sim p_{i', n_h}} [u_i(\theta_i, \theta_{i'}; \tilde{a}_i, s_{i'}(\theta_{i'})) - u_i(\theta_i, \theta_{i'}; s_i(\theta_i), s_{i'}(\theta_{i'}))] \geq 16\epsilon.$$
 - 11: **end if**
 - 12: **end for**
 - 13: **end for**
-

Lemma 18 Suppose Algorithm 5 does not find an (every-type- 16ϵ)-BNE, i.e. the if-statement in Line 7 is always false, For every $t \in [T]$, we have

$$\frac{1}{H} \sum_{h \in [H]} \mathbb{E}_{x_K^\dagger \times x_{-K}^{(t)}} \left[\text{TV}(p_{-K, n_h, a_{K, h}^\dagger}; x_{-K, n_h, a_{K, h}^\dagger}^{(t)}) \right] \leq 2\sqrt{\frac{\log(T)}{H}}.$$

Here n_1, \dots, n_H is the (random) sequence of gadgets visited under $x_K^\dagger \times x_{-K}^\dagger$ and $a_{K, h}^\dagger = a_{K, n_h}^\dagger$ is the action of player K at gadget n_h .

Proof For any gadget $n_h \in N_h$, define

$$w_{-K, n_h, a_{K, h}^\dagger} = \mathbb{E}_{t' \sim q_{-K, n_h}} \left[x_{-K, n_h, a_{K, h}^\dagger}^{(t')} \right] \in \Delta(O_{-K}).$$

We first bound the (expected) TV distance between $w_{-K, n_h, a_{K, h}^\dagger}$ and $x_{-K, n_h, a_{K, h}^\dagger}^{(t)}$. Consider the following online density estimation procedure. The outcome space is O_{-K} , the set of experts are T indices $\{x_{-K}^{(t)}\}_{t \in [T]}$, and the context comes from $\cup_{h \in [H]} N_h \times A_K$. At round $h \in [H]$,

- Nature provides the context $(n_h, a_{K, h}^\dagger)$;
- The algorithm predicts $w_{-K, n_h, a_{K, h}^\dagger} \in \Delta(O_{-K})$;
- Nature samples the outcome $o_{-K, h} \sim x_{-K, n_h, a_{K, h}^\dagger}^{(t)}$;
- The next context $n_{h+1} = n_h | (a_{K, h}^\dagger, o_{-K, h})$ and $a_{K, h+1}^\dagger = a_{K, n_{h+1}}^\dagger$

By the guarantee of Lemma 9, we have

$$\frac{1}{H} \sum_{h \in [H]} \mathbb{E}_{x_K^\dagger \times x_{-K}^{(t)}} \left[\text{TV}(w_{-K, n_h, a_{K, h}^\dagger}; x_{-K, n_h, a_{K, h}^\dagger}^{(t)}) \right] \leq \sqrt{\frac{\log(T)}{H}}. \quad (5)$$

Now consider the three distributions $p_{-K, n_h, a_{K, h}^\dagger}, w_{-K, n_h, a_{K, h}^\dagger}, x_{-K, n_h, a_{K, h}^\dagger}^{(t)} \in \Delta(\Theta_{i'} \times A_1 \times A_2)$. Intuitively, $x_{-K, n_h, a_{K, h}^\dagger}^{(t)}$ is the outcome under the product distribution $x_{-i}^{(t)}$; $w_{-K, n_h, a_{K, h}^\dagger}$ is the predicted outcome under Vovk's algorithm; $p_{-K, n_h, a_{K, h}^\dagger}$ is the outcome after taking marginal of $w_{-K, n_h, a_{K, h}^\dagger}$. Formally, these three distributions have the same marginal distribution on $\Theta_{i'}$ due to the independence of chance player. Conditioning on the first coordinates being $\theta_{i'} \in \Theta_{i'}$

- $x_{-K, n_h, a_{K, h}^\dagger}^{(t)} | \theta_{i'} = x_{1, n_h, \theta_1}^{(t)} \times x_{2, n_h, \theta_1}^{(t)} \in \Delta(A_1) \times \Delta(A_2)$ is a product distribution;
- $p_{-K, n_h, a_{K, h}^\dagger} | \theta_{i'} = p_{1, n_h, \theta_1} \times p_{2, n_h, \theta_2} \in \Delta(A_1) \times \Delta(A_2)$ is a product distribution; while
- $w_{-K, n_h, a_{K, h}^\dagger} | \theta_{i'} = \mathbb{E}_{t \sim q_{-K, n_h}} [x_{1, n_h, \theta_1}^{(t)} \times x_{2, n_h, \theta_2}^{(t)}] \in \Delta(A_1 \times A_2)$ is a joint distribution over $A_1 \times A_2$

