

# Bayes correlated equilibria and no-regret dynamics

Kaito Fujii\*

April 11, 2023

## Abstract

This paper explores equilibrium concepts for Bayesian games, which are fundamental models of games with incomplete information. We aim at three desirable properties of equilibria. First, equilibria can be naturally realized by introducing a mediator into games. Second, an equilibrium can be computed efficiently in a distributed fashion. Third, any equilibrium in that class approximately maximizes social welfare, as measured by the price of anarchy, for a broad class of games. These three properties allow players to compute an equilibrium and realize it via a mediator, thereby settling into a stable state with approximately optimal social welfare. Our main result is the existence of an equilibrium concept that satisfies these three properties.

Toward this goal, we characterize various (non-equivalent) extensions of correlated equilibria, collectively known as Bayes correlated equilibria. In particular, we focus on communication equilibria (also known as coordination mechanisms), which can be realized by a mediator who gathers each player's private information and then sends correlated recommendations to the players. We show that if each player minimizes a variant of regret called untruthful swap regret in repeated play of Bayesian games, the empirical distribution of these dynamics converges to a communication equilibrium. We present an efficient algorithm for minimizing untruthful swap regret with a sublinear upper bound, which we prove to be tight up to a multiplicative constant. As a result, by simulating the dynamics with our algorithm, we can efficiently compute an approximate communication equilibrium. Furthermore, we extend existing lower bounds on the price of anarchy based on the smoothness arguments from Bayes Nash equilibria to equilibria obtained by the proposed dynamics.

**The latest version is available at [arXiv:2304.05005](https://arxiv.org/abs/2304.05005).**

---

\*National Institute of Informatics. Email: [fujiik@nii.ac.jp](mailto:fujiik@nii.ac.jp)

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Preliminaries . . . . .	5
1.2	Overview of our results . . . . .	7
1.3	Further related work . . . . .	13
1.4	Organization of this paper . . . . .	16
<b>2</b>	<b>Bayes correlated equilibria and no-regret dynamics</b>	<b>16</b>
2.1	Strategy representability . . . . .	16
2.2	Online learning with stochastic types . . . . .	18
2.3	Strategic-form correlated equilibria . . . . .	20
2.4	Agent-normal-form correlated equilibria . . . . .	23
2.5	Communication equilibria . . . . .	26
2.6	Relations among classes of Bayes correlated equilibria . . . . .	29
<b>3</b>	<b>Algorithm for minimizing untruthful swap regret</b>	<b>34</b>
3.1	Reduction to $\Phi$ -regret minimization . . . . .	35
3.2	Reduction to online linear optimization . . . . .	37
3.3	Decomposition into small external regret minimization problems . . . . .	39
3.4	Application to equilibrium computation . . . . .	42
<b>4</b>	<b>Lower bound for untruthful swap regret</b>	<b>42</b>
4.1	Analysis for randomly branching rewards . . . . .	43
4.2	Assignment of types . . . . .	50
4.3	Analysis of the expected reward in each block . . . . .	53
4.4	Proof of lower bound . . . . .	61
<b>5</b>	<b>Smoothness and price of anarchy</b>	<b>64</b>
5.1	POA bounds for the sum of payoffs . . . . .	65
5.2	POA bounds for conditionally smooth mechanisms . . . . .	67
5.3	POA bounds for type-dependent action sets . . . . .	69
<b>6</b>	<b>Concluding remarks</b>	<b>70</b>
<b>A</b>	<b>On conditional independence property</b>	<b>74</b>
<b>B</b>	<b>Algorithm for minimizing strategy swap regret</b>	<b>75</b>
<b>C</b>	<b>On linear swap regret minimization</b>	<b>78</b>
<b>D</b>	<b>On Bayes coarse correlated equilibria</b>	<b>81</b>

# 1 Introduction

In this paper, we study equilibrium concepts for *Bayesian games*,<sup>1</sup> in which players do not share complete information about the game. Since introduced by Harsanyi [1967, 1968a,b], Bayesian games have been studied as one of the most significant models in game theory.

The best-known equilibrium concept for Bayesian games must be a *Bayes Nash equilibrium*, which is a natural generalization of a (mixed) Nash equilibrium in games with complete information. However, this equilibrium concept has a crucial issue concerning efficient computability. Even for a special case with complete information, computing an approximate Nash equilibrium is PPAD-complete [Daskalakis et al., 2009, Chen et al., 2009] and requires exponentially many queries in the number of players [Babichenko, 2016]. Even for two-player Bayesian games with a constant number of actions, computing an approximate Bayes Nash equilibrium is PPAD-complete [Rubinstein, 2018]. This is why players cannot be expected to reach a Bayes Nash equilibrium in practical situations.

A possible remedy for this issue is to introduce a *mediator* to the game. For games with complete information, if a mediator can recommend (possibly correlated) actions to players, a broader class of equilibria called *correlated equilibria* is a natural equilibrium concept [Aumann, 1974]. Correlated equilibria allow players’ actions to be correlated via the mediator’s recommendation, while their actions must be independently decided in a Nash equilibrium. An example of correlated equilibria in our society is realized by traffic signals, which randomly recommend cars in one direction to “proceed” and cars in the other direction to “stop.” As we can see from this example, players can coordinate and achieve a better outcome in a correlated equilibrium.

In contrast to Nash equilibria, correlated equilibria can be efficiently computed even for multi-player non-zero-sum games. A standard approach to computing a correlated equilibrium is to simulate a variant of *no-regret dynamics*, in which players repeat the same game many times and independently decide their own action in each round. If each player decides their action according to an algorithm that exhibits no internal regret [Foster and Vohra, 1997, Hart and Mas-Colell, 2000] or swap regret [Blum and Mansour, 2007], the dynamics are known to converge to a correlated equilibrium.<sup>2</sup> An advantage of this approach is that the dynamics are *uncoupled*, that is, they can be simulated by independent players in a distributed fashion. Furthermore, the existence of such dynamics can explain why people automatically reach an equilibrium in our society. No-regret dynamics in Bayesian games are known to converge to a *Bayes coarse correlated equilibrium*<sup>3</sup> [Hartline et al., 2015]. However, this equilibrium concept is usually unnatural because it is assumed that once observing the recommendation by the mediator, each player must follow it, even if it is more beneficial to deviate. The convergence of a variant of no-regret dynamics to more natural equilibrium concepts is still unknown.

In addition to efficient computability, another desirable property of equilibrium concepts is that any equilibrium in that class approximately maximizes social welfare. As the famous example of the prisoners’ dilemma illustrates, players can fall into a bad equilibrium by trying to maximize their own

---

<sup>1</sup>Some existing studies do not distinguish “games with incomplete information” and “Bayesian games.” In this paper, we always assume that players share a common prior distribution that generates a type profile. To put emphasis on this assumption, we use the term “Bayesian games” throughout this paper.

<sup>2</sup>Formally, the empirical distribution of action profiles in all past rounds converges to a correlated equilibrium. In this paper, we say “dynamics converge” in this sense.

<sup>3</sup>Formally, Bayes coarse correlated equilibria have various (non-equivalent) definitions. The dynamics proposed by Hartline et al. [2015] converge to an agent-normal-form coarse correlated equilibrium. See Section D for formal definitions of Bayes coarse correlated equilibria.

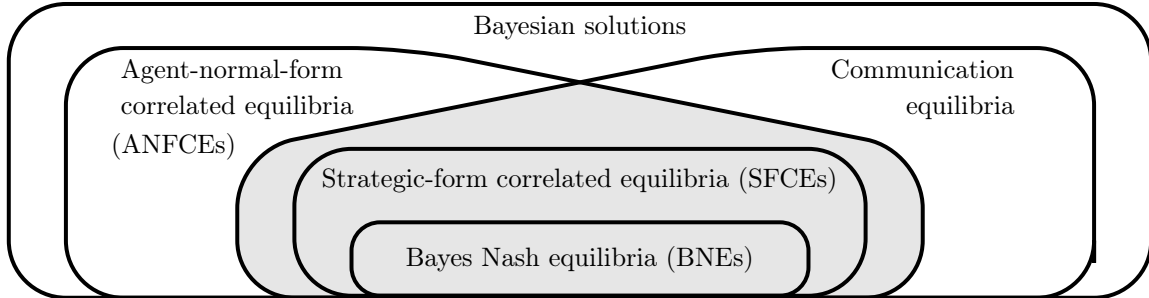


Figure 1: Relations among classes of Bayes correlated equilibria. Formally, SFCEs and ANFCEs in this figure are their counterparts in the space of type-wise distributions (see Section 2.3 and Section 2.4). The shaded region represents communication equilibria with strategy representability (equivalently, the intersection of communication equilibria and ANFCEs). This equilibrium concept can be realized by a mediator acting as a communication device, efficiently computable via dynamics minimizing untruthful swap regret, and guaranteed to have price-of-anarchy lower bounds for (conditionally) smooth games.

payoffs. To measure how bad equilibria could be, Koutsoupias and Papadimitriou [1999] proposed the *price of anarchy* (POA), which is defined as the ratio of the social welfare in the worst equilibrium to the maximum social welfare. An argument using *smoothness* is a versatile framework for bounding the POA for a broad class of games. For games with complete information, Roughgarden [2015a] proved that smoothness of games implies a lower bound on the POA of correlated equilibria,<sup>4</sup> and various classes of games satisfy smoothness. Syrgkanis and Tardos [2013] later developed a variant of smoothness for mechanism design and provided lower bounds on the POA of correlated equilibria for a large class of mechanisms. For Bayesian games, Roughgarden [2015b] and Syrgkanis [2012] derived bounds on the POA of Bayes Nash equilibria by using the smoothness arguments, but it remains to be seen whether these bounds can be extended to a broader class of equilibria. There are several exceptions that considered Bayes coarse correlated equilibria [Caragiannis et al., 2015, Hartline et al., 2015], but their applications are limited (see Section 1.3 for more detailed discussions).

If there exists an equilibrium concept that satisfies these desirable properties, we can compute an equilibrium efficiently in a distributed fashion and then realize it by introducing a mediator, which enables players to settle into an equilibrium with high social welfare. Thus, the following question regarding Bayesian games naturally arises.

*Is there any equilibrium concept for Bayesian games that can be realized by a mediator, efficiently computable via uncoupled dynamics, and guaranteed to have POA bounds for various games?*

To answer this question, we investigate the concept of *communication equilibria*, which can be realized by introducing a mediator to Bayesian games [Myerson, 1982, Forges, 1986]. In Bayesian games, each player has private information represented by a *type* generated from a commonly known prior distribution. To deal with this private information, a mediator in a communication equilibrium communicates with players bidirectionally. A mediator first gathers the types from players

<sup>4</sup>More precisely, Roughgarden [2015a] proved this result for a broader class of equilibria called *coarse correlated equilibria*. However, this equilibrium concept is unnatural in many situations, as is the concept of Bayes coarse correlated equilibria.

and then sends recommendations to each player based on the received types. This is an equilibrium in the sense that players have no incentive to deviate from reporting the type truthfully or following the action recommendation. Note that Myerson [1982] originally called communication equilibria *coordination mechanisms* with an emphasis on their relations to Bayesian incentive-compatible mechanisms. Although these two concepts are similar in that truthful type reporting is incentive-compatible, a mediator in communication equilibria is required to recommend incentive-compatible actions to players, while a mechanism determines an outcome in a top-down fashion. A communication equilibrium can be interpreted as a solution concept combining truthful type reporting of Bayesian incentive-compatible mechanisms and action recommendations of correlated equilibria. These two incentive constraints complicate the structure of this concept and make its computation challenging.

Our central result is an efficient algorithm for computing a communication equilibrium. To derive this algorithm, we first show that if each player minimizes a variant of regret, which we call *untruthful swap regret*, the dynamics converge to a communication equilibrium. We then develop an efficient algorithm for minimizing untruthful swap regret with a sublinear upper bound. We also prove a lower bound that shows the tightness of the upper bound. As a byproduct, we characterize various extensions of correlated equilibria, collectively called *Bayes correlated equilibria* [Forges, 1993, 2014] (Figure 1). This reveals an additional property called *strategy representability*, which other extensions have but communication equilibria lack. The limit of dynamics minimizing untruthful swap regret also satisfies strategy representability. Furthermore, we show that most existing POA bounds of Bayes Nash equilibria based on the smoothness arguments can be directly extended to communication equilibria with strategy representability. Thus, the concept of communication equilibria with strategy representability satisfies the desirable properties.

## 1.1 Preliminaries

**Basic notations** The sets of reals and integers are denoted by  $\mathbb{R}$  and  $\mathbb{Z}$ , respectively. For any  $a, b \in \mathbb{R}$ , we define  $[a, b] = \{x \in \mathbb{R} \mid a \leq x \leq b\}$ . For any positive integer  $n \in \mathbb{Z}$ , we define  $[n] = \{i \in \mathbb{Z} \mid 1 \leq i \leq n\} = \{1, 2, \dots, n\}$ . For any finite set  $A$ , we denote the probability simplex by  $\Delta(A) = \{\pi \in [0, 1]^A \mid \sum_{a \in A} \pi(a) = 1\}$ . For any distribution  $\pi \in \Delta(A)$ , we write  $a \sim \pi$  when  $a \in A$  is generated from the distribution  $\pi$ . By abuse of notation, we write  $a \sim A$  when  $a$  is generated from the uniform distribution over  $A$ . To specify a random variable that a probability or an expected value is concerned with, we write  $\Pr_{a \sim \pi}(\cdot)$  or  $\mathbb{E}_{a \sim \pi}[\cdot]$ . If the distribution is obvious, we simply write  $\Pr_a(\cdot)$  or  $\mathbb{E}_a[\cdot]$ . For vectors  $x, y \in \mathbb{R}^n$ , the inner product is written as  $\langle x, y \rangle = \sum_{i=1}^n x(i)y(i)$ . Similarly, for matrices  $P, Q \in \mathbb{R}^{n \times m}$ , the matrix inner product is written as  $\langle P, Q \rangle = \sum_{i=1}^n \sum_{j=1}^m P(i, j)Q(i, j)$ .

**Bayesian games** Let  $n \in \mathbb{Z}$  be the number of players and  $N = [n]$  a set of players. Let  $A_i$  be a finite set of actions for player  $i \in N$  and  $A = A_1 \times \dots \times A_n$  the set of action profiles of all players. Let  $\Theta_i$  be a set of finite possible types for each player  $i \in N$  and  $\Theta = \Theta_1 \times \dots \times \Theta_n$  the set of type profiles of all players. Let  $v_i: \Theta \times A \rightarrow [0, 1]$  be a utility function that maps a type profile and action profile to a payoff value for player  $i$ .<sup>5</sup> Since the size of  $\Theta$  and  $A$  is exponential in  $n$ , we

<sup>5</sup>If each player's type represents their own preference, it is more natural to assume that their payoff depends only on their own type, not the other players' types. In this case, the utility function is defined to be  $v_i: \Theta_i \times A \rightarrow [0, 1]$  whose value is determined by their own type  $\theta_i$ , not by the other player's types  $\theta_{-i}$ . We make this assumption to obtain POA bounds in Section 5.

assume that each  $v_i$  is given by an oracle that returns the value of  $v_i(\theta; a)$  for any  $\theta \in \Theta$  and  $a \in A$ . For each player  $i \in N$ , subscript  $-i$  represents the other players  $N \setminus \{i\}$ , i.e.,  $a_{-i} = (a_j)_{j \in N \setminus \{i\}}$  and  $\theta_{-i} = (\theta_j)_{j \in N \setminus \{i\}}$ .

In this paper, we distinguish between the two words “strategies” and “actions.” While actions are defined as above, strategies are decisions that determine an action for every type. Formally, a strategy  $s_i: \Theta_i \rightarrow A_i$  of player  $i \in N$  determines action  $s_i(\theta_i) \in A_i$  to be selected given their own type  $\theta_i \in \Theta_i$ . For each player  $i \in N$ , let  $S_i = A_i^{\Theta_i}$  be the set of all strategies. We write a strategy profile as  $s = (s_1, \dots, s_n)$ , and let  $S = S_1 \times \dots \times S_n$  be the set of all strategy profiles. For notational simplicity, we write  $s(\theta) = (s_j(\theta_j))_{j \in N}$  and  $s_{-i}(\theta_{-i}) = (s_j(\theta_j))_{j \in N \setminus \{i\}}$  for each  $s \in S$ ,  $\theta \in \Theta$ , and  $i \in N$ .

In Bayesian games, we assume that the type profile is generated from the probability distribution  $\rho \in \Delta(\Theta)$  that every player knows in advance. Since the size of  $\Theta$  is exponential in  $n$  in general, the distribution  $\rho$  cannot be expressed in polynomial size. In most parts of this paper, we assume that it is possible to efficiently sample  $\theta_i$  from the marginal distribution of  $\rho$ , which is denoted by  $\rho_i$ , and  $\theta_{-i}$  from the distribution  $\rho$  conditioned on any  $\theta_i \in \Theta_i$ , which is denoted by  $\rho|\theta_i$ . This assumption holds for a special case where  $\rho$  is a product distribution, i.e.,  $\rho(\theta) = \prod_{i \in N} \rho_i(\theta_i)$  holds for each  $\theta \in \Theta$ . For POA bounds in Section 5, we assume that  $\rho$  is a product distribution.

A Bayesian game proceeds in the following steps: (1) First, a type profile  $\theta$  is generated from  $\rho \in \Delta(\Theta)$ , and each player  $i \in N$  is notified of their own type  $\theta_i \in \Theta_i$ . (2) Each player then decides their action  $a_i \in A_i$  without knowing the other players’ types  $\theta_{-i}$ . (3) Finally, each player obtains the payoff  $v_i(\theta; a)$  depending on all players’ actions and types.