- $p_{-K,n_h,a_{K,h}^\dagger} | \theta_{i'}$ and $w_{-K,n_h,a_{K,h}^\dagger} | \theta_{i'}$ have the same marginal distribution over $a_1 \in A_1$ (resp. $a_2 \in A_2$), because $p_{1,n_h,\theta_1} = \mathbb{E}_{t \sim q_{-K,n_h}} [x_{1,n_h,\theta_1}^{(t)}]$ and $p_{2,n_h,\theta_2} = \mathbb{E}_{t \sim q_{-K,n_h}} [x_{2,n_h,\theta_2}^{(t)}]$ (see Line 6 of Algorithm 5)

Then we have

$$\begin{aligned} \text{TV}(p_{-K,n_h,a_{K,h}^\dagger}; x_{-K,n_h,a_{K,h}^\dagger}^{(t)}) &= \mathbb{E}_{\theta_{i'} \sim \rho | \theta_i} \left[\text{TV} \left(p_{1,n_h,\theta_1} \times p_{2,n_h,\theta_2} ; x_{1,n_h,\theta_1}^{(t)} \times x_{2,n_h,\theta_2}^{(t)} \right) \right] \\ &\leq 2 \mathbb{E}_{\theta_{i'} \sim \rho | \theta_i} \left[\text{TV} \left(w_{-K,n_h,a_{K,h}^\dagger} | \theta_{i'} ; x_{1,n_h,\theta_1}^{(t)} \times x_{2,n_h,\theta_2}^{(t)} \right) \right] \\ &= 2 \text{TV} \left(w_{-K,n_h,a_{K,h}^\dagger} ; x_{-K,n_h,a_{K,h}^\dagger}^{(t)} \right). \end{aligned}$$

Here the first step follows from Fact 20, the second follows from Fact 21 and the last step follows from Fact 20.

Taking an expectation and combining with Eq. (5), we have

$$\begin{aligned} &\frac{1}{H} \sum_{h \in [H]} \mathbb{E}_{x_K^\dagger \times x_{-K}^{(t)}} \left[\text{TV}(p_{-K,n_h,a_{K,h}^\dagger}; x_{-K,n_h,a_{K,h}^\dagger}^{(t)}) \right] \\ &\leq \frac{2}{H} \sum_{h \in [H]} \mathbb{E}_{x_K^\dagger \times x_{-K}^{(t)}} \left[\text{TV}(w_{-K,n_h,a_{K,h}^\dagger}; x_{-K,n_h,a_{K,h}^\dagger}^{(t)}) \right] \leq 2 \sqrt{\frac{\log(T)}{H}}. \end{aligned}$$

■

Now we can prove the correctness of the reduction.

Lemma 19 *Algorithm 5 returns an (every-type- 16ϵ)-BNE.*

Proof We prove by contradiction and assume Algorithm 5 never returns. We compute the utility under deviation x_K^\dagger . For any product distribution $t \in [T]$, we have

$$\begin{aligned} u_K(x_K^\dagger, x_{-K}^{(t)}) &= \sum_{h=1}^H \mathbb{E}_{x_K^\dagger \times x_{-K}^{(t)}} \left[u_{K,n_h}(a_{K,h}^\dagger, x_{-K,n_h,a_{K,h}^\dagger}^{(t)}) \right] \\ &\geq \sum_{h=1}^H \mathbb{E}_{x_K^\dagger \times x_{-K}^{(t)}} \left[-\frac{2}{H} \text{TV}(p_{-K,n_h,a_{K,h}^\dagger}; x_{-K,n_h,a_{K,h}^\dagger}^{(t)}) + u_{K,n_h}(a_{K,h}^\dagger, p_{-K,n_h,a_{K,h}^\dagger}^{(t)}) \right] \\ &\geq \sum_{h=1}^H \mathbb{E}_{x_K^\dagger \times x_{-K}^{(t)}} \left[-\frac{2}{H} \text{TV}(p_{-K,n_h,a_{K,h}^\dagger}; x_{-K,n_h,a_{K,h}^\dagger}^{(t)}) + \frac{16\epsilon}{H} \right] \\ &\geq -4 \sqrt{\frac{\log(T)}{H}} + 16\epsilon \geq 12\epsilon. \end{aligned}$$

The second step holds since u_{K,n_h} is between $[-\frac{1}{H}, \frac{1}{H}]$, the third step follows from the choice $a_{K,h}^\dagger$ (see Line 10 of Algorithm 5). The fourth step follows from Lemma 18.

Taking an expectation over $t \sim [T]$, we have

$$u_{\kappa}(x_{\kappa}^{\dagger}, \mu_{-\kappa}) = \frac{1}{T} \sum_{t \in [T]} u_i(x_{\kappa}^{\dagger}, x_{-\kappa}^{(t)}) \geq 12\epsilon.$$

Note μ is an ϵ -CCE, we have $u_{\kappa}(\mu) \geq u_{\kappa}(x_{\kappa}^{\dagger}, \mu_{-\kappa}) - \epsilon \geq 11\epsilon$. This contradicts Lemma 17, hence, we have proved Algorithm 5 must return an (every-type- 16ϵ)-BNE at some step. \blacksquare

Now we can finish the proof of Theorem 11.

Proof [Proof of Theorem 11] Let $\delta > 0$ be the obsolete constant in Theorem 10. We prove by contradiction and assume that there exists a constant $\alpha > 0$, such that one can find an rank- $2^{\log_2^{1-\alpha}(|\mathcal{I}|)} \frac{\delta}{16}$ -CCE of a three-player extensive-form game with runtime polynomial in $|\mathcal{I}|$. Given a two-player, n types, two-action Bayesian game \mathcal{G} , we construct the extensive-form game Γ with $H = (\frac{\log_2(n)}{\delta^2})^{1/\alpha}$ repetitions. The number of information sets at each gadget is $O(n^2)$, and therefore, the total number of information sets $|\mathcal{I}| = n^{O(H)}$. Moreover, the number of repetition H satisfies

$$\frac{\log(2^{\log_2^{1-\alpha}(|\mathcal{I}|)})}{(\delta/16)^2} = O(\delta^{-2} \log_2^{1-\alpha}(|\mathcal{I}|)) = O(H^{1-\alpha} \log_2^{1-\alpha}(n) \delta^{-2}) \leq H.$$

Hence, by Lemma 19, one can find an (every-type- δ)-BNE in time $\text{poly}(|\mathcal{I}|) = n^{O(H)} = n^{\text{polylog}(n)}$, which, by Theorem 10 contradicts the assumption that $\text{PPAD} \not\subseteq \text{TIME}(n^{\text{polylog}(n)})$. \blacksquare

We use the following facts in the proof.

Fact 20 *Let X, Y be finite domains. Let $p, q \in \Delta(X \times Y)$ be two probability distributions over $X \times Y$. Define $p(x) := \sum_{y \in Y} p(x, y)$ be the marginal probability of $x \in X$ and $p(y|x) = p(x, y)/p(x)$ be the conditional probability of $y \in Y$ on $x \in X$. Suppose $p(x) = q(x)$ holds for any $x \in X$, then*

$$\text{TV}(p; q) = \sum_{x \in X} p(x) \cdot \text{TV}(p(\cdot|x); q(\cdot|x)).$$

Proof We have

$$\begin{aligned} \text{TV}(p; q) &= \frac{1}{2} \sum_{(x,y) \in X \times Y} |p(x, y) - q(x, y)| = \sum_{x \in X} \left(\frac{1}{2} \sum_{y \in Y} |p(x, y) - q(x, y)| \right) \\ &= \sum_{x \in X} p(x) \left(\frac{1}{2} \sum_{y \in Y} |p(y|x) - q(y|x)| \right) = \sum_{x \in X} p(x) \cdot \text{TV}(p(\cdot|x); q(\cdot|x)). \end{aligned}$$

The first step follows from the definition of TV distance, the third step holds since $p(x, y) = p(x) \cdot p(y|x)$, $q(x, y) = q(x)q(y|x) = p(x)q(y|x)$, the last step follows the definition of TV distance. \blacksquare

Fact 21 *Let X, Y be finite domains. Let $w \in \Delta(X \times Y)$ be a joint distribution over $X \times Y$, and w_X, w_Y be the marginal distribution of w on X, Y . Let $p = w_X \times w_Y \in \Delta(X) \times \Delta(Y)$ be a product distribution over $X \times Y$ and it has the same marginal distribution as w . Then for any product distribution $q \in \Delta(X) \times \Delta(Y)$, we have*

$$\text{TV}(p; q) \leq 2\text{TV}(w; q)$$

Proof We have

$$\text{TV}(p; q) \leq \text{TV}(w_X; q_X) + \text{TV}(w_Y; q_Y) \leq 2\text{TV}(w; q).$$