There are two widely-used interpretations of Bayesian games: the strategic form (also called the random-vector model or the normal form) and the agent normal form (also called the posterior-lottery model, the Selten model, or the population interpretation). Both interpretations reduce a Bayesian game to a normal-form game with complete information as follows.

- In the strategic form of a Bayesian game, the strategies  $S_i$  are considered as a decision space for player  $i \in N$ . Each player  $i \in N$  obtains the expected payoff  $\sum_{\theta \in \Theta} \rho(\theta) v_i(\theta; s(\theta))$  for strategy profile  $s \in S$ .
- In the agent normal form, we hypothetically consider the same player with different types as different players, that is, the set of players are  $\{(i, \theta_i) \mid i \in N, \theta_i \in \Theta_i\}$ , which we denote by  $N'$ . Each player  $(i, \theta_i) \in N'$  selects an action from  $A_i$ . A type profile  $\theta \in \Theta$  is randomly generated from  $\rho$ , and then only one player  $(i, \theta_i)$  becomes active for each  $i \in N$ . Only the active players can obtain non-zero payoffs, which are determined only by the active players’ actions. The inactive players always obtain payoff 0 and do not affect the other players’ payoffs. If we denote player  $(i, \theta_i)$ ’s action by  $a_{i, \theta_i}$ , the expected payoff for player  $(i, \theta_i)$  is  $\sum_{\theta' \in \Theta: \theta'_i = \theta_i} \rho(\theta') v_i(\theta'; (a_{j, \theta'_j})_{j \in N})$ .

**Correlated equilibria** A normal-form game with complete information is specified by players  $N = \{1, 2, \dots, n\}$ , action profiles  $A = \prod_{i \in N} A_i$ , and a utility function  $v_i: A \rightarrow [0, 1]$  for each player  $i \in N$ . In this game, a distribution  $\pi \in \Delta(A)$  is an  $\epsilon$ -approximate correlated equilibrium if

$$\mathbb{E}_{a \sim \pi} [v_i(a)] \geq \mathbb{E}_{a \sim \pi} [v_i(\phi(a_i), a_{-i})] - \epsilon.$$

holds for any  $i \in N$  and  $\phi: A_i \rightarrow A_i$ , where  $\epsilon \in [0, 1]$  represents an additive error. In an  $\epsilon$ -approximate correlated equilibrium, each player has gain at most  $\epsilon$  in expectation by taking an

Table 1: A comparison of the classes of Bayes correlated equilibria (or their counterparts in  $\Delta(A)^\ominus$ ). SR stands for “strategy representability” defined in Section 2.1.

Class	Characterization
Strategic-form correlated equilibria	Incentive constraints w.r.t. strategies ( $\text{IC}'_{\text{SF}}$ ) & SR
Agent-normal-form correlated equilibria	Incentive constraints w.r.t. actions ( $\text{IC}'_{\text{ANF}}$ ) & SR
Communication equilibria	Incentive constraints w.r.t. types and actions ( $\text{IC}_{\text{Com}}$ )
Bayesian solutions	Incentive constraints w.r.t. actions ( $\text{IC}'_{\text{ANF}}$ )

action  $\phi(a_i) \in A_i$  instead of the recommended action  $a_i \in A_i$ .

Suppose that the same players play this game  $T$  times and decide a randomized action in each round. Let  $\pi_i^t \in \Delta(A_i)$  be the decision by each player  $i \in N$  for each round  $t \in [T]$ . Then the payoff obtained by each player  $i \in N$  in each round  $t \in [T]$  is  $\mathbb{E}_{a \sim \pi^t} [v_i(a)]$ , where  $\pi^t \in \Delta(A)$  is the product distribution that independently generates  $a_j \sim \pi_j^t$  for each  $j \in N$ . This payoff can be expressed as  $\mathbb{E}_{a_i \sim \pi_i^t} [u_i^t(a_i)]$  if we set  $u_i^t(a_i) = \mathbb{E}_{a_{-i} \sim \pi_{-i}^t} [v_i(a)]$  for each action  $a_i \in A_i$ , where  $\pi_{-i}^t \in \Delta(A_{-i})$  is the product distribution that independently generates  $a_j \sim \pi_j^t$  for each  $j \in N \setminus \{i\}$ . Thus, the online decision-making problem that each player  $i \in N$  faces can be regarded as an online learning problem with decision space  $A_i$  and rewards  $(u_i^t)_{t=1}^T$ . For this problem, swap regret is defined as

$$R_{\text{S},i} = \max_{\phi: A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{a_i \sim \pi_i^t} [u_i^t(\phi(a_i))] - \sum_{t=1}^T \mathbb{E}_{a_i \sim \pi_i^t} [u_i^t(a_i)].$$

This quantity expresses regret for how much the player would have gained by playing  $\phi(a_i)$  instead of each  $a_i \in A_i$ . Blum and Mansour [2007] proved that if each player minimizes their swap regret, the empirical distribution of the dynamics for  $T$  rounds is a  $(\max_{i \in N} R_{\text{S},i}/T)$ -approximate correlated equilibrium. Blum and Mansour [2007] also proposed an efficient algorithm that achieves  $O(\sqrt{T|A_i| \log |A_i|})$  swap regret.

## 1.2 Overview of our results

**Bayes correlated equilibria and no-regret dynamics** Our first goal is to define untruthful swap regret and prove that dynamics minimizing it converge to a communication equilibrium with strategy representability. In the process leading to this, we define approximate versions of Bayes correlated equilibria and consider extensions of no-swap-regret dynamics.

To discuss the convergence of dynamics, we need to characterize each class as a set of distributions (Table 1). For this purpose, we first clarify the difference between  $\Delta(S)$  and  $\Delta(A)^\ominus$  by considering two corresponding scenarios. In the first scenario, a mediator generates  $s \sim \sigma \in \Delta(S)$  without observing type profile  $\theta \sim \rho$ , and then each player  $i \in N$  takes action  $s_i(\theta_i)$ . In the second scenario, a mediator generates  $a \sim \pi(\theta) \in \Delta(A)$  based on type profile  $\theta \in \Theta$ . Thus, given a type profile distribution  $\rho \in \Delta(\Theta)$ , both distributions  $\sigma \in \Delta(S)$  and  $\pi \in \Delta(A)^\ominus$  determine a distribution over action profiles  $A$  for each  $\theta \in \Theta$ . However, not all distributions obtained from  $\Delta(A)^\ominus$  can be realized by  $\Delta(S)$ . There exists a simple example  $\pi \in \Delta(A)^\ominus$  that cannot be represented by any distribution  $\sigma \in \Delta(S)$  (see Theorem 2.2 and Figure 3). Strategic-form correlated equilibria (SFCEs) and agent-normal-form correlated equilibria (ANFCEs) are defined as a subset of  $\Delta(S)$ , while communication equilibria and Bayesian solutions are defined as a subset of  $\Delta(A)^\ominus$ . To compare these

classes, we consider the set of distributions in  $\Delta(A)^\Theta$  that can be realized by SFCEs or ANFCEs, and we call  $\pi \in \Delta(A)^\Theta$  strategy-representable if there exists some  $\sigma \in \Delta(S)$  that yields the same distribution over action profiles for each  $\theta \in \Theta$  (Theorem 2.1).

SFCEs and ANFCEs are defined as correlated equilibria of the strategic form and agent normal form, respectively. Their incentive constraints are natural extensions of those for correlated equilibria. In an SFCE, each player  $i \in N$  cannot gain by selecting strategy  $\phi_{\text{SF}}(s_i) \in S_i$  instead of the recommended strategy  $s_i \in S_i$  according to any  $\phi_{\text{SF}}: S_i \rightarrow S_i$ , while in an ANFCE, by selecting action  $\phi(\theta_i, a_i) \in A_i$  instead of the recommended action  $a_i \in A_i$  when the type is  $\theta_i \in \Theta_i$  according to any  $\phi: \Theta_i \times A_i \rightarrow A_i$ . No-swap-regret dynamics converging to a correlated equilibrium can be directly extended to these classes. Note that this naive extension does not immediately imply efficient computability of an SFCE because it requires swap regret minimization with decision space  $S_i$ , whose size is exponential in  $|\Theta_i|$ . See Section 2.3 for SFCEs and Section 2.4 for ANFCEs.

Our main focus is the notion of communication equilibria. The incentive constraints for communication equilibria are specified by two different kinds of deviations: reporting an untruthful type  $\psi(\theta_i) \in \Theta_i$  instead of the true type  $\theta_i \in \Theta_i$  according to  $\psi: \Theta_i \rightarrow \Theta_i$  and taking an action instead of the recommended action according to  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Formally, we say  $\pi \in \Delta(A)^\Theta$  is an  $\epsilon$ -approximate communication equilibrium if the incentive constraint

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\psi(\theta_i), \theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon. \quad (\text{IC}_{\text{Com}})$$

holds for each  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . As it allows untruthful type reporting, this constraint is stronger than that for ANFCEs. On the other hand, ANFCEs, defined on  $\Delta(S)$ , do not include communication equilibria, defined on  $\Delta(A)^\Theta$ .

To deal with these two kinds of deviations, we introduce a new variant of regret called untruthful swap regret. In repeated play of Bayesian games, each player faces an online learning problem with rewards determined by a stochastic type. In each round  $t \in [T]$  of this problem, each player  $i \in N$  decides a distribution  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$  that determines a (randomized) action for each type and receives reward<sup>6</sup>

$$u_i^t(\theta_i, a_i) = \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{a_{-i} \sim \pi_{-i}^t(\theta_{-i})} [v_i(\theta; a)] \right]$$

depending on their realized type  $\theta_i \sim \rho_i$  and action  $a_i \sim \pi_i^t(\theta_i)$ , where  $\pi_{-i}^t(\theta_{-i})$  is the product distribution that independently generates  $a_j \sim \pi_j^t(\theta_j)$  for each  $j \in N \setminus \{i\}$ . For this online learning problem, the untruthful swap regret  $R_{\text{US},i}^T$  for player  $i \in N$  is defined as

$$R_{\text{US},i}^T = \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta_i, \phi(\theta_i, a_i))] - \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, a_i)] \right].$$

Note that this definition coincides with swap regret when  $|\Theta| = 1$ . We show that if every player achieves sublinear untruthful swap regret, then the dynamics converge to a communication equilibrium. Furthermore, the empirical distribution of uncoupled dynamics is always strategy-representable (Theorem 2.4).

<sup>6</sup>We use the term ‘‘reward’’ for online learning problems and distinguish it from the term ‘‘payoff’’ for games.

**Theorem 1.1** (Informal version of Theorem 2.17). *The empirical distribution of the dynamics for  $T$  rounds is a  $\max_{i \in N} R_{\text{US},i}/T$ -approximate communication equilibrium with strategy representability, where  $R_{\text{US},i}$  is the untruthful swap regret for each player  $i \in N$ .*

Note that the set of all communication equilibria with strategy representability equals the intersection of communication equilibria and ANFCEs (Theorem 2.18). We present several other results: we clarify the relations between these classes as well as Bayesian solutions and Bayes Nash equilibria (see Section 2.6), we slightly improve an upper bound on strategy swap regret from the direct application of the existing swap regret minimization algorithm (Section B), and we provide various definitions of Bayes coarse correlated equilibria for comparison (Section D).

**Algorithm for minimizing untruthful swap regret** To approximately compute a communication equilibria with strategy representability, we propose an efficient algorithm for minimizing untruthful swap regret. We construct our algorithm in the following three steps to obtain an upper bound of the optimal order.

In the first step (Section 3.1), the problem of minimizing untruthful swap regret is interpreted as an instance of  $\Phi$ -regret minimization. To our knowledge, the definition of  $\Phi$ -regret was first provided by Greenwald and Jafari [2003] when the decision space is a probability simplex and then extended to a general decision space by Stoltz and Lugosi [2007]. For a set of transformations  $\Phi$ , the  $\Phi$ -regret measures how the algorithm gains by transforming their decisions with an optimal  $\phi \in \Phi$  in hindsight. To interpret untruthful swap regret as  $\Phi$ -regret, we change the decision space from the set of distributions for each type  $\Delta(A_i)^{\Theta_i}$  to the set of vectors  $\mathcal{X}$  defined by

$$\mathcal{X} = \left\{ x \in [0, 1]^{\Theta_i \times A_i} \mid \sum_{a_i \in A_i} x(\theta_i, a_i) = 1 \ (\forall \theta_i \in \Theta_i) \right\},$$

where we set  $x^t(\theta_i, a_i) = \pi_i^t(\theta_i, a_i)$  for each  $t \in [T]$ ,  $\theta_i \in \Theta_i$ , and  $a_i \in A_i$ . Then the space of all transformations expressed by  $\psi: \Theta_i \rightarrow \Theta_i$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$  can be written as  $\mathcal{Q}$  defined by

$$\mathcal{Q} = \left\{ Q \in [0, 1]^{(\Theta_i \times A_i) \times (\Theta_i \times A_i)} \mid \begin{array}{l} \text{there exists some } W \in [0, 1]^{\Theta_i \times \Theta_i} \text{ such that} \\ \sum_{\theta'_i \in \Theta_i} W(\theta_i, \theta'_i) = 1 \ (\forall \theta_i \in \Theta_i) \text{ and} \\ \sum_{a_i \in A_i} Q((\theta_i, a_i), (\theta'_i, a'_i)) = W(\theta_i, \theta'_i) \ (\forall \theta_i, \theta'_i \in \Theta_i, a'_i \in A_i) \end{array} \right\},$$

where  $W$  represents a convex combinations of all possible  $\psi$  and each block matrix  $Q_{\theta_i, \psi(\theta_i)}$  represents a convex combination of all possible  $\phi(\theta_i, \cdot)$ . Then we obtain

$$R_{\text{US},i} = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Qx^t, \bar{u}^t \rangle - \sum_{t=1}^T \langle x^t, \bar{u}^t \rangle,$$

where  $\bar{u}^t$  is the reward vector weighted by prior probabilities  $\rho_i$ , i.e.,  $\bar{u}^t(\theta_i, a_i) = \rho_i(\theta_i)u_i^t(\theta_i, a_i)$  for each  $\theta_i \in \Theta_i$  and  $a_i \in A_i$  (Theorem 3.1).

In the second step (Section 3.2), we apply an efficient reduction framework that reduces  $\Phi$ -regret minimization to online linear optimization with decision space  $\Phi$ , which was proposed by Gordon et al. [2008]. This framework internally uses fixed point computation of each  $Q \in \mathcal{Q}$ . We show that the fixed point  $x \in \mathcal{X}$  such that  $Qx = x$  can be obtained by using eigenvector computation. Now our problem is reduced to online linear optimization with decision space  $\mathcal{Q}$ . Since for online linear

optimization with a polytope constraint, there exist efficient algorithms with sublinear external regret (e.g., Follow-the-Perturbed-Leader [Kalai and Vempala, 2005] or Component Hedge [Koolen et al., 2010]), this problem is already tractable. However, they do not provide an upper bound of the optimal order. We thus need the following third step.

In the third step (Section 3.3), we decompose online linear optimization with decision space  $\mathcal{Q}$  into small external regret minimization problems. Since  $\mathcal{Q}$  can be regarded as a product of block stochastic matrices and a stochastic matrix in some sense, it can be decomposed into  $|\Theta_i|^2|A_i| + |\Theta_i|$  probability simplices. Then we can apply algorithms for external regret minimization with a probability simplex. The final algorithm is described in Algorithm 3, whose regret upper bound is the following.

**Theorem 1.2** (Informal version of Theorem 3.8). *There exists an efficiently computable algorithm with  $O(\sqrt{T} \max\{\log |\Theta_i|, |A_i| \log |A_i|\} + |A_i| \log |A_i|)$  untruthful swap regret.*

As a corollary, we obtain an algorithm that computes an approximate communication equilibrium with strategy representability in polynomial time with high probability (Section 3.4). This theorem also answers the problem of minimizing *linear swap regret* proposed by Mansour et al. [2022]. See Section 1.3 and Section C for more details on linear swap regret.

**Lower bound for untruthful swap regret** For swap regret, Ito [2020] already proved an  $\Omega(\sqrt{T}|A_i| \log |A_i|)$  lower bound, which can apply to untruthful swap regret. To show the tightness of the leading term in our upper bound, it is sufficient to prove an  $\Omega(\sqrt{T} \log |\Theta_i|)$  lower bound.

The proof idea for the swap regret lower bounds [Blum and Mansour, 2007, Ito, 2020] can be summarized as follows. A problem instance is constructed such that the reward for each action  $a_i \in A_i$  in each round  $t \in [T]$  is independently chosen from the uniform distribution over  $\{0, 1\}$ . We can assume that the algorithm selects most of the actions  $A_i$  for at least  $\Omega(T/|A_i|)$  rounds (This can be guaranteed by “blocking” operations, which set the rewards of actions selected more than a certain number of times to 0 for the remaining rounds). For the rounds when the algorithm selects such an action  $a_i \in A_i$ , the gap between the total expected rewards for  $a_i$  and an optimal action  $\phi(a_i) \in A_i$  is  $\Omega\left(\sqrt{\frac{T}{|A_i|} \log |A_i|}\right)$  due to an anti-concentration bound for binomial distributions. The swap regret is at least the sum of  $\Omega\left(\sqrt{\frac{T}{|A_i|} \log |A_i|}\right)$  for these actions, hence  $\Omega(\sqrt{T}|A_i| \log |A_i|)$ .