■

Appendix C. Missing proof from Section 4

First, we prove

Proof [Proof of Lemma 13] We have

$$\begin{aligned} \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} \sum_{s_i \in S_i} q_{t,\ell}(s_i) r_{t,k}(s_i(k)) &= \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} \sum_{j \in [n]} \sum_{s_i \in S_i, s_i(k)=j} q_{t,\ell}(s_i) r_{t,k}(j) \\ &= \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} \sum_{j \in [n]} w_{t,\ell,k}(j) r_{t,k}(j) \\ &\geq \max_{j \in [n]} \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} r_{t,k}(j) - 2\sqrt{H \log(n)} H^{\ell-1} \\ &= \max_{j \in [n]} \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} r_{t,k}(j) - 2\epsilon H^\ell. \end{aligned}$$

Here the second step holds since the strategy $q_{t,\ell} = w_{t,\ell,1} \times \dots \times w_{t,\ell,K}$ is the product distribution of $\{w_{t,\ell,k}\}_{k \in [K]}$. The third step holds due to the regret guarantee of MWU, the last step holds due to $H = \log(n)/\epsilon^2$. ■

Then we prove

Proof [Proof of Lemma 14] First, we have

$$\begin{aligned} &\sum_{t \in [T]} \sum_{s_i \in S_i} p_t(s_i) r_{t,k}(\phi(s_i)(k)) - p_t(s_i) r_{t,k}(s_i(k)) \\ &= \frac{1}{L} \sum_{\ell \in [L]} \sum_{t \in [T]} \sum_{s_i \in S_i} q_{t,\ell}(s_i) r_{t,k}(\phi(s_i)(k)) - q_{t,\ell}(s_i) r_{t,k}(s_i(k)) \end{aligned} \quad (6)$$

since $p_t = \frac{1}{L} \sum_{\ell \in [L]} q_{t,\ell}$.

For the first term, for each thread $\ell \in [L]$, we have

$$\begin{aligned} \sum_{t \in [T]} \sum_{s_i \in S_i} q_{t,\ell}(s_i) r_{t,k}(\phi(s_i)(k)) &= \sum_{\beta \in [T/H^\ell]} \sum_{h \in [H]} \sum_{t=(\beta-1)H^\ell+(h-1)H^{\ell-1}+1}^{(\beta-1)H^\ell+hH^{\ell-1}} \sum_{s_i \in S_i} q_{t,\ell}(s_i) r_{t,k}(\phi(s_i)(k)) \\ &\leq \sum_{\beta \in [T/H^\ell]} \sum_{h \in [H]} \left\| \sum_{t=(\beta-1)H^\ell+(h-1)H^{\ell-1}+1}^{(\beta-1)H^\ell+hH^{\ell-1}} r_{t,k} \right\|_\infty \end{aligned} \quad (7)$$

where the second step holds since $q_{t,\ell} = w_{t,\ell,1} \times \cdots \times w_{t,\ell,K}$ is fixed during the interval $t \in [(\beta - 1)H^\ell + (h - 1)H^{\ell-1} + 1 : (\beta - 1)H^\ell + hH^{\ell-1}]$.

For the second term in Eq. (6), for each thread $\ell \in [L]$ we have

$$\begin{aligned}
\sum_{t \in [T]} \sum_{s_i \in S_i} q_{t,\ell}(s_i) r_{t,k}(s_i(k)) &= \sum_{\beta \in [T/H^\ell]} \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} \sum_{s_i \in S_i} q_{t,\ell}(s_i) r_{t,k}(s_i(k)) \\
&\geq \sum_{\beta \in [T/H^\ell]} \left(\left\| \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} r_{t,k} \right\|_\infty - 2\epsilon H^\ell \right) \\
&= \sum_{\beta \in [T/H^\ell]} \left\| \sum_{t=(\beta-1)H^\ell+1}^{\beta H^\ell} r_{t,k} \right\|_\infty - 2\epsilon T \tag{8}
\end{aligned}$$

where the second step follows from Lemma 13.

Combining Eq. (6)(7)(8), we have

$$\sum_{t \in [T]} \sum_{s_i \in S_i} p_t(s_i) r_{t,k}(\phi(s_i)(k)) - p_t(s_i) r_{t,k}(s_i(k)) \leq \frac{1}{L} \sum_{t \in [T]} \|r_{t,k}\|_\infty + 2\epsilon T \leq 3\epsilon T$$

Here the second step follows from $L = 1/\epsilon$ and $\|r_{t,k}\|_\infty \leq 1$. ■