Based on a similar idea, for untruthful swap regret, we can easily prove a lower bound if the algorithm’s decisions are guaranteed to be sufficiently different for different types. For example, we consider a problem instance with two actions  $A_i = \{\alpha_0, \alpha_1\}$ , in which for each type  $\theta_i \in \Theta_i$  and round  $t \in [T]$ , the reward for  $\alpha_0$  is chosen from  $\{0, 1\}$  uniformly at random, and the reward for  $\alpha_1$  is decided by flipping the corresponding reward for  $\alpha_0$ . Then the total expected reward for any algorithm is  $T/2$ . Assume that the algorithm’s decisions are completely different for different types. For example, we assume that for each type  $\theta_i \in \Theta_i$  and round  $t \in [T]$ , the algorithm chooses  $\alpha_0$  or  $\alpha_1$  uniformly at random. Recall that the competitor of untruthful swap regret can apply any type swap  $\psi: \Theta_i \rightarrow \Theta_i$ . Since for each pair of  $\theta_i$  and  $\psi(\theta_i) \in \Theta_i$ , the expected total reward obtained by applying the decision for  $\psi(\theta_i)$  to the rewards for  $\theta_i$  follows a binomial distribution, the optimal reward is  $T/2 + \Omega(\sqrt{T} \log |\Theta_i|)$  again from an anti-concentration bound, which leads to a desired lower bound.

However, it is not easy to guarantee that the algorithm’s decisions are sufficiently different for different types. For any payoffs, if the algorithm’s decisions are completely the same for all types,

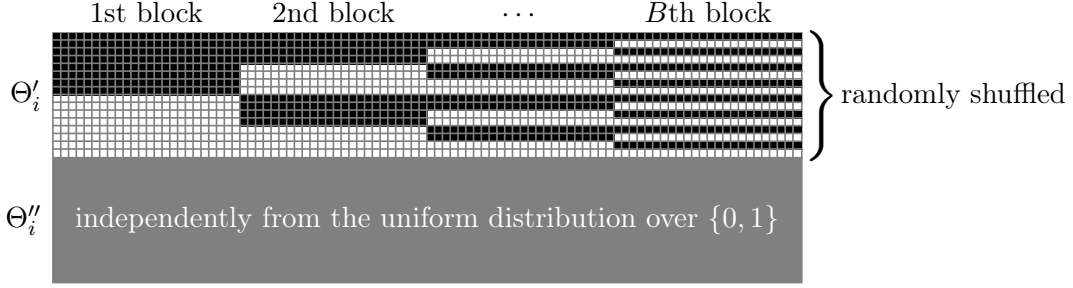


Figure 2: A problem instance for proving a lower bound on untruthful swap regret. The types  $\Theta_i$  are partitioned into  $\Theta'_i$  and  $\Theta''_i$ . The rewards for  $\Theta'_i$  are constant within each block of rounds and randomly branch at the beginning of each block. The reward for  $\alpha_0$  is 1 for the black cells and 0 for the white cells. For each  $\theta''_i \in \Theta''_i$ , the rewards for  $\alpha_0$  are generated independently. Every reward for  $\alpha_1$  is determined by flipping the corresponding reward for  $\alpha_0$ .

$\psi$  does not yield any untruthful swap regret. We thus need to consider a combination of  $\psi$  and  $\phi$ . For example, if there is a type  $\theta_i \in \Theta_i$  for which the payoff for  $\alpha_0$  is always better than  $\alpha_1$ , the algorithm suffers untruthful swap regret due to  $\phi(\theta_i, \alpha_0) = \phi(\theta_i, \alpha_1) = \alpha_0$  unless always choosing  $\alpha_0$  (as formally stated in Theorem 4.2).

Based on this observation, we create a problem instance with  $A_i = \{\alpha_0, \alpha_1\}$  and the uniform distribution  $\rho_i$  (Figure 2). We assume  $|\Theta_i| = 2^{B+1}$  for some  $B \in \mathbb{Z}$  for simplicity. The set of types  $\Theta_i$  is partitioned into  $\Theta'_i$  and  $\Theta''_i$  of equal size. The set of rounds  $[T]$  is partitioned into  $B$  blocks of equal length. The rewards for  $\Theta''_i$  are randomly determined as in the problem instance described above. Each  $\theta'_i \in \Theta'_i$  is associated with a binary sequence of length  $B$  by a random bijection  $\zeta: \Theta'_i \rightarrow \{0, 1\}^B$ , and the  $b$ th bit of this sequence indicates which of  $\alpha_0$  or  $\alpha_1$  has reward 1 for  $\theta'_i$  in the  $b$ th block.

The rewards for  $\Theta'_i$  are designed so that the algorithm must use different decisions for different types. Let  $\theta'_i{}^0, \theta'_i{}^1 \in \Theta'_i$  be the types for which  $\alpha_0$ 's and  $\alpha_1$ 's rewards are always 1 (the completely black row and completely white row), respectively. As mentioned above,  $\phi(\theta_i, \alpha_0) = \phi(\theta_i, \alpha_1) = \alpha_0$  forces the algorithm to almost always select  $\alpha_0$  for the type  $\theta'_i{}^0$ . Similarly, the algorithm almost always selects  $\alpha_1$  for the type  $\theta'_i{}^1 \in \Theta'_i$ . Since the rewards for  $\Theta'_i$  are randomly shuffled, the algorithm does not know which are  $\theta'_i{}^0$  and  $\theta'_i{}^1$  before reaching the final block. In the first block, since the algorithm does not know which type in the black rows is  $\theta'_i{}^0$  and which type in the white rows is  $\theta'_i{}^1$ , the algorithm must almost always choose  $\alpha_0$  for all the black cells and  $\alpha_1$  for all the white cells.

We are tempted to expect that for the second and later blocks, the algorithm also almost always selects  $\alpha_0$  for the black cells and  $\alpha_1$  for the white cells. Unfortunately, there is a counterexample to this expectation. Suppose that the algorithm uses the same action for each contiguous *two* blocks. For example, if the first, second, third, and fourth blocks are black, white, white, and black, respectively, the algorithm uses  $\alpha_0$  for the first two blocks and  $\alpha_1$  for the next two blocks. Then the algorithm successfully selects  $\alpha_0$  for  $\theta'_i{}^0$  and  $\alpha_1$  for  $\theta'_i{}^1$  in almost all rounds. Furthermore, even though the algorithm obtains the total reward only  $T/2$  for some type, this is optimal among all decisions that select the same action for each contiguous two blocks, and then the algorithm suffers no untruthful swap regret.<sup>7</sup> We prove that this example is tight, that is, if the untruthful swap

<sup>7</sup>Note that a similar algorithm focusing on each contiguous *three* blocks yields  $\Omega(T)$  untruthful swap regret. For example, this algorithm always selects  $\alpha_0$  for the row with black, white, and white cells for the first, second, and

regret is small, then the difference of the algorithm’s decisions between the “adjacent” black and white rows must be at least  $T/2$  approximately (see Theorem 4.3 for a formal statement). By using this fact, we show that the algorithm’s decisions must be significantly different at each branching point with some constant probability (Theorem 4.4).

If the algorithm always selects  $\alpha_0$  for the black cells and  $\alpha_1$  for the white cells, by choosing a better action separately for  $\theta''_i \in \Theta''_i$  in each block, we can show that the algorithm suffers  $B \cdot \Omega(\sqrt{T/B}) = \Omega(\sqrt{T \log |\Theta_i|})$  untruthful swap regret. Since we can control the algorithm’s decisions only for a constant fraction of all rounds, we need to carefully determine  $\psi$  so that the algorithm cannot manipulate the untruthful swap regret (Section 4.2) and trace the difference of cumulative rewards by using martingale analysis (Section 4.3). These analyses lead to the following desired lower bound (Section 4.4).

**Theorem 1.3** (Informal version of Theorem 4.1). *For any algorithm, there exists a problem instance such that the untruthful swap regret is  $\Omega(\sqrt{T \log |\Theta_i|})$ .*

**Smoothness and price of anarchy** Given an equilibrium class  $\Pi \subseteq \Delta(A)^\Theta$ , the price of anarchy (POA) of a Bayesian game is defined as

$$\text{POA}_\Pi = \frac{\inf_{\pi \in \Pi} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right]}{\mathbb{E}_{\theta \sim \rho} \left[ \max_{a \in A} v_{\text{SW}}(\theta; a) \right]},$$

where  $v_{\text{SW}}: \Theta \times A \rightarrow [0, 1]$  is the social welfare function. For analyzing the POA, we make two assumptions:  $\rho$  is a product distribution and the value of  $v_i(\theta; a)$  does not depend on  $\theta_{-i}$  for each  $i \in N$ . These two assumptions are natural when a type represents a preference or attribute of each player. They were assumed in existing studies [Syrkkanis, 2012, Syrgkanis and Tardos, 2013, Roughgarden, 2015b].

Our analysis is based on smoothness, which is a standard assumption for bounding the POA of games with complete information. Intuitively, smoothness is a condition that for any action profile, there exists some deviation that significantly increases the social welfare toward the optimum. By assuming  $(\lambda, \mu)$ -smoothness, Roughgarden [2015a] proved that the POA of coarse correlated equilibria is bounded below by  $\lambda/(1 + \mu)$ . Furthermore, Roughgarden [2015a] showed that a broad class of games satisfies smoothness.

The smoothness argument cannot be extended directly to Bayesian games, even for Bayes Nash equilibria. Note that the POA of Bayesian games does not coincide with the POA of the strategic form or agent normal form. This is because the denominator takes an optimal action profile for each  $\theta \in \Theta$ , which can be better than the expected social welfare for the optimal strategy profile  $\max_{s \in S} \mathbb{E}_{\theta \sim \rho} [v_{\text{SW}}(\theta; s(\theta))]$ . Therefore, we cannot directly apply the smoothness analysis of the strategic form or agent normal form. Existing studies analyze the POA of Bayes Nash equilibria by making an additional assumption. We will show that their results can be extended from Bayes Nash equilibria to communication equilibria with strategy representability.

In Section 5.1, we consider the case where the social welfare is the sum of all players’ payoffs, which is often assumed from the viewpoint of utilitarianism. We extend the proof given by

---

third blocks, but selecting  $\alpha_1$  all the time obtains  $T/3$  better reward.

Syrkkanis and Tardos [2013] from Bayes Nash equilibria to communication equilibria with strategy representability and obtain a lower bound of  $\lambda/(1 + \mu)$ . We make an assumption called  $(\lambda, \mu)$ -*conditional smoothness*, which considers deviations of each player depending on the current action. Since stronger deviations can be used, conditional smoothness is a weaker assumption than (unconditional) smoothness. This assumption seems to correspond to correlated equilibria, and Syrgkanis and Tardos [2013] actually used it to bound the POA of correlated equilibria for games with complete information. However, since our current proof uses the property of communication equilibria with strategy representability also for another purpose, even if we assume (unconditional) smoothness, it cannot extend to Bayes coarse correlated equilibria.

In Section 5.2, we consider Bayesian games derived from Bayesian mechanisms. Syrgkanis and Tardos [2013] defined  $(\lambda, \mu)$ -conditional smoothness<sup>8</sup> for mechanisms, which is satisfied by the simultaneous and sequential composition of first-price, second-price, all-pay, and many other auctions. We extend their lower bound of  $\lambda/\max\{1, \mu\}$  to communication equilibria with strategy representability.

In Section 5.3, we consider the case where the set of actions for each player changes depending on their type. For example, in a routing game, each player selects their route out of all paths from an origin to a destination, but if the origin-destination pair is determined by their type, then the set of possible actions changes depending on the type. A standard assumption for this setting is  $(\lambda, \mu)$ -*universal smoothness*,<sup>9</sup> which was satisfied by routing games and Bayesian effort games [Roughgarden, 2015b, Syrgkanis, 2012]. Under the conditional version of this assumption, we prove a lower bound of  $\lambda/(1 + \mu)$  on the POA of communication equilibria with strategy representability. We also briefly discuss how we can extend the definition of communication equilibria, the dynamics, and the untruthful swap regret minimization algorithm to this setting.

### 1.3 Further related work

**Computational studies on correlated equilibria** There exist many studies for computing a correlated equilibrium in games with *complete* information. A pioneering study on this topic was given by Gilboa and Zemel [1989], but the time complexity of the proposed algorithm is polynomial in the number of action profiles  $|A|$ , which is exponential in our setting. There are mainly two approaches to designing efficient algorithms. One approach is based on no-regret dynamics. No-internal-regret dynamics converging to correlated equilibria were developed by Foster and Vohra [1997] and Hart and Mas-Colell [2000] originally in the economics community. It was later applied to efficiently computing correlated equilibria [Nisan et al., 2007, Chapter 4] and extended to stronger regret notions [Blum and Mansour, 2007, Hazan and Kale, 2007, Rakhlin et al., 2011]. Another approach is *Ellipsoid Against Hope* proposed by Papadimitriou and Roughgarden [2008] (and modified by Jiang and Leyton-Brown [2015]), which computes an exact correlated equilibrium of succinctly represented games by the ellipsoid method. On the hardness side, Hart and Nisan [2018] proved that there is no deterministic or exact algorithm for computing a correlated equilibrium with polynomial query complexity.

---

<sup>8</sup>They call it smoothness, but we here explicitly call it conditional smoothness to distinguish unconditional and conditional versions.

<sup>9</sup>Roughgarden [2015b] simply call it smoothness, but we follow the terminology used by Syrgkanis [2012] for distinguishing it from (non-universal) smoothness.

**Computational studies on Bayes correlated equilibria** We can find definitions of Bayes correlated equilibria in several computer science papers, but their varieties have not been considered. In Appendix D.2.2 of their paper, Mansour et al. [2022] mentioned a class of correlated equilibria of two-player Bayesian games, in which *a single player* cannot benefit from any type swap *or* action swap but not their combination. Their definition is somewhat similar to communication equilibria, but the incentive constraints are weak. For Bayesian first-price auctions, Jin and Lu [2023] considered a class of Bayes correlated equilibria, which is equivalent to Bayesian solutions. Both of them did not provide any algorithm for computing an equilibrium in these classes.

Some studies mentioned Bayes coarse correlated equilibria. Hartline et al. [2015] considered no-external-regret dynamics in the agent normal form of Bayesian games and their convergence to an agent-normal-form coarse correlated equilibrium (ANFCCE). Since communication equilibria with strategy representability are included in ANFCCEs, our proposed dynamics converge to a narrower equilibrium concept. The definition of Bayes coarse correlated equilibria used by Caragiannis et al. [2015] and Jin and Lu [2023] is similar to ANFCCEs, but it is defined on  $\Delta(A)^\Theta$ , not  $\Delta(S)$ . We briefly discuss definitions of Bayes coarse correlated equilibria in Section D.

**Bayesian persuasion** Bayesian persuasion was proposed as a framework for information design by Kamenica and Gentzkow [2011]. As discussed by Bergemann and Morris [2016], Bayesian persuasion is optimization over *universal Bayesian solutions*, which is a variant of Bayes correlated equilibria for the setting where a mediator has more information than players. The set of universal Bayesian solutions is the largest among the variants described in this paper as shown by Forges [1993]. Computing a universal Bayesian solution is easier than the other classes considered in this paper, but computing an *optimal* equilibrium is rather difficult. Even for games with complete information, the problem of computing an optimal correlated equilibrium is in general harder than finding a correlated equilibrium [Papadimitriou and Roughgarden, 2008]. Most of the algorithmic studies of Bayesian persuasion make some assumption to alleviate complexity of equilibria: there is only one player [Dughmi and Xu, 2021] or players’ actions do not affect other players’ payoffs (no-externalities assumption) [Babichenko and Barman, 2017]. There are only a small number of exceptions, which deal with specific games such as atomic singleton congestion games [Zhou et al., 2022].

**Correlated equilibria in extensive-form games with imperfect information** Forges [1986] extended the notion of communication equilibria to extensive-form games, in which the mediator receives signals from players and then sends signals to players at the beginning of each node of the game tree. Bayesian games can be regarded as a special case of extensive-form games with imperfect information by regarding types as an action of a chance node. Note that the size of the game tree obtained by this reduction is exponential: the number of actions of this chance node is  $|\Theta|$ , and even in the case of complete information, the size of the game tree is at least the number of action profiles  $|A|$ . There is a large amount of literature on computing extensions of correlated equilibrium to extensive-form games with imperfect information [von Stengel and Forges, 2008, Celli and Gatti, 2018, Farina et al., 2022, Zhang and Sandholm, 2022], but they consider algorithms running in polynomial time in the size of the game tree, which is exponential in our formulation. Zhang and Sandholm [2022] mentioned how to design efficiently computable dynamics converging to a communication equilibrium as an open problem, which we partially resolve.

**Linear swap regret and polytope swap regret** Mansour et al. [2022] considered repeated play of Bayesian games with two players: an optimizer and a learner who follows an online learning algorithm. Their goal is to determine which concept of regret corresponds to the optimizer’s benefit compared with the Stackelberg value. Although their goal is different from ours, they proposed similar variants of regret: *linear swap regret* and *polytope swap regret*. For any polytope, linear swap regret is defined by the competitor that can use *any* valid linear transformation<sup>10</sup> for swapping the algorithm’s decision. Note that Gordon et al. [2008] already proposed the same concept for general convex compact sets as *linear regret*. Since untruthful swap regret is defined by a specific class of valid linear transformations represented by  $\mathcal{Q}$ , it is easier to minimize than linear swap regret. Mansour et al. [2022] proposed an algorithm for minimizing linear swap regret, but they did not prove any rigorous regret bound, and this algorithm was not proved to be efficient. In Section C, we prove that linear swap regret is equivalent to untruthful swap regret when the decision space is  $\mathcal{X}$  defined in Section 1.2, and then our algorithm can be used for minimizing linear swap regret.

Polytope swap regret is defined by the competitor that can use any swap of the vertices of the polytope, while the vertex decomposition of the algorithm’s decision in each round can be optimally decided by the learner in hindsight. Since our strategy swap regret requires the algorithm to submit vertex decomposition in each round irrevocably, it is more difficult to minimize than polytope swap regret.

**Price of anarchy of Bayesian games** The definition of the POA for games with complete information was provided by Koutsoupias and Papadimitriou [1999] and then extended from Nash equilibria to (coarse) correlated equilibria by Blum et al. [2008]. Most existing studies for Bayesian games focus on the POA of Bayes Nash equilibria [Christodoulou et al., 2016, Leme and Tardos, 2010, Roughgarden, 2015b, Syrgkanis, 2012, Syrgkanis and Tardos, 2013, Hartline et al., 2014, Feldman et al., 2020, Jin and Lu, 2022]. An exception is a study by Hartline et al. [2015], which extended the smoothness framework for mechanisms developed by Syrgkanis and Tardos [2013] to ANFCCEs. However, its applications are not clear due to the following reason. Hartline et al. [2015] showed that if a Bayesian mechanism satisfies (unconditional) smoothness, the POA of Bayes coarse correlated equilibria is bounded. Syrgkanis and Tardos [2013] showed that a wide class of mechanisms satisfies *conditional* smoothness, which assumes that the deviation can depend on the current action (see Section 5.2 for a formal definition). By carefully reading proofs of conditional smoothness of mechanisms presented by Syrgkanis and Tardos [2013], we can see that some mechanisms introduced in their paper satisfy (unconditional) smoothness, but not all. We prove that conditional smoothness is sufficient for extending the POA bounds to communication equilibria with strategy representability in Section 5.2.

As mentioned above, Caragiannis et al. [2015] and Jin and Lu [2023] considered a variant of ANFCCEs defined on  $\Delta(A)^\ominus$ . Caragiannis et al. [2015] proved POA bounds of this class for generalized second-price auctions. For first-price auctions, Jin and Lu [2023] considered the price of stability, which is defined to be the ratio of the *optimal* social welfare achieved by equilibria to the maximum social welfare.

In Section 5, we assume that the players’ types are independent, i.e., the prior distribution  $\rho$  is a product distribution. Roughgarden [2015b] called the POA for the correlated case *cPOA* (correlated POA) and showed a lower bound if the deviation of each player  $i \in N$  can be determined only by  $\theta_i$  in the definition of smoothness.

---

<sup>10</sup>We say a linear transformation is *valid* when it maps any vector in the polytope to a vector in the polytope.

## 1.4 Organization of this paper

The remaining part of this paper is organized as follows.

- Section 2 provides various definitions and characterizations of approximate Bayes correlated equilibria and then shows the convergence of variants of no-regret dynamics. Section 2.1 defines strategy representability, which is used for the characterizations. Section 2.2 specifies the online learning setting that each player is faced with in dynamics. Section 2.3 provides a definition of approximate SFCEs and shows that dynamics minimizing a variant of regret called strategy swap regret converge to them. Section 2.4 provides a definition of approximate ANFCEs and shows that dynamics minimizing a variant of regret called type-wise swap regret converge to them. Section 2.5, the main part of this section, provides a definition of approximate communication equilibria and shows that dynamics minimizing untruthful swap regret converge to communication equilibria with strategy representability. Section 2.6 shows several inclusion relations between the classes of Bayes correlated equilibria and Bayes Nash equilibria.
- Section 3 derives an efficient algorithm for minimizing untruthful swap regret in three steps, each of which is presented in Sections 3.1 to 3.3. Section 3.4 applies this algorithm to computing a communication equilibrium with strategy representability.
- Section 4 proves a lower bound on untruthful swap regret. The proof consists of four steps, each of which is presented in Sections 4.1 to 4.4.
- Section 5 shows lower bounds on the price of anarchy for each of three cases: Section 5.1 assumes that the social welfare is the sum of payoffs, Section 5.2 assumes that the game is derived from a Bayesian mechanism with a quasilinear utility function, and Section 5.3 assumes a stronger version of smoothness called universal smoothness.

## 2 Bayes correlated equilibria and no-regret dynamics

In this section, we review various classes of Bayes correlated equilibria surveyed by Forges [1993] and propose variants of no-regret dynamics converging to them. The highlight of this section is the definition of untruthful swap regret and the convergence of dynamics minimizing it to communication equilibria with strategy representability (Section 2.5). As a byproduct, we characterize each class of Bayes correlated equilibria, propose dynamics converging to SFCEs and ANFCEs, and show relations among these classes.

### 2.1 Strategy representability

To classify Bayes correlated equilibria in terms of distributions, we need to compare different kinds of distributions. As we will see, SFCEs and ANFCEs are defined as distributions over strategy profiles, that is,  $\sigma \in \Delta(S)$ . On the other hand, communication equilibria and Bayesian solutions are defined as distributions over action profiles for each type profile, which we call *type-wise distributions*. A type-wise distribution  $\pi \in \Delta(A)^\Theta$  determines a distribution  $\pi(\theta) \in \Delta(A)$  for each type  $\theta \in \Theta$ . We write  $\pi(\theta; a)$  instead of  $\pi(\theta)(a)$  for readability.

		$\theta_2$		$\theta'_2$	
		$a_2$	$a'_2$	$a_2$	$a'_2$
$\theta_1$	$a_1$	1/2	0	1/2	0
	$a'_1$	0	1/2	0	1/2
$\theta'_1$	$a_1$	1/2	0	0	1/2
	$a'_1$	0	1/2	1/2	0

Figure 3: An example of a type-wise distribution that is not strategy-representable.

To compare these distributions, we adopt the approach of transforming  $\sigma \in \Delta(S)$  into  $\pi \in \Delta(A)^\Theta$  by considering the marginal distribution. Suppose that the mediator samples  $s \sim \sigma$  independently of  $\theta \sim \rho$  and then recommends  $s_i(\theta_i)$  to each player  $i \in N$ . Then, for each  $\sigma \in \Delta(S)$ , we can define its corresponding distribution  $\pi \in \Delta(A)^\Theta$  by

$$\pi(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a) \quad (1)$$

for each  $\theta \in \Theta$  and  $a \in A$ . This operation defines a function  $\eta: \Delta(S) \rightarrow \Delta(A)^\Theta$ .

We define the strategy representability as the property that  $\pi \in \eta(\Delta(S))$ . This property implies that the mediator can realize  $\pi$  by sending a randomized recommendation of  $s \in S$  without observing the type profile  $\theta \in \Theta$ .

**Definition 2.1** (Strategy representability). We call a type-wise distribution  $\pi \in \Delta(A)^\Theta$  *strategy-representable* if there exists  $\sigma \in \Delta(S)$  such that  $\pi(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a)$  for each  $\theta \in \Theta$  and  $a \in A$ .

As we show below, not all type-wise distributions  $\Delta(A)^\Theta$  are strategy-representable. That is, if we generate  $s \sim \sigma$  independently of  $\theta \sim \rho$ , we cannot realize all distributions in  $\Delta(A)^\Theta$ .

**Example 2.2.** We give an example  $\pi \in \Delta(A)^\Theta$  that cannot be represented by any  $\sigma \in \Delta(S)$ . Suppose that  $A_1 = \{a_1, a'_1\}$ ,  $A_2 = \{a_2, a'_2\}$ ,  $\Theta_1 = \{\theta_1, \theta'_1\}$ , and  $\Theta_2 = \{\theta_2, \theta'_2\}$ . Let  $\rho$  be the uniform distribution on  $\Theta_1 \times \Theta_2$ . As illustrated in Figure 3, we define  $\pi \in \Delta(A)^\Theta$  by  $\pi(\theta_1, \theta_2; a_1, a_2) = \pi(\theta_1, \theta_2; a'_1, a'_2) = \pi(\theta_1, \theta'_2; a_1, a_2) = \pi(\theta_1, \theta'_2; a'_1, a'_2) = \pi(\theta'_1, \theta_2; a_1, a_2) = \pi(\theta'_1, \theta_2; a'_1, a'_2) = \pi(\theta'_1, \theta'_2; a_1, a_2) = \pi(\theta'_1, \theta'_2; a'_1, a'_2) = \frac{1}{2}$ . We show that  $\pi$  is not strategy-representable. If there exists  $\sigma \in \Delta(S)$  that satisfies  $\pi(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a)$  for any  $\theta \in \Theta$  and  $a \in A$ , then  $\Pr(s_1(\theta_1) = a_1, s_2(\theta_2) = a_2) = \Pr(s_1(\theta_1) = a'_1, s_2(\theta_2) = a'_2) = \frac{1}{2}$ . This implies that the event  $s_1(\theta_1) = a_1$  coincides with the event  $s_2(\theta_2) = a_2$ , and  $s_1(\theta_1) = a'_1$  coincides with  $s_2(\theta_2) = a'_2$ . Similarly, from  $\Pr(s_1(\theta_1) = a_1, s_2(\theta'_2) = a_2) = \Pr(s_1(\theta_1) = a'_1, s_2(\theta'_2) = a'_2) = \frac{1}{2}$ , we can see that  $s_1(\theta_1) = a_1$  coincides with  $s_2(\theta'_2) = a_2$ , and  $s_1(\theta_1) = a'_1$  coincides with  $s_2(\theta'_2) = a'_2$ . Therefore,  $s_2(\theta_2) = s_2(\theta'_2)$  holds with probability 1. By symmetry, we can see that  $s_1(\theta_1) = s_1(\theta'_1)$  holds with probability 1. Hence, the distributions of  $(s_1(\theta_1), s_2(\theta_2))$  and  $(s_1(\theta'_1), s_2(\theta'_2))$  must be identical, but they are not, which leads to a contradiction. This implies that  $\pi$  is not strategy-representable.

**Remark 2.3.** Forges [1993] did not formally define the strategy representability but instead claimed that a property called *conditional independence property* is important to classify Bayes correlated

equilibria. As Forges [1993] claimed, the conditional independence property holds for any strategy-representable distribution. Although it was not stated explicitly, the conditional independence property is not equivalent to the strategy representability (see Section A).<sup>11</sup> As discussed below, the strategy representability is the exact difference between several classes of Bayes correlated equilibria.

We say  $\pi \in \Delta(A)^\Theta$  is a *type-wise product distribution* if there exists some  $\pi_{i,\theta_i} \in \Delta(A_i)$  for each  $i \in N$  and  $\theta_i \in \Theta_i$  such that  $\pi(\theta; a) = \prod_{i \in N} \pi_{i,\theta_i}(a_i)$  for every  $\theta \in \Theta$  and  $a \in A$ . Let  $\Pi_{\text{prod}} \subseteq \Delta(A)^\Theta$  be the set of all type-wise product distributions. As we show below, a finite mixture of type-wise product distributions is always strategy-representable. We will use this fact to prove that the average of finite rounds of uncoupled dynamics in the agent normal form is strategy-representable since the distribution in each round is type-wise product. Furthermore, the inverse also holds.

**Proposition 2.4.** *A distribution  $\pi \in \Delta(A)^\Theta$  is strategy-representable if and only if  $\pi$  is a finite mixture of type-wise product distributions, that is, there exist a positive integer  $T$ ,  $\sigma' \in \Delta([T])$ , and  $\pi^t \in \Pi_{\text{prod}}$  for each  $t \in [T]$  such that  $\pi(\theta; a) = \sum_{t \in [T]} \sigma'(t) \pi^t(\theta; a)$  for each  $\theta \in \Theta$  and  $a \in A$ .*

*Proof.* Assume  $\pi$  is strategy-representable, that is, there exists  $\sigma \in \Delta(S)$  such that  $\pi(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a)$  for each  $\theta \in \Theta$  and  $a \in A$ . Let  $\sigma' = \sigma$  and define  $\pi^s \in \Delta(A)^\Theta$  by  $\pi^s(\theta; a) = \mathbf{1}_{\{s(\theta)=a\}}$ . We can check  $\pi^s \in \Pi_{\text{prod}}$  by setting  $\pi_{i,\theta_i}(a_i) = \mathbf{1}_{\{s_i(\theta_i)=a_i\}}$ . Then

$$\pi(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a) = \sum_{s \in S} \sigma(s) \mathbf{1}_{\{s(\theta)=a\}} = \sum_{s \in S} \sigma'(s) \pi^s(\theta; a)$$

for each  $\theta \in \Theta$  and  $a \in A$ .

Assume there exists a positive integer  $T$ ,  $\sigma' \in \Delta(T)$ , and  $\pi^t \in \Pi_{\text{prod}}$  for each  $t \in [T]$  such that  $\pi(\theta; a) = \sum_{t \in [T]} \sigma'(t) \pi^t(\theta; a)$ . Since  $\pi^t \in \Pi_{\text{prod}}$  for each  $t \in [T]$ , there exists  $\pi_{i,\theta_i}^t \in \Delta(A_i)$  for each  $i \in N$  and  $\theta_i \in \Theta_i$  such that  $\pi^t(\theta; a) = \prod_{i \in N} \pi_{i,\theta_i}^t(a_i)$  for each  $\theta \in \Theta$  and  $a \in A$ . We define  $\sigma^t \in \Delta(S)$  by  $\sigma^t(s) = \prod_{i \in N} \prod_{\theta_i \in \Theta_i} \pi_{i,\theta_i}^t(s_i(\theta_i))$  for each  $s \in S$ . Then for each  $\theta \in \Theta$  and  $a \in A$ , it holds that

$$\Pr_{s \sim \sigma^t}(s(\theta) = a) = \sum_{\substack{s \in S: \\ s(\theta)=a}} \prod_{i \in N} \prod_{\theta_i \in \Theta_i} \pi_{i,\theta_i}^t(s_i(\theta_i)) = \prod_{i \in N} \pi_{i,\theta_i}^t(a_i) = \pi^t(\theta; a).$$

Define  $\sigma \in \Delta(S)$  by  $\sigma(s) = \sum_{t \in [T]} \sigma'(t) \sigma^t(s)$  for every  $s \in S$ . Then

$$\Pr_{s \sim \sigma}(s(\theta) = a) = \sum_{t \in [T]} \sigma'(t) \Pr_{s \sim \sigma^t}(s(\theta) = a) = \sum_{t \in [T]} \sigma'(t) \pi^t(\theta; a) = \pi(\theta; a)$$

for each  $\theta \in \Theta$  and  $a \in A$ , which implies that  $\pi$  is strategy-representable.  $\square$

## 2.2 Online learning with stochastic types

To discuss dynamics converging to Bayes correlated equilibria, we define *online learning with stochastic types*, which is an online decision-making problem that each player is faced with in repeated play of Bayesian games.

<sup>11</sup>A related discussion about *belief invariant Bayesian solutions* can be found in existing studies [Lehrer et al., 2010, Forges, 2006].

This problem is defined as an extension of an ordinary online learning setting. Recall that in each round of online learning, a learner with a set of finite actions  $A_i$  decides a randomized action  $a_i \in A_i$  and then obtains payoff  $u_i^t(a_i)$  according to a reward vector  $u_i^t \in [0, 1]^{A_i}$ . In online learning with stochastic types, in addition to actions  $A_i$ , the learner is associated with a set of finite types  $\Theta_i$  and a prior distribution  $\rho_i \in \Delta(\Theta_i)$ . In each round  $t \in [T]$ , the learner decides a strategy distribution  $\sigma_i^t \in \Delta(S_i)$ . This distribution randomly determines a strategy  $s_i \in A_i^{\Theta_i}$  that assigns an action  $a_i \in A_i$  to each type  $\theta_i \in \Theta_i$ . Then the learner observes a reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  that assigns a reward value for each type  $\theta_i \in \Theta_i$  and action  $a_i \in A_i$ . The learner's action  $s_i(\theta_i) \in A_i$  is randomly determined by strategy  $s_i \sim \sigma_i^t$  and type  $\theta_i \sim \rho_i$ . The expected reward for randomized strategy  $\sigma_i^t \in \Delta(A_i^{\Theta_i})$  is therefore  $\mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, s_i(\theta_i))] \right]$ .

Suppose that players are repeatedly involved in the same Bayesian game and decide their own randomized strategies in each round. Then the online decision-making problem that each player is faced with can be formulated as online learning with stochastic types. In each round  $t \in [T]$ , each player  $i \in N$  decides a randomized strategy  $\sigma_i^t \in \Delta(S_i)$  and then obtains the expected payoff  $\mathbb{E}_{\theta \sim \rho} [\mathbb{E}_{s \sim \sigma^t} [v_i(\theta; s(\theta))]]$ , where  $\sigma^t$  is the product distribution that independently generates  $s_j \sim \sigma_j^t$  for each  $j \in N$ . We define the reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  by

$$u_i^t(\theta_i, a_i) = \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}^t} [v_i(\theta; a_i, s_{-i}(\theta_{-i}))] \right] \quad (2)$$

for each type  $\theta_i \in \Theta_i$  and action  $a_i \in A_i$ , where  $\sigma_{-i}^t$  is the product distribution that independently generates  $s_j \sim \sigma_j^t$  for each  $j \in N \setminus \{i\}$ . By using this reward vector, the expected payoff for player  $i \in N$  can be expressed as

$$\begin{aligned} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma^t} [v_i(\theta; s(\theta))] \right] &= \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} \left[ \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}^t} [v_i(\theta; s_i(\theta_i), s_{-i}(\theta_{-i}))] \right] \right] \right] \\ &= \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, s_i(\theta_i))] \right]. \end{aligned}$$

Thus, the problem for player  $i \in N$  in repeated play of Bayesian games is online learning with stochastic types specified by actions  $A_i$ , types  $\Theta_i$ , prior distribution  $\rho_i \in \Delta(\Theta_i)$ , and reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  for each  $t \in [T]$  as defined above. Note that the decision space is  $S_i$ , but the rewards for them can be succinctly represented by  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$ . If the goal is to minimize external regret, this problem can be solved separately for each type. By considering various extensions of swap regret to this setting, we will define dynamics converging to each class of Bayes correlated equilibria.

For ANFCEs (Section 2.4) and communication equilibria (Section 2.5), we restrict the decision  $\sigma_i^t \in \Delta(S_i)$  of each player  $i \in N$  to a product distribution, i.e., there exists some  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$  such that  $\sigma_i^t(s_i) = \prod_{\theta_i \in \Theta_i} \pi_i^t(\theta_i; s_i(\theta_i))$  for each  $s_i \in S_i$ . Then the definition of reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  is

$$u_i^t(\theta_i, a_i) = \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{a_{-i} \sim \pi_{-i}^t(\theta_{-i})} [v_i(\theta; a)] \right] \quad (3)$$

for each type  $\theta_i \in \Theta_i$  and action  $a_i \in A_i$ , where  $\pi_{-i}^t(\theta_{-i})$  is the product distribution that independently generates  $a_j \sim \pi_j^t(\theta_j)$  for each  $j \in N \setminus \{i\}$ . While a strategy distribution  $\sigma_i^t \in \Delta(S_i)$  requires space exponential in  $|\Theta_i|$ , a product distribution can be succinctly represented by  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$

in polynomial space. For ANFCEs and communication equilibria, we will define type-wise swap regret and untruthful swap regret, for which a product distribution is sufficient. On the other hand, for SFCEs, we define strategy swap regret, for which we need to consider a strategy distribution (Section 2.3).

### 2.3 Strategic-form correlated equilibria

*Strategic-form correlated equilibria* (SFCEs; also known as normal-form correlated equilibria) are defined as correlated equilibria of the strategic form of a Bayesian game. Recall that in the strategic form, the set of strategies  $S_i = A_i^{\Theta_i}$  is considered as a decision space for each player  $i \in N$ . This class is defined as a set of joint distributions  $\sigma \in \Delta(S)$  on strategy profiles in which each player  $i \in N$  cannot gain by deviating from recommended strategy  $s_i \in S_i$  in expectation. This deviation is represented by a swap  $\phi_{\text{SF}}: S_i \rightarrow S_i$  that maps recommended strategy  $s_i \in S_i$  to another strategy  $\phi_{\text{SF}}(s_i)$ . Here we define an approximate version in which each player can gain at most  $\epsilon$  in expectation by deviating from the recommendation.

**Definition 2.5** ( $\epsilon$ -Approximate strategic-form correlated equilibria ( $\epsilon$ -SFCEs)). For any  $\epsilon \geq 0$ , a distribution  $\sigma \in \Delta(S)$  is an  $\epsilon$ -approximate strategic-form correlated equilibrium if for any  $i \in N$  and any  $\phi_{\text{SF}}: S_i \rightarrow S_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i}))] \right] - \epsilon. \quad (\text{ICSF})$$

Let  $\Sigma_{\text{SF}}^\epsilon \subseteq \Delta(S)$  be the set of all  $\epsilon$ -SFCEs.

The notion of SFCE is a natural equilibrium concept for the following scenario. A mediator samples  $s \in S$  according to a distribution  $\sigma \in \Delta(S)$  and then informs each player  $i \in N$  of  $s_i$ . Independently, a type profile  $\theta$  is sampled from  $\rho$  and each player  $i \in N$  observes their own type  $\theta_i$ . Then, each player  $i \in N$  decides their action  $s_i(\theta_i) \in A_i$ . If every player cannot gain by deviating from action  $s_i(\theta_i)$  in expectation, the distribution  $\sigma$  is an SFCE.

The role of this mediator can be interpreted as a correlation device. A correlation device privately sends (possibly correlated) signals to each player according to some distribution. An SFCE can be defined as a Nash equilibrium of an extended game equipped with a correlation device. A correlation device can send any signal in general, but from the revelation principle, we can assume that each signal corresponds to a recommendation of each strategy without loss of generality.

Since SFCEs are correlated equilibria of the strategic form of Bayesian games, no-swap-regret dynamics converging to a correlated equilibrium can be directly extended to dynamics converging to an SFCE. Recall that the decision space for each player  $i \in N$  in the strategic form is  $S_i$ . Then each player  $i \in N$  decides a strategy distribution  $\sigma_i^t \in \Delta(S_i)$  in each round  $t \in [T]$  and obtains the expected payoff  $\mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, s_i(\theta_i))] \right]$  according to reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  defined as (2), which is the setting of online learning with stochastic types. We consider an extension of swap regret to this online learning problem with decision space  $S_i$ , which we call *strategy swap regret*.

**Definition 2.6** (Strategy swap regret). For online learning with stochastic types specified by actions  $A_i$ , types  $\Theta_i$ , prior distribution  $\rho_i$ , and reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  for every time round  $t \in [T]$ , strategy swap regret is defined as

$$R_{\text{SS},i}^T = \max_{\phi_{\text{SF}}: S_i \rightarrow S_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, (\phi_{\text{SF}}(s_i))(\theta_i)) - u_i^t(\theta_i, s_i(\theta_i))] \right].$$

---

**Algorithm 1:** Dynamics for strategic-form correlated equilibria

---

For each player  $i \in N$ , let  $\mathcal{A}_i$  be a subroutine that minimizes strategy swap regret for online learning with stochastic types;

**for** each round  $t = 1, \dots, T$  **do**

Each player  $i \in N$  decides their randomized strategy  $\sigma_i^t \in \Delta(S_i)$  according to  $\mathcal{A}_i$  and shares it with the other players;

Each player  $i \in N$  computes reward  $u_i^t(\theta_i, a_i) = \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}^t} [v_i(\theta; a_i, s_{-i}(\theta_{-i}))] \right]$  for every  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ , where  $\sigma_{-i}^t \in \Delta(S_{-i})$  is the product distribution that independently generates  $s_j \sim \sigma_j^t$  for each  $j \in N \setminus \{i\}$ ;

Feed  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  to  $\mathcal{A}_i$  as a reward vector for round  $t$ ;

**end**

---

Theorem 3 of Blum and Mansour [2007] relates each player's swap regret with the convergence to a correlated equilibrium. By directly applying it to the strategic form of Bayesian games, we can show that if the strategy swap regret of each player  $i \in N$  grows sublinear in  $T$ , we can guarantee the convergence of the dynamics to an SFCE. This is claimed in the following proposition formally.

**Proposition 2.7.** *Let  $\sigma_i^t \in \Delta(S_i)$  be the randomized strategy of each player  $i \in N$  for each round  $t \in [T]$  in Algorithm 1. Then, the empirical distribution  $\sigma \in \Delta(S)$  defined by  $\sigma(s) = \frac{1}{T} \sum_{t=1}^T \prod_{i \in N} \sigma_i^t(s_i)$  for each  $s \in S$  is a  $\frac{\max_{i \in N} R_{SS,i}^T}{T}$ -SFCE, where  $R_{SS,i}^T$  is the strategy swap regret for each subroutine  $\mathcal{A}_i$ .*

*Proof.* From the definition of  $\epsilon$ -SFCEs, it is sufficient to prove

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i}))] \right] - \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \leq \frac{\max_{j \in N} R_{SS,j}^T}{T}$$

for each  $i \in N$  and  $\phi_{\text{SF}}: S_i \rightarrow S_i$ . Let  $\sigma^t \in \Delta(S_{-i})$  be the product distribution that independently generates  $s_j \sim \sigma_j^t$  for each  $j \in N$  and  $\sigma_{-i}^t \in \Delta(S)$  the product distribution that independently generates  $s_j \sim \sigma_j^t$  for each  $j \in N \setminus \{i\}$ . Then the left-hand side can be bounded as

$$\begin{aligned} & \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i}))] \right] - \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \\ &= \mathbb{E}_{t \sim [T]} \left[ \mathbb{E}_{s \sim \sigma^t} \left[ \mathbb{E}_{\theta \sim \rho} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i})) - v_i(\theta; s(\theta))] \right] \right] \quad (\text{from the definition of } \sigma) \\ &= \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{s_i \sim \sigma_i^t} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}^t} \left[ \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i})) - v_i(\theta; s(\theta))] \right] \right] \right] \\ & \hspace{15em} (\text{since } \sigma^t \text{ is the product distribution}) \\ &= \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{s_i \sim \sigma_i^t} \left[ \mathbb{E}_{\theta_i \sim \rho_i} [u_i^t(\theta_i, (\phi_{\text{SF}}(s_i))(\theta_i)) - u_i^t(\theta_i, s_i(\theta_i))] \right] \quad (\text{from the definition of } u_i^t) \\ &\leq \frac{R_{SS,i}^T}{T} \end{aligned}$$

$$\leq \frac{\max_{j \in N} R_{SS,j}^T}{T},$$

which completes the proof.  $\square$

We provide a detailed description of the dynamics in Algorithm 1. We should note that this proposition does not imply the existence of a polynomial-time algorithm for computing an approximate SFCE. This is because the size of the decision space  $S_i$  for each player  $i \in N$  is exponential. In Section B, we provide an algorithm that has an  $O\left(\sqrt{T|A_i|^{|\Theta_i|} \log |A_i|} + |A_i|^{|\Theta_i|} \log |A_i|\right)$  upper bound on strategy swap regret, which is slightly better than a direct application of the swap regret minimization algorithm proposed by Blum and Mansour [2007] but still exponential in  $|\Theta_i|$ . Computing an  $\epsilon$ -SFCE by simulating the dynamics with this algorithm requires time exponential in  $|\Theta_i|$ . Whether there exists an algorithm that computes an  $\epsilon$ -SFCE in time polynomial in  $n$ ,  $|A_i|$ ,  $|\Theta_i|$ , and  $1/\epsilon$  with an oracle for utility functions  $(v_i)_{i=1}^n$  is an open problem left for future work.

**Remark 2.8.** We define approximate SFCEs based on swap regret, but it is also possible to define them based on internal regret. Note that we have two different definitions of approximate correlated equilibria that are based on internal regret or swap regret.<sup>12</sup> These definitions coincide when the error is 0, but different when the error is non-zero. In this paper, we adopt the extension based on swap regret, but it is also possible to use the extension based on internal regret, in which  $\phi_{\text{SF}}: S_i \rightarrow S_i$  in the definition is restricted to a swap for a single strategy:

$$\phi_{\text{SF}}(s_i) = \begin{cases} s_i'' & \text{if } s_i = s_i' \\ s_i & \text{otherwise} \end{cases}$$

for some  $s_i', s_i'' \in S_i$ . However, the problem of computing this version of approximate SFCEs is rather easy. It is because if we set  $\sigma \in \Delta(S)$  to the uniform distribution over  $S$ , then the error is always at most  $1/(\min_{i \in N} |S_i|)$ . Hence, for  $\epsilon \geq 1/(\min_{i \in N} |S_i|)$ , computing it can be done easily. On the other hand, for  $\epsilon \leq 1/(\min_{i \in N} |S_i|)$ , if we assume  $|S_i|$  has the same order for every  $i \in N$ , then we can straightforwardly run an internal regret minimization algorithm with decision space  $S_i$  in time polynomial in  $1/\epsilon = \min_{i \in N} |S_i|$ .

To compare SFCEs with other classes defined on  $\Delta(A)^\Theta$ , we define the set of distributions obtained by mapping SFCEs to  $\Delta(A)^\Theta$  by  $\eta: \Delta(S) \rightarrow \Delta(A)^\Theta$  (see Section 2.1 for the definition of  $\eta$ ). Formally, it is defined as  $\Pi_{\text{SF}}^\epsilon = \{\eta(\sigma) \in \Delta(A)^\Theta \mid \sigma \in \Sigma_{\text{SF}}^\epsilon\}$ . From the definition, every  $\pi \in \Pi_{\text{SF}}^\epsilon$  is strategy-representable.

As we will see,  $\Pi_{\text{SF}}^\epsilon$  is a subset of approximate communication equilibria. Informally, computing an  $\epsilon$ -SFCE is harder than computing an  $\epsilon$ -approximate communication equilibrium. If an  $\epsilon$ -SFCE  $\sigma \in \Delta(S)$  can be obtained, then we can compute its corresponding distribution  $\eta(\sigma) \in \Delta(A)^\Theta$ , which is also an  $\epsilon$ -approximate communication equilibrium. Note that this reduction is informal because distributions in  $\Delta(S)$  and  $\Delta(A)^\Theta$  have an exponentially large support in general, and the computational complexity of computing  $\eta(\sigma)$  depends on the succinct representation of  $\sigma$ .

Whether  $\sigma \in \Delta(S)$  is an SFCE cannot be determined by its corresponding  $\eta(\sigma) \in \Delta(A)^\Theta$ . As in the following example, even if  $\sigma, \sigma' \in \Delta(S)$  correspond to the same  $\pi \in \Delta(A)^\Theta$ , it is possible that  $\sigma$  is an SFCE, but  $\sigma'$  is not.

<sup>12</sup>Ganor and Karthik C. S. [2018] called them *approximate correlated equilibria* and *approximate rule correlated equilibria*, respectively.

**Example 2.9.** Suppose  $\Theta_1 = \{\theta_1, \theta'_1\}$  and  $\Theta_2 = \{\theta_2\}$  with  $A_1 = \{a_1, a'_1\}$  and  $A_2 = \{a_2, a'_2\}$ . The prior distribution  $\rho \in \Delta(\Theta)$  is the uniform distribution over  $\Theta$ . Let  $v_1(\theta; a_1, a_2) = v_1(\tilde{\theta}; a'_1, a'_2) = 1$  and  $v_1(\tilde{\theta}; a'_1, a_2) = v_1(\tilde{\theta}; a_1, a'_2) = 0$  for any  $\tilde{\theta} \in \Theta$ , while the second player's payoff is always 0, i.e.,  $v_2 \equiv 0$ . We consider two distinct distributions  $\sigma, \sigma' \in \Delta(S)$  with the same corresponding  $\pi \in \Delta(A)^\Theta$ . The first one  $\sigma$  is the uniform distribution over  $S$ . The second one  $\sigma'$  is the uniform distribution over  $s^1, s^2, s^3, s^4$  defined as

$$\begin{aligned} s_1^1(\theta_1) &= a_1, & s_1^1(\theta'_1) &= a_1, & s_2^1(\theta_2) &= a_2, \\ s_1^2(\theta_1) &= a'_1, & s_1^2(\theta'_1) &= a'_1, & s_2^2(\theta_2) &= a_2, \\ s_1^3(\theta_1) &= a_1, & s_1^3(\theta'_1) &= a'_1, & s_2^3(\theta_2) &= a'_2, \\ s_1^4(\theta_1) &= a'_1, & s_1^4(\theta'_1) &= a_1, & s_2^4(\theta_2) &= a'_2. \end{aligned}$$

Both  $\pi(\theta_1, \theta_2)$  and  $\pi(\theta'_1, \theta_2)$  are the uniform distribution over  $A_1 \times A_2$ . The expected payoff for each type of  $\theta_1$  and  $\theta'_1$  is  $\frac{1}{2}$ .

First, we show that  $\sigma$  is an SFCE. Since  $v_2 \equiv 0$ , the second player does not have any incentive to deviate. Since  $\sigma$  is the uniform distribution over  $S$ , when  $s_1 \in S_1$  is recommended, the posterior distribution of  $s_2$  is the uniform distribution over  $S_2$ . Hence, the expected payoff of each action of  $A_1$  is  $\frac{1}{2}$ , and the first player does not have any incentive to deviate.

Next, we show that  $\sigma'$  is not an SFCE. Since  $s_1^1, s_1^2, s_1^3, s_1^4$  are all distinct, the first player can infer the recommendation  $s_2(\theta_2)$  for the second player. The first player can increase the payoff to 1 by taking the same action as the second player. Therefore,  $\sigma'$  is not an SFCE.

## 2.4 Agent-normal-form correlated equilibria

Agent-normal-form correlated equilibria (ANFCEs) are defined as correlated equilibria of the agent normal form of a Bayesian game. Recall that in the agent normal form, the same player with different types is hypothetically considered as distinct players. In an ANFCE, the mediator recommends possibly correlated actions to all hypothetical players  $(i, \theta_i) \in N'$ . Since an action profile for the agent normal form assigns some  $a_i \in A_i$  to every  $i \in N$  and  $\theta_i \in \Theta_i$ , it can be identified with a strategy profile  $s \in S$ . Hence, we can define ANFCEs as the set of distributions over  $S$ . The only difference from SFCEs is the information that each player can use for deviations. While each player  $i \in N$  can decide their action according to the recommended actions for all types  $\Theta_i$  in SFCEs, each player  $i \in N$  can use only the recommended action  $s_i(\theta_i)$  for their realized type  $\theta_i \in \Theta_i$  in ANFCEs. Here we define an approximate version of ANFCEs with additive error  $\epsilon$ .

**Definition 2.10** ( $\epsilon$ -Approximate agent-normal-form correlated equilibria ( $\epsilon$ -ANFCEs)). For any  $\epsilon \geq 0$ , a distribution  $\sigma \in \Delta(S)$  is an  $\epsilon$ -approximate agent-normal-form correlated equilibrium if for any  $i \in N$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))] \right] - \epsilon. \quad (\text{IC}_{\text{ANF}})$$

Let  $\Sigma_{\text{ANF}}^\epsilon \subseteq \Delta(S)$  be the set of all  $\epsilon$ -ANFCEs.

In contrast to SFCEs, each player  $i \in N$  in ANFCEs is informed of only the action for the realized type  $s_i(\theta_i)$ , not the full strategy  $s_i$ . In an SFCE, each player can use recommendations for unrealized types, which differentiates SFCEs from ANFCEs. In an ANFCE, each player  $i \in N$

with type  $\theta_i \in \Theta_i$  can use only  $s_i(\theta_i)$  to decide deviation  $\phi(\theta_i, s_i(\theta_i))$ . This difference makes  $\Sigma_{\text{ANF}}^\epsilon$  broader than  $\Sigma_{\text{SF}}^\epsilon$ .

**Remark 2.11.** The incentive constraint in our definition is slightly stronger than the incentive constraint for each hypothetical player  $(i, \theta'_i) \in N'$ , which can be written as

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbf{1}_{\{\theta_i = \theta'_i\}} \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbf{1}_{\{\theta_i = \theta'_i\}} \mathbb{E}_{s \sim \sigma} \left[ v_i(\theta; \phi_{\theta'_i}(s_i(\theta_i)), s_{-i}(\theta_{-i})) \right] \right] - \epsilon$$

for any  $\phi_{\theta'_i}: A_i \rightarrow A_i$ . This can be checked by letting  $\phi(\theta'_i, \cdot) = \phi_{\theta'_i}$  and  $\phi(\theta''_i, \cdot)$  be the identity map for each  $\theta''_i \in \Theta_i \setminus \{\theta'_i\}$  in  $(\text{IC}_{\text{ANF}})$ . Note that this difference is only for the error parameter, and these two incentive constraints are equivalent for  $\epsilon = 0$ .

Since the setting of the agent normal form is hypothetical, it is difficult to consider a natural scenario for ANFCEs. In an ANFCE, the mediator determines a strategy profile  $s \in S$  according to  $\sigma \in \Delta(S)$  but recommends to each player only  $s_i(\theta_i)$ . To realize it, the mediator must know the type profile  $\theta \in \Theta$  in advance. However, if it is the case, then it is possible to realize any distribution  $\pi(\theta) \in \Delta(A)$  for each  $\theta \in \Theta$ , for which the notion of Bayesian solutions introduced later is more appropriate.

By extending no-swap-regret-dynamics to the agent normal form, we can obtain uncoupled dynamics converging to an ANFCE. In these dynamics, for every time round  $t \in [T]$ , each hypothetical player  $(i, \theta_i) \in N'$  decides a distribution  $\pi_i^t(\theta_i) \in \Delta(A_i)$  over their actions. Hence, swap regret for each hypothetical player  $(i, \theta'_i) \in N'$  considers an action swap  $\phi_{\theta'_i}: A_i \rightarrow A_i$ . As with the incentive constraints discussed above, we consider the sum of swap regrets for all hypothetical players corresponding to  $i \in N$ , which we call *type-wise swap regret*.

**Definition 2.12** (Type-wise swap regret). For online learning with stochastic types specified by actions  $A_i$ , types  $\Theta_i$ , prior distribution  $\rho_i$ , and reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  for each round  $t \in [T]$ , type-wise swap regret is defined as

$$R_{\text{TS},i}^T = \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, \phi(\theta_i, a_i)) - u_i^t(\theta_i, a_i)] \right].$$

As with SFCEs, by directly applying Theorem 3 of Blum and Mansour [2007] to the agent normal form of Bayesian games, we obtain the following claim.

**Proposition 2.13.** *Let  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$  be the type-wise distribution of each player  $i \in N$  for each round  $t \in [T]$  in Algorithm 2. Then, the empirical distribution  $\sigma \in \Delta(S)$ , which is defined by  $\sigma(s) = \frac{1}{T} \sum_{t=1}^T \prod_{i \in N} \prod_{\theta_i \in \Theta_i} \pi_i^t(\theta_i; s_i(\theta_i))$  for each  $s \in S$ , is a  $\frac{\max_{i \in N} R_{\text{TS},i}^T}{T}$ -ANFCE, where  $R_{\text{TS},i}^T$  is the type-wise swap regret for each subroutine  $A_i$ .*

*Proof.* From the definition of  $\epsilon$ -ANFCEs, it is sufficient to prove

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))] \right] - \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \leq \frac{\max_{j \in N} R_{\text{TS},j}^T}{T}$$

for each  $i \in N$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Let  $\sigma^t \in \Delta(S)$  be the product distribution that independently generates  $s_j(\theta_j) \sim \pi_j^t(\theta_j)$  for each  $j \in N$  and  $\theta_j \in \Theta_j$ . For each  $\theta_{-i} \in \Theta_{-i}$ , let  $\pi_{-i}^t(\theta_{-i}) \in \Delta(A_{-i})$

---

**Algorithm 2:** Dynamics for ANFCEs or communication equilibria

---

For each  $i \in N$ , let  $\mathcal{A}_i$  be a subroutine that minimizes type-wise swap regret (ANFCEs) or untruthful swap regret (communication equilibria with strategy representability) for online learning with stochastic types;

**for** each round  $t = 1, \dots, T$  **do**

Each player  $i \in N$  decides their randomized strategy  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$  according to  $\mathcal{A}_i$  and shares it with the other players;

Each player  $i \in N$  computes reward  $u_i^t(\theta_i, a_i) = \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{a_{-i} \sim \pi_{-i}^t(\theta_{-i})} [v_i(\theta; a)] \right]$  for every  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ , where  $\pi_{-i}^t(\theta_{-i}) \in \Delta(A_{-i})$  is the product distribution that independently generates  $a_j \sim \pi_j^t(\theta_j)$  for each  $j \in N \setminus \{i\}$ ;

Feed  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  to each  $\mathcal{A}_i$  as a reward vector for round  $t$ ;

**end**

---

be the product distribution that independently generates  $a_j \sim \pi_j^t(\theta_j)$  for each  $j \in N \setminus \{i\}$ . Then the left-hand side can be bounded as

$$\begin{aligned}
& \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))] \right] - \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \\
&= \mathbb{E}_{t \sim [T]} \left[ \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma^t} [v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i})) - v_i(\theta; s(\theta))] \right] \right] \quad (\text{from the definition of } \sigma) \\
&= \mathbb{E}_{t \sim [T]} \left[ \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} \left[ \mathbb{E}_{a_{-i} \sim \pi_{-i}^t(\theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i}) - v_i(\theta; a)] \right] \right] \right] \quad (\text{from the definition of } \sigma^t) \\
&= \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} \left[ \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{a_{-i} \sim \pi_{-i}^t(\theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i}) - v_i(\theta; a)] \right] \right] \right] \\
&= \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, \phi(\theta_i, a_i)) - u_i^t(\theta_i, a_i)] \right] \quad (\text{from the definition of } u_i^t) \\
&\leq \frac{R_{\text{TS},i}^T}{T} \\
&\leq \frac{\max_{j \in N} R_{\text{TS},j}^T}{T},
\end{aligned}$$

which completes the proof.  $\square$

A detailed description of the dynamics can be found in Algorithm 2. We can obtain an algorithm for minimizing type-wise swap regret by running the swap regret minimization algorithm proposed by Blum and Mansour [2007] separately for each  $(i, \theta_i) \in N'$ . Since the rewards for each  $(i, \theta_i) \in N'$  are always bounded above by  $\rho_i(\theta_i)$ , an upper bound of  $O(\sqrt{T|A_i| \log |A_i|})$  on swap regret directly applies to type-wise swap regret.

As with  $\epsilon$ -SFCEs, we can define its counterpart  $\Pi_{\text{ANF}}^\epsilon \subseteq \Delta(A)^\Theta$  by  $\Pi_{\text{ANF}}^\epsilon = \eta(\Sigma_{\text{ANF}}^\epsilon)$  (see Section 2.1 for the definition of  $\eta$ ). Since the incentive constraints for ANFCEs only rely on the action profile  $s(\theta)$  instead of the strategy profile  $s$ , we can characterize  $\Pi_{\text{ANF}}^\epsilon$  as follows.

**Proposition 2.14.** For any type-wise distribution  $\pi \in \Delta(A)^\Theta$ , the following are equivalent:

(i)  $\pi \in \Pi_{\text{ANF}}^\epsilon$ .

(ii)  $\pi$  is strategy-representable, and for any  $i \in N$  and any  $\phi: \Theta_i \times A_i \rightarrow A_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon. \quad (\text{IC}'_{\text{ANF}})$$

*Proof.* First, we prove (ii) assuming (i). Suppose  $\pi \in \Pi_{\text{ANF}}^\epsilon$ . Since  $\Pi_{\text{ANF}}^\epsilon = \eta(\Sigma_{\text{ANF}}^\epsilon)$ , there exists  $\sigma \in \Sigma_{\text{ANF}}^\epsilon$  such that  $\eta(\sigma) = \pi$ . From the definition of  $\Sigma_{\text{ANF}}^\epsilon$ , the incentive constraint  $(\text{IC}_{\text{ANF}})$  holds for every  $i \in N$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Since  $\pi = \eta(\sigma)$ , when  $s$  is generated from  $\sigma$ , the distribution of  $s(\theta)$  equals  $\pi(\theta)$ . Then  $(\text{IC}_{\text{ANF}})$  implies  $(\text{IC}'_{\text{ANF}})$ .

Next, we prove (i) assuming (ii). Suppose  $\pi$  is strategy-representable and satisfies the incentive constraint  $(\text{IC}'_{\text{ANF}})$  for every  $i \in N$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . From the strategy representability, there exists  $\sigma \in \Delta(S)$  such that  $\pi = \eta(\sigma)$ . As discussed above, when  $s$  is generated from  $\sigma$ , the distribution of  $s(\theta)$  equals  $\pi(\theta)$ . Then  $(\text{IC}'_{\text{ANF}})$  implies  $(\text{IC}_{\text{ANF}})$ .  $\square$

## 2.5 Communication equilibria

A communication equilibria is defined as a type-wise distribution  $\pi \in \Delta(A)^\Theta$  that is incentive-compatible in the following scenario. First, type profile  $\theta \in \Theta$  is jointly generated from prior distribution  $\rho$ , and then each player observes their own type  $\theta_i \in \Theta_i$ . Next, the players privately report their own types to the mediator. Then the mediator generates an action profile  $a \in A$  from  $\pi(\theta)$  and privately recommends  $a_i$  to each player  $i \in N$ . If each player  $i \in N$  has no incentive to deviate from reporting the true type  $\theta_i$  and following the recommendation  $a_i$ , then this type-wise distribution  $\pi \in \Delta(A)^\Theta$  is called a communication equilibrium.

Formally, the incentive constraints for each player  $i \in N$  are defined for the following two mappings. The first one  $\psi: \Theta_i \rightarrow \Theta_i$  specifies a type  $\psi(\theta_i)$  to be reported when the true type is  $\theta_i$ . As a result, the distribution  $\pi(\psi(\theta_i), \theta_{-i})$  that generates the recommended action profile  $a \in A$  may be different from the original one  $\pi(\theta)$ . The second one  $\phi: \Theta_i \times A_i \rightarrow A_i$  specifies an action  $\phi(\theta_i, a_i)$  when  $a_i$  is recommended by the mediator and the true type is  $\theta_i$ . If each player cannot gain by using any combination of  $\psi$  and  $\phi$ , a distribution  $\pi$  is a communication equilibrium. Here we define its approximate version.

**Definition 2.15** ( $\epsilon$ -Approximate communication equilibria). For any  $\epsilon \geq 0$ , a type-wise distribution  $\pi \in \Delta(A)^\Theta$  is an  $\epsilon$ -approximate communication equilibrium if for any  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\psi(\theta_i), \theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon. \quad (\text{IC}_{\text{Com}})$$

Let  $\Pi_{\text{Com}}^\epsilon \subseteq \Delta(A)^\Theta$  be the set of all  $\epsilon$ -approximate communication equilibria.

A communication equilibrium can be defined as a Nash equilibrium for the extended game equipped with a *communication device*, which is a device that privately receives a signal from each player and then privately sends a stochastic (possibly correlated) signal to each player based on

the received signals. In general, the inputs and outputs for the communication device may be any signals different from type profiles  $\Theta$  and action profiles  $A$ . By the revelation principle, it is sufficient to consider *direct* (or *canonical*) ones, in which the sets of inputs and outputs coincide with  $\Theta$  and  $A$  [Myerson, 1982]. In this paper, we consider direct communication equilibria only.

There are two differences between communication equilibria and ANFCEs. First, in an ANFCE, the mediator must generate a strategy profile  $s \in S$  before observing the type profile  $\theta \in \Theta$ , while in a communication equilibrium, the mediator can generate an action profile  $a \in A$  after observing the type profile. Since there exists a type-wise distribution  $\pi$  that cannot be expressed by any strategy distribution (Section 2.1), ANFCEs are more restrictive than communication equilibria in this aspect. Second, in an ANFCE, the players must inform the mediator of their true types, while in a communication equilibrium, the players can untruthfully report their types. If the mediator can validate the reported types in some way, ANFCEs are more natural, but communication equilibria are more natural otherwise.

To introduce dynamics converging to a communication equilibrium, we present the definition of untruthful swap regret. In addition to  $\phi: \Theta_i \times A_i \rightarrow A_i$  that represents the deviation from the recommended action, the competitor can use  $\psi: \Theta_i \rightarrow \Theta_i$  that represents untruthful type reporting.

**Definition 2.16** (Untruthful swap regret). For online learning with stochastic types specified by actions  $A_i$ , types  $\Theta_i$ , prior distribution  $\rho_i$ , and reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  for every round  $t \in [T]$ , untruthful swap regret is defined as

$$R_{\text{US},i}^T = \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta_i, \phi(\theta_i, a_i))] - \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, a_i)] \right].$$

If we restrict  $\psi$  to the identity map, untruthful swap regret is reduced to type-wise swap regret. Then untruthful swap regret is no less than type-wise swap regret.

Now we consider dynamics in which each player  $i \in N$  minimizes untruthful swap regret  $R_{\text{US},i}$ . In the following theorem, we show that these dynamics converge to a communication equilibrium. Moreover, the empirical distribution is always strategy-representable from Theorem 2.4.

**Theorem 2.17.** Let  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$  be the type-wise distribution of each player  $i \in N$  for each round  $t \in [T]$  in Algorithm 2. Then, the empirical distribution  $\pi \in \Delta(A)^\Theta$ , which is defined by  $\pi(\theta; a) = \frac{1}{T} \sum_{t=1}^T \prod_{i \in N} \pi_i^t(\theta_i; a_i)$  for each  $\theta \in \Theta$  and  $a \in A$ , is a  $\frac{\max_{i \in N} R_{\text{US},i}^T}{T}$ -approximate communication equilibrium that satisfies strategy representability, where  $R_{\text{US},i}^T$  is the untruthful swap regret for each subroutine  $\mathcal{A}_i$ .

*Proof.* From Theorem 2.4,  $\pi$  is strategy-representable. From the definition of  $\epsilon$ -approximate communication equilibria, it is sufficient to prove

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\psi(\theta_i), \theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \leq \frac{\max_{j \in N} R_{\text{US},j}^T}{T}$$

for each  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . For each  $t \in [T]$  and  $\theta \in \Theta$ , let  $\pi^t(\theta) \in \Delta(A)$  be the product distribution that independently generates  $a_j \sim \pi_j^t(\theta_j)$  for each  $j \in N$ . Similarly, let  $\pi_{-i}^t(\theta) \in \Delta(A_{-i})$  be the product distribution that independently generates  $a_j \sim \pi_j^t(\theta_j)$  for each  $j \in N \setminus \{i\}$ . The left-hand side can be bounded as

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\psi(\theta_i), \theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right]$$

$$\begin{aligned}
&= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{t \sim [T]} \left[ \mathbb{E}_{a \sim \pi^t(\psi(\theta_i), \theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] - \mathbb{E}_{a \sim \pi^t(\theta)} [v_i(\theta; a)] \right] \right] \\
&= \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{\theta_{-i} \sim \rho_{-i} | \theta_i} \left[ \mathbb{E}_{a_{-i} \sim \pi_{-i}^t(\theta_{-i})} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] - \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [v_i(\theta; a)] \right] \right] \right] \\
&\hspace{15em} \text{(since } \pi^t(\theta) \text{ is a product distribution)} \\
&= \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta_i, \phi(\theta_i, a_i))] - \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, a_i)] \right] \quad \text{(from the definition of } u_i^t) \\
&\leq \frac{R_{\text{US},i}^T}{T} \\
&\leq \frac{\max_{j \in N} R_{\text{US},j}^T}{T},
\end{aligned}$$

which completes the proof.  $\square$

We present characterizations of communication equilibria with strategy representability, to which the proposed dynamics converge.

**Proposition 2.18.** *For any type-wise distribution  $\pi \in \Delta(A)^\Theta$ , the following are equivalent:*

- (i)  $\pi \in \Pi_{\text{Com}}^\epsilon$  and  $\pi$  is strategy-representable.
- (ii)  $\pi \in \Pi_{\text{Com}}^\epsilon \cap \Pi_{\text{ANF}}^\epsilon$ .
- (iii) There exists some  $\sigma \in \Delta(S)$  such that  $\eta(\sigma) = \pi$  and

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; \phi(\theta_i, s_i(\psi(\theta_i))), s_{-i}(\theta_{-i}))] \right] - \epsilon. \quad (\text{IC}_{\text{ComSR}})$$

holds for any  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ .

*Proof.* First, we prove (ii) assuming (i). Suppose  $\pi \in \Pi_{\text{Com}}^\epsilon$  and  $\pi$  is strategy-representable. Since  $(\text{IC}_{\text{ANF}}^\epsilon)$  is a weaker condition than  $(\text{IC}_{\text{Com}}^\epsilon)$ ,  $\pi$  satisfies  $(\text{IC}_{\text{ANF}}^\epsilon)$  for any  $\phi: \Theta_i \times A_i \rightarrow A_i$ . We thus obtain  $\pi \in \Pi_{\text{ANF}}^\epsilon$  from Theorem 2.14.

Next, we prove (iii) assuming (ii). Suppose  $\pi \in \Pi_{\text{Com}}^\epsilon \cap \Pi_{\text{ANF}}^\epsilon$ . Since  $\Pi_{\text{ANF}}^\epsilon = \eta(\Sigma_{\text{ANF}}^\epsilon)$  from the definition,  $\pi \in \Pi_{\text{ANF}}^\epsilon$  implies that there exists  $\sigma \in \Sigma_{\text{ANF}}^\epsilon$  such that  $\eta(\sigma) = \pi$ . Moreover, since  $\pi \in \Pi_{\text{Com}}^\epsilon$ ,  $(\text{IC}_{\text{Com}}^\epsilon)$  holds for any  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Since  $\eta$  is defined by  $\pi(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a)$  for each  $\theta \in \Theta$  and  $a \in A$ , the left-hand side of  $(\text{IC}_{\text{Com}}^\epsilon)$  is equal to the left-hand side of  $(\text{IC}_{\text{ComSR}}^\epsilon)$ . Then it is sufficient to show the equality of their right-hand sides. Again from the definition of  $\eta$ , for any  $\theta \in \Theta$ , we have

$$\pi(\psi(\theta_i), \theta_{-i}; a) = \Pr_{s \sim \sigma}(s_i(\psi(\theta_i)) = a_i, s_{-i}(\theta_{-i}) = a_{-i}).$$

Hence, for each fixed  $\theta \in \Theta$ , the distribution of  $a$  in the right-hand side of  $(\text{IC}_{\text{Com}}^\epsilon)$  equals the distribution of  $(s_i(\psi(\theta_i)), s_{-i}(\theta_{-i}))$  in the right-hand side of  $(\text{IC}_{\text{ComSR}}^\epsilon)$ , and then the right-hand sides are equal.

Finally, we prove (i) assuming (iii). The existence of  $\sigma \in \Delta(S)$  such that  $\eta(\sigma) = \pi$  implies the strategy representability of  $\pi$ . It is sufficient to prove that  $(\text{IC}_{\text{Com}}^\epsilon)$  holds for any  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Since  $(\text{IC}_{\text{ComSR}}^\epsilon)$  and  $(\text{IC}_{\text{Com}}^\epsilon)$  are equivalent as proved above, we obtain  $\pi \in \Pi_{\text{Com}}^\epsilon$ .  $\square$

## 2.6 Relations among classes of Bayes correlated equilibria

Here we compare the classes of Bayes correlated equilibria and show their relations. First, we define an approximate version of *Bayesian solutions* (also called *partial Bayesian approach*) [Forges, 1993].

**Definition 2.19** ( $\epsilon$ -approximate Bayesian solutions). For any  $\epsilon \geq 0$ , a type-wise distribution  $\pi \in \Delta(A)^\Theta$  is an  $\epsilon$ -approximate Bayesian solution if for any  $i \in N$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon. \quad (\text{IC}_{\text{BS}})$$

Let  $\Pi_{\text{BS}}^\epsilon \subseteq \Delta(A)^\Theta$  be the set of all  $\epsilon$ -approximate Bayesian solutions.

Forges [1993] showed the relations among the classes in terms of the set of players' payoff values. Formally, if we denote the expected payoff vector achieved by  $\pi \in \Delta(A)^\Theta$  by

$$v_{\text{EP}}(\pi) = \left( \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \right)_{i \in N} \in [0, 1]^N,$$

then Forges [1993] showed  $v_{\text{EP}}(\Pi_{\text{SF}}^0) \subseteq v_{\text{EP}}(\Pi_{\text{ANF}}^0) \cap v_{\text{EP}}(\Pi_{\text{Com}}^0)$  and  $v_{\text{EP}}(\Pi_{\text{ANF}}^0) \cup v_{\text{EP}}(\Pi_{\text{Com}}^0) \subseteq v_{\text{EP}}(\Pi_{\text{BS}}^0)$ . Moreover, Forges [1993] showed that these inclusions are strict, and  $v_{\text{EP}}(\Pi_{\text{ANF}}^0) \not\subseteq v_{\text{EP}}(\Pi_{\text{Com}}^0)$  and  $v_{\text{EP}}(\Pi_{\text{Com}}^0) \not\subseteq v_{\text{EP}}(\Pi_{\text{ANF}}^0)$  for some Bayesian games.

Here we show the same inclusion relations among the approximate versions in terms of type-wise distributions, that is,

$$\Pi_{\text{SF}}^\epsilon \subseteq \Pi_{\text{ANF}}^\epsilon \cap \Pi_{\text{Com}}^\epsilon \quad \text{and} \quad \Pi_{\text{ANF}}^\epsilon \cup \Pi_{\text{Com}}^\epsilon \subseteq \Pi_{\text{BS}}^\epsilon$$

for each  $\epsilon \geq 0$ . We also show that these inclusion relations are strict. First, we show the following relation.

**Proposition 2.20.** *For any Bayesian game, it holds that  $\Pi_{\text{SF}}^\epsilon \subseteq \Pi_{\text{ANF}}^\epsilon \cap \Pi_{\text{Com}}^\epsilon$  for any  $\epsilon \geq 0$ .*

*Proof.* Let  $\sigma \in \Sigma_{\text{SF}}^\epsilon$  be an arbitrary  $\epsilon$ -SFCE and  $\pi = \eta(\sigma) \in \Pi_{\text{SF}}^\epsilon$  its corresponding type-wise distribution. From Theorem 2.18, it is sufficient to prove the incentive constraint ( $\text{IC}_{\text{ComSR}}$ ) for each  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Since  $\sigma$  is an  $\epsilon$ -SFCE, the incentive constraint ( $\text{IC}_{\text{SF}}$ ) holds for any  $\phi_{\text{SF}}: S_i \rightarrow S_i$ . If we set  $(\phi_{\text{SF}}(s_i))(\theta_i) = \phi(\theta_i, s_i(\psi(\theta_i)))$  for each  $\theta_i \in \Theta_i$ , then it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; \phi(\theta_i, s_i(\psi(\theta_i))), s_{-i}(\theta_{-i}))] \right] - \epsilon,$$

which completes the proof.  $\square$

Next, we provide an example of a Bayesian game in which  $\pi \in (\Pi_{\text{ANF}}^0 \cap \Pi_{\text{Com}}^0) \setminus \Pi_{\text{SF}}^0$  exists. While deviations for SFCEs allow players to use recommendations for all other types, the deviations for communication equilibria allow players to use recommendations for at most one other type and ANFCE no other type. We design a distribution such that a player can gain by using recommendations for two other types. Then this distribution is a communication equilibrium and ANFCE but not SFCE.

**Proposition 2.21.** *For some Bayesian game,  $(\Pi_{\text{ANF}}^0 \cap \Pi_{\text{Com}}^0) \setminus \Pi_{\text{SF}}^0 \neq \emptyset$ .*

*Proof.* The distribution is defined as follows. Let  $\Theta_1 = \{\theta_1, \theta'_1, \theta''_1\}$  and  $\Theta_2 = \{\theta_2\}$  with the uniform prior distribution  $\rho$  over  $\Theta_1 \times \Theta_2$ . Let  $A_1 = \{0, 1, 2, 3, 4\}$  and  $A_2 = \{1, 2, 3, 4\}$ . We define the first player's payoffs for  $\theta_1, \theta'_1, \theta''_1$  by

$$\begin{aligned} v_1(\theta_1, \theta_2; a_1, a_2) &= \begin{cases} \frac{1}{2} & \text{if } a_1 = 0 \\ 1 & \text{if } a_1 \neq 0 \text{ and } a_1 = a_2 \\ 0 & \text{if } a_1 \neq 0 \text{ and } a_1 \neq a_2. \end{cases} \\ v_1(\theta'_1, \theta_2; a_1, a_2) &= \begin{cases} \frac{1}{2} & \text{if } (a_1, a_2) \in \{(4, 1), (4, 2), (1, 3), (1, 4)\} \\ 0 & \text{otherwise} \end{cases} \\ v_1(\theta''_1, \theta_2; a_1, a_2) &= \begin{cases} \frac{1}{2} & \text{if } (a_1, a_2) \in \{(2, 1), (2, 3), (3, 2), (3, 4)\} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

The second player's payoff  $v_2$  is defined as  $v_2(\tilde{\theta}_1, \theta_2; a_1, a_2) = \mathbf{1}_{\{a_1 \neq a_2\}}$  for any  $\tilde{\theta}_1 \in \Theta_1$ . We consider  $\pi \in \Delta(A)^\Theta$  such that  $\pi(\theta_1, \theta_2)$ ,  $\pi(\theta'_1, \theta_2)$ , and  $\pi(\theta''_1, \theta_2)$  are the uniform distributions over  $\{0\} \times \{1, 2, 3, 4\}$ ,  $\{(4, 1), (4, 2), (1, 3), (1, 4)\}$ , and  $\{(2, 1), (2, 3), (3, 2), (3, 4)\}$ , respectively.

- First, we show that  $\pi$  is strategy-representable by providing  $\sigma \in \Delta(S)$  such that  $\pi = \eta(\sigma)$ . From the definition of  $\pi(\theta_1, \theta_2)$ , it must hold that  $s_1(\theta_1) = 0$  with probability 1. From the definition of  $\pi(\theta'_1, \theta_2)$ , it holds that

$$\begin{aligned} \Pr_{s \sim \sigma}(s_1(\theta'_1) = 4 \mid s_2(\theta_2) = 1) &= 1 & \Pr_{s \sim \sigma}(s_1(\theta'_1) = 4 \mid s_2(\theta_2) = 2) &= 1 \\ \Pr_{s \sim \sigma}(s_1(\theta'_1) = 1 \mid s_2(\theta_2) = 3) &= 1 & \Pr_{s \sim \sigma}(s_1(\theta'_1) = 1 \mid s_2(\theta_2) = 4) &= 1. \end{aligned}$$

Similarly, from the definition of  $\pi(\theta''_1, \theta_2)$ , it holds that

$$\begin{aligned} \Pr_{s \sim \sigma}(s_1(\theta''_1) = 2 \mid s_2(\theta_2) = 1) &= 1 & \Pr_{s \sim \sigma}(s_1(\theta''_1) = 3 \mid s_2(\theta_2) = 2) &= 1 \\ \Pr_{s \sim \sigma}(s_1(\theta''_1) = 2 \mid s_2(\theta_2) = 3) &= 1 & \Pr_{s \sim \sigma}(s_1(\theta''_1) = 3 \mid s_2(\theta_2) = 4) &= 1. \end{aligned}$$

Hence,  $\sigma$  can be defined as

$$\begin{aligned} \Pr_{s \sim \sigma}(s_1(\theta_1) = 0, s_1(\theta'_1) = 4, s_1(\theta''_1) = 2, s_2(\theta_2) = 1) &= \frac{1}{4} \\ \Pr_{s \sim \sigma}(s_1(\theta_1) = 0, s_1(\theta'_1) = 4, s_1(\theta''_1) = 3, s_2(\theta_2) = 2) &= \frac{1}{4} \\ \Pr_{s \sim \sigma}(s_1(\theta_1) = 0, s_1(\theta'_1) = 1, s_1(\theta''_1) = 2, s_2(\theta_2) = 3) &= \frac{1}{4} \\ \Pr_{s \sim \sigma}(s_1(\theta_1) = 0, s_1(\theta'_1) = 1, s_1(\theta''_1) = 3, s_2(\theta_2) = 4) &= \frac{1}{4}. \end{aligned} \tag{4}$$

Therefore,  $\pi$  is strategy-representable.

- Next, we show that for any  $\psi: \Theta_i \rightarrow \Theta_i$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ , distribution  $\pi$  satisfies the incentive constraint for communication equilibria ( $\text{IC}_{\text{Com}}$ ). Since the second player's payoff is at most 1 and the expected payoff is 1 under  $\pi$ , this player does not have any incentive

to deviate from the recommendation. Under distribution  $\pi$ , the first player always obtains payoff  $\frac{1}{2}$ , and then the expected payoff is  $\frac{1}{2}$ . Then we show that the first player cannot gain by using any combination of  $\psi$  and  $\phi$ . When the type is  $\theta'_1$  or  $\theta''_1$ , the payoff is at most  $\frac{1}{2}$  and the expected payoff is  $\frac{1}{2}$  under  $\pi$ , the first player does not have any incentive to deviate from the recommendation. When the type is  $\theta_1$ , we consider the case of  $\psi(\theta_1) = \theta_1$  and the others separately. If  $\psi(\theta_1) = \theta_1$ , then the first player is always recommended taking the action 0, and the posterior distribution on the second player's action is still the uniform distribution. If  $\psi(\theta_1) = \theta'_1$  (or  $\psi(\theta_1) = \theta''_1$ ), then the first player can learn whether  $s_2(\theta_2) \in \{1, 2\}$  or  $s_2(\theta_2) \in \{3, 4\}$  ( $s_2(\theta_2) \in \{1, 3\}$  or  $s_2(\theta_2) \in \{2, 4\}$ , respectively). By using this information, the first player can take the same action  $s_2(\theta_2)$  with probability 1/2, but the expected payoff is  $\frac{1}{2} \cdot 1 = \frac{1}{2}$ , which is not better than the recommendation. Therefore, the right-hand side is at most  $\frac{1}{2}$  for any  $\psi$  and  $\phi$ .

These two facts imply  $\pi \in \Pi_{\text{ANF}}^0 \cap \Pi_{\text{Com}}^0$  from Theorem 2.18. Finally, we show  $\pi \notin \Pi_{\text{SF}}^0$ . As discussed above, it is sufficient to consider  $\sigma$  that satisfies the conditions (4). When the first player's type is  $\theta_1$ , the first player can learn  $s_2(\theta_2)$  from  $s_1(\theta'_1)$  and  $s_1(\theta''_1)$ . Then  $\phi_{\text{SF}}$  can be defined such that  $(\phi_{\text{SF}}(s_1))(\theta_1) = s_2(\theta_2)$  always holds. This deviation provides the expected payoff  $\frac{1}{3} \cdot 1 + \frac{1}{3} \cdot \frac{1}{2} + \frac{1}{3} \cdot \frac{1}{2} = \frac{2}{3}$  to the first player, which is better than  $\frac{1}{2}$  obtained by the recommendation  $\sigma$ . Therefore,  $\sigma$  is not an SFCE and  $\pi \notin \Pi_{\text{SF}}^0$ .  $\square$

**Remark 2.22.** The example used in the proof also proves that  $v_{\text{EP}}(\Pi_{\text{ANF}}^0 \cap \Pi_{\text{Com}}^0) \setminus v_{\text{EP}}(\Pi_{\text{SF}}^0) \neq \emptyset$  for some Bayesian game. This is stronger than the claim that  $(v_{\text{EP}}(\Pi_{\text{ANF}}^0) \cap v_{\text{EP}}(\Pi_{\text{Com}}^0)) \setminus v_{\text{EP}}(\Pi_{\text{SF}}^0) \neq \emptyset$  for some Bayesian game, which was proved by Forges [1993] with distinct  $\pi \in \Pi_{\text{ANF}}^0$  and  $\pi' \in \Pi_{\text{Com}}^0$  such that  $v_{\text{EP}}(\pi) = v_{\text{EP}}(\pi')$ .

We can check  $v_{\text{EP}}(\pi) = (\frac{1}{2}, 1) \notin v_{\text{EP}}(\Pi_{\text{SF}}^0)$  as follows. If the second player's expected payoff is 1, then  $a_1 \neq a_2$  holds with probability 1. Hence, the first player with type  $\theta_1$  never obtains the payoff 1. To achieve the expected payoff  $\frac{1}{2}$ , the first player must achieve the payoff  $\frac{1}{2}$  with probability 1, but then  $\sigma$  must satisfy the conditions (4). As proved above,  $\sigma$  is not an SFCE.

**Bayesian solutions** Next, we show that  $\epsilon$ -approximate Bayesian solutions contain the union of  $\epsilon$ -ANFCEs and  $\epsilon$ -approximate communication equilibria.

**Proposition 2.23.** *For any Bayesian game, it holds that  $\Pi_{\text{ANF}}^\epsilon \cup \Pi_{\text{Com}}^\epsilon \subseteq \Pi_{\text{BS}}^\epsilon$  for any  $\epsilon \geq 0$ .*

*Proof.* Assume  $\pi \in \Pi_{\text{ANF}}^\epsilon$  or  $\pi \in \Pi_{\text{Com}}^\epsilon$ . If  $\pi \in \Pi_{\text{ANF}}^\epsilon$ , then  $\pi$  satisfies the incentive constraint  $(\text{IC}'_{\text{ANF}})$  for any  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Since  $(\text{IC}_{\text{BS}})$  is identical to  $(\text{IC}'_{\text{ANF}})$ , this implies  $\pi \in \Pi_{\text{BS}}^\epsilon$ . If  $\pi \in \Pi_{\text{Com}}^\epsilon$ , then  $\pi$  satisfies the incentive constraint  $(\text{IC}_{\text{Com}})$  for any  $\psi: \Theta_i \rightarrow \Theta_i$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . By considering  $\psi$  as the identity map, we can obtain  $(\text{IC}_{\text{BS}})$ , which implies  $\pi \in \Pi_{\text{BS}}^\epsilon$ .  $\square$

This inclusion for  $\epsilon = 0$  is strict for some Bayesian game. Forges [1993] provided an example of a Bayesian game for which  $v_{\text{EP}}(\Pi_{\text{BS}}^0) \setminus (v_{\text{EP}}(\Pi_{\text{ANF}}^0) \cup v_{\text{EP}}(\Pi_{\text{Com}}^0)) \neq \emptyset$  holds. Since  $v_{\text{EP}}(\Pi_{\text{ANF}}^0) \cup v_{\text{EP}}(\Pi_{\text{Com}}^0) = v_{\text{EP}}(\Pi_{\text{ANF}}^0 \cup \Pi_{\text{Com}}^0)$  holds in general,  $v_{\text{EP}}(\Pi_{\text{BS}}^0) \setminus v_{\text{EP}}(\Pi_{\text{ANF}}^0 \cup \Pi_{\text{Com}}^0) \neq \emptyset$  in this game, which implies  $\Pi_{\text{BS}}^0 \setminus (\Pi_{\text{ANF}}^0 \cup \Pi_{\text{Com}}^0) \neq \emptyset$ .

**Bayes Nash equilibria** Finally, we consider the relations between Bayes Nash equilibria and the classes of Bayes correlated equilibria. Recall that there are two interpretations of a Bayesian game: the strategic form and the agent normal form. For each of these two different normal-form

games corresponding to the same Bayesian game, we can consider the concept of Nash equilibria. A well-known interesting fact is that these two concepts are equivalent. This equilibrium concept is called *Bayes Nash equilibrium*. Here we define approximate versions of Bayes Nash equilibria in the strategic form and the agent normal form, respectively.

First, we define Bayes Nash equilibria as Nash equilibria of the strategic form. In the strategic form, each player  $i \in N$  independently decides a strategy  $s_i \in S_i$  according to some distribution  $\sigma_i \in \Delta(S_i)$ . Here we denote this Nash equilibrium by the product distribution of  $\sigma_1, \dots, \sigma_n$ . Let  $\Sigma_{\text{Prod}} \subseteq \Delta(S)$  be the set of all product distributions, that is, for each  $\sigma \in \Sigma_{\text{Prod}}$ , there exists some  $\sigma_i \in \Delta(S_i)$  for each  $i \in N$  such that  $\sigma(s) = \prod_{i=1}^n \sigma_i(s_i)$  holds for every  $s \in S$ . We then define an approximate version of Bayes Nash equilibria in the strategic form as follows.

**Definition 2.24** ( $\epsilon$ -Approximate Bayes Nash equilibria ( $\epsilon$ -BNEs) in the strategic form). For any  $\epsilon \geq 0$ , a distribution  $\sigma \in \Sigma_{\text{Prod}} \subseteq \Delta(S)$  is an  $\epsilon$ -approximate Bayes Nash equilibrium in the strategic form if for any  $i \in N$  and any  $\phi_{\text{SF}}: S_i \rightarrow S_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i}))] \right] - \epsilon. \quad (\text{IC}_{\text{BNE}})$$

Let  $\Sigma_{\text{BNE}}^\epsilon \subseteq \Delta(S)$  be the set of all  $\epsilon$ -approximate Bayes Nash equilibria in the strategic form.

Next, we define Bayes Nash equilibria as Nash equilibria in the agent normal form. In the agent normal form, each hypothetical player  $(i, \theta_i) \in N'$  decides a distribution  $\pi_{i, \theta_i} \in \Delta(A_i)$ . A type profile  $\theta \in \Theta$  sampled from  $\rho \in \Delta(\Theta)$  determines the active players, whose actions  $a_i \in A_i$  are independently sampled from  $\pi_{i, \theta_i}$  for each  $i \in N$ . Therefore, the action profile of the active players follows a type-wise product distribution  $\pi \in \Pi_{\text{Prod}}$ . Recall that we call  $\pi \in \Delta(A)^\Theta$  a type-wise product distribution if there exists some  $\pi_{i, \theta_i} \in \Delta(A_i)$  for each  $i \in N$  and  $\theta_i \in \Theta_i$  such that  $\pi(\theta; a) = \prod_{i \in N} \pi_{i, \theta_i}(a_i)$  for every  $\theta \in \Theta$  and  $a \in A$ . We define an approximate version of Bayes Nash equilibria in the agent normal form as follows.

**Definition 2.25** ( $\epsilon$ -Approximate Bayes Nash equilibria ( $\epsilon$ -BNEs) in the agent normal form). For any  $\epsilon \geq 0$ , a type-wise product distribution  $\pi \in \Pi_{\text{Prod}} \subseteq \Delta(A)^\Theta$  is an  $\epsilon$ -approximate Bayes Nash equilibria in the agent normal form if for any  $i \in N$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon. \quad (\text{IC}'_{\text{BNE}})$$

Let  $\Pi_{\text{BNE}}^\epsilon \subseteq \Delta(A)^\Theta$  be the set of all  $\epsilon$ -approximate Bayes Nash equilibria in the agent normal form.

Here we show the equivalence of  $\Sigma_{\text{BNE}}^\epsilon$  and  $\Pi_{\text{BNE}}^\epsilon$ . This is an approximate version of the well-known fact that Nash equilibria in the strategic form and the agent normal form are equivalent [Harsanyi, 1967], but we provide the proof for completeness.

**Proposition 2.26.** *For any distribution  $\sigma \in \Sigma_{\text{BNE}}^\epsilon$ , its corresponding type-wise distribution satisfies  $\eta(\sigma) \in \Pi_{\text{BNE}}^\epsilon$ . Conversely, for any type-wise distribution  $\pi \in \Pi_{\text{BNE}}^\epsilon$ , there exists a distribution  $\sigma \in \Sigma_{\text{BNE}}^\epsilon$  such that  $\eta(\sigma) = \pi$ .*

*Proof.* Let  $\sigma \in \Sigma_{\text{BNE}}^\epsilon$  be any  $\epsilon$ -BNE in the strategic form and  $\pi = \eta(\sigma)$ . First, we show that  $\pi$  is a type-wise product distribution. Since  $\sigma \in \Sigma_{\text{Prod}}$ , there exists  $\sigma_i \in \Delta(S_i)$  for each  $i \in N$  such that

$\sigma(s) = \prod_{i=1}^n \sigma_i(s_i)$  holds for every  $s \in S$ . For each  $i \in N$  and  $\theta_i \in \Theta_i$ , define  $\pi_{i,\theta_i} \in \Delta(A_i)$  by  $\pi_{i,\theta_i}(a_i) = \Pr_{s_i \sim \sigma_i}(s_i(\theta_i) = a_i)$  for each  $a_i \in A_i$ . Then

$$\pi(\theta; a) = \eta(\sigma)(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a) = \prod_{i \in N} \Pr_{s_i \sim \sigma_i}(s_i(\theta_i) = a_i) = \prod_{i \in N} \pi_{i,\theta_i}(a_i)$$

for every  $\theta \in \Theta$  and  $a \in A$ , which implies  $\pi \in \Pi_{\text{Prod}}$ .

Next, we show the incentive constraint ( $\text{IC}'_{\text{BNE}}$ ) for each  $i \in N$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . We use the incentive constraint ( $\text{IC}_{\text{BNE}}$ ) for  $\phi_{\text{SF}}: S_i \rightarrow S_i$  defined by  $(\phi_{\text{SF}}(s_i))(\theta_i) = \phi(\theta_i, s_i(\theta_i))$  for each  $\theta_i \in \Theta_i$ . Since  $\eta(\sigma) = \pi$  implies that  $s(\theta)$  conforms to  $\pi(\theta) \in \Delta(A)$ , the left-hand side of ( $\text{IC}_{\text{BNE}}$ ) equals the left-hand side of ( $\text{IC}'_{\text{BNE}}$ ). The right-hand side of ( $\text{IC}_{\text{BNE}}$ ) equals

$$\begin{aligned} & \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i}))] \right] - \epsilon \\ &= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))] \right] - \epsilon \quad (\text{since } (\phi_{\text{SF}}(s_i))(\theta_i) = \phi(\theta_i, s_i(\theta_i))) \\ &= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon, \quad (\text{since } s(\theta) \text{ conforms to } \pi(\theta)) \end{aligned}$$

which coincides with the right-hand side of ( $\text{IC}'_{\text{BNE}}$ ). Therefore,  $\pi$  is an  $\epsilon$ -BNE in the agent normal form.

Conversely, let  $\pi \in \Pi_{\text{BNE}}^\epsilon$  be any  $\epsilon$ -BNE in the agent normal form. Since  $\pi$  is a type-wise product distribution, there exists  $\pi_{i,\theta_i} \in \Delta(A_i)$  for each  $i \in N$  and  $\theta_i \in \Theta_i$  such that  $\pi(\theta; a) = \prod_{i \in N} \pi_{i,\theta_i}(a_i)$  for every  $\theta \in \Theta$  and  $a \in A$ . We can define the corresponding  $\sigma \in \Delta(S)$  by  $\sigma(s) = \prod_{i \in N} \prod_{\theta_i \in \Theta_i} \pi_{i,\theta_i}(s_i(\theta_i))$  for every  $s \in S$ . Then

$$\eta(\sigma)(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a) = \prod_{i \in N} \pi_{i,\theta_i}(a_i) = \pi(\theta; a)$$

for every  $\theta \in \Theta$  and  $a \in A$ , which implies  $\eta(\sigma) = \pi$ .

Next, we show  $\sigma \in \Sigma_{\text{BNE}}^\epsilon$  by proving  $\sigma \in \Sigma_{\text{Prod}}$  and the incentive constraint ( $\text{IC}_{\text{BNE}}$ ) holds for any  $i \in N$  and  $\phi_{\text{SF}}: S_i \rightarrow S_i$ . We can see  $\sigma \in \Sigma_{\text{Prod}}$  by expressing  $\sigma$  as  $\sigma(s) = \prod_{i \in N} \sigma_i(s_i)$  for every  $s \in S$  where  $\sigma_i \in \Delta(S_i)$  is defined by  $\sigma_i(s_i) = \prod_{\theta_i \in \Theta_i} \pi_{i,\theta_i}(s_i(\theta_i))$  for every  $s_i \in S_i$ . We then prove that ( $\text{IC}_{\text{BNE}}$ ) holds for any  $i \in N$  and  $\phi_{\text{SF}}: S_i \rightarrow S_i$  by using ( $\text{IC}'_{\text{BNE}}$ ) for  $\phi: \Theta_i \times A_i \rightarrow A_i$  defined by

$$\phi(\theta_i, a_i) \in \operatorname{argmax}_{a'_i \in A_i} \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}} [v_i(\theta; a'_i, s_{-i}(\theta_{-i}))] \right]$$

for each  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ , where  $\sigma_{-i} \in \Delta(S_{-i})$  is the product distribution that independently generates  $s_j \sim \sigma_j$  for each  $j \in N \setminus \{i\}$ . Note that  $\phi(\theta_i, \cdot)$  can be a constant function. Since  $\eta(\sigma) = \pi$ , the left-hand side of ( $\text{IC}'_{\text{BNE}}$ ) equals the left-hand side of ( $\text{IC}_{\text{BNE}}$ ). The right-hand side of ( $\text{IC}'_{\text{BNE}}$ ) is bounded as

$$\begin{aligned} & \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon \\ &= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}} [v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))] \right] - \epsilon \quad (\text{since } s(\theta) \text{ conforms to } \pi(\theta)) \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}} [v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))] \right] \right] - \epsilon \\
&= \mathbb{E}_{\theta_i \sim \rho_i} \left[ \max_{a'_i \in A_i} \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}} [v_i(\theta; a'_i, s_{-i}(\theta_{-i}))] \right] \right] - \epsilon \quad (\text{from the definition of } \phi) \\
&\geq \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i} \left[ \mathbb{E}_{\theta_{-i} \sim \rho | \theta_i} \left[ \mathbb{E}_{s_{-i} \sim \sigma_{-i}} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i}))] \right] \right] \right] - \epsilon \\
&= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; (\phi_{\text{SF}}(s_i))(\theta_i), s_{-i}(\theta_{-i}))] \right] - \epsilon,
\end{aligned}$$

which implies (IC<sub>BNE</sub>) for any  $\phi_{\text{SF}}: S_i \rightarrow S_i$ .  $\square$

Finally, we relate the classes of Bayes correlated equilibria to Bayes Nash equilibria. The following proposition shows that if a type-wise product distribution is contained in each class of Bayes correlated equilibria, it is a Bayes Nash equilibrium.

**Proposition 2.27.**  $\Pi_{\text{SF}}^\epsilon \cap \Pi_{\text{Prod}} = \Pi_{\text{ANF}}^\epsilon \cap \Pi_{\text{Prod}} = \Pi_{\text{Com}}^\epsilon \cap \Pi_{\text{Prod}} = \Pi_{\text{BS}}^\epsilon \cap \Pi_{\text{Prod}} = \Pi_{\text{BNE}}^\epsilon$ .

*Proof.* From Theorems 2.20 and 2.23,  $\Pi_{\text{SF}}^\epsilon$  is a subset of  $\Pi_{\text{ANF}}^\epsilon$ ,  $\Pi_{\text{Com}}^\epsilon$ , and  $\Pi_{\text{BS}}^\epsilon$ . Moreover,  $\Pi_{\text{BS}}^\epsilon$  is a superset of  $\Pi_{\text{SF}}^\epsilon$ ,  $\Pi_{\text{ANF}}^\epsilon$ , and  $\Pi_{\text{Com}}^\epsilon$ . Therefore, it is sufficient to show that  $\Pi_{\text{BNE}}^\epsilon \subseteq \Pi_{\text{SF}}^\epsilon \cap \Pi_{\text{Prod}}$  and  $\Pi_{\text{BS}}^\epsilon \cap \Pi_{\text{Prod}} \subseteq \Pi_{\text{BNE}}^\epsilon$ .

First, we prove  $\Pi_{\text{BNE}}^\epsilon \subseteq \Pi_{\text{SF}}^\epsilon \cap \Pi_{\text{Prod}}$ . Suppose  $\pi \in \Pi_{\text{BNE}}^\epsilon$ . From Theorem 2.26, there exists  $\sigma \in \Sigma_{\text{BNE}}^\epsilon$  such that  $\eta(\sigma) = \pi$ . Since this  $\sigma$  satisfies (IC<sub>BNE</sub>) for each  $i \in N$  and  $\phi_{\text{SF}}: S_i \rightarrow S_i$ , which is equivalent to (IC<sub>SF</sub>). Hence,  $\sigma \in \Sigma_{\text{SF}}^\epsilon$ , and then  $\pi = \eta(\sigma) \in \Pi_{\text{SF}}^\epsilon$ . Moreover, from the definition of  $\epsilon$ -BNEs in the agent normal form,  $\pi$  is a type-wise product distribution, which implies  $\Pi_{\text{BNE}}^\epsilon \subseteq \Pi_{\text{SF}}^\epsilon \cap \Pi_{\text{Prod}}$ .

Next, we prove  $\Pi_{\text{BS}}^\epsilon \cap \Pi_{\text{Prod}} \subseteq \Pi_{\text{BNE}}^\epsilon$ . Suppose  $\pi \in \Pi_{\text{BS}}^\epsilon \cap \Pi_{\text{Prod}}$ . Since  $\pi \in \Pi_{\text{BS}}^\epsilon$ , the incentive constraint (IC<sub>BS</sub>) holds for each  $i \in N$  and  $\phi: \Theta_i \times A_i$ , which is equivalent to (IC'<sub>BNE</sub>). Moreover,  $\pi$  is a type-wise product distribution since  $\pi \in \Pi_{\text{Prod}}$ . From the definition of  $\epsilon$ -BNEs in the agent normal form, we obtain  $\pi \in \Pi_{\text{BNE}}^\epsilon$ .  $\square$

### 3 Algorithm for minimizing untruthful swap regret

In this section, we describe an efficient algorithm for minimizing untruthful swap regret, which leads to efficient computation of approximate communication equilibria. In this section, we focus on the problem that a player  $i \in N$  is faced with, and therefore, the subscript  $i$  for  $A_i$ ,  $a_i$ ,  $\Theta_i$ ,  $\theta_i$ ,  $\pi_i^t$ , and  $R_{\text{US},i}^T$  is not essential in this section. We put the subscript  $i$  just for notational consistency.

We derive the algorithm in three steps: First, we reduce untruthful swap regret minimization to a  $\Phi$ -regret minimization problem. Then we apply the  $\Phi$ -regret minimization framework developed by Gordon et al. [2008] and reduce it to an online linear optimization problem. Finally, to obtain a regret upper bound optimal up to a multiplicative constant, we further decompose it to small external regret minimization problems. The final algorithm is described in Algorithm 3.

---

**Algorithm 3:** Algorithm for minimizing untruthful swap regret
 

---

**Input:** The set of types  $\Theta_i$  and the set of actions  $A_i$  are specified in advance. The reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  is given at the end of each round  $t \in [T]$ .

Initialize subroutines as follows:

- let  $\mathcal{E}_{\theta_i}$  be a multiplicative weights algorithm with decision space  $\Theta_i$  for each  $\theta_i \in \Theta_i$ ;
- let  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$  be a multiplicative weights algorithm with the doubling trick [Cesa-Bianchi et al., 1997, Section 4.6] with decision space  $A_i$  for each  $\theta_i, \theta'_i \in \Theta_i$  and  $a'_i \in A_i$ ;

**for** each round  $t = 1, \dots, T$  **do**

Let  $w_{\theta_i}^t \in \Delta(\Theta_i)$  be the output of  $\mathcal{E}_{\theta_i}$  in round  $t$  for each  $\theta_i \in \Theta_i$ ;  
 Let  $y_{\theta_i, \theta'_i, a'_i}^t \in \Delta(A_i)$  be the output of  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$  in round  $t$  for each  $\theta_i, \theta'_i \in \Theta_i$  and  $a'_i \in A_i$ ;  
 Define  $Q^t \in [0, 1]^{(\Theta_i \times A_i) \times (\Theta_i \times A_i)}$  by  $Q^t((\theta_i, a_i), (\theta'_i, a'_i)) = w_{\theta_i}^t(\theta'_i) y_{\theta_i, \theta'_i, a'_i}^t(a_i)$  for each  $\theta_i, \theta'_i \in \Theta_i$  and  $a_i, a'_i \in A_i$ ;  
 Compute an eigenvector  $x^t \in \mathbb{R}^{\Theta_i \times A_i}$  of  $Q^t$  such that  $Q^t x^t = x^t$  and  $(x^t)^\top \mathbf{1} = |\Theta_i|$ ;  
 Decide the output  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$  by  $\pi_i^t(\theta_i; a_i) = x^t(\theta_i, a_i)$  for each  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ ;  
 Observe reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  and feed reward vectors to subroutines as follows:

- feed  $\sum_{a_i, a'_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) \pi_i^t(\theta'_i; a'_i) \rho_i(\theta_i) u_i^t(\theta_i, a_i)$  as the reward for decision  $\theta'_i \in \Theta_i$  to subroutine  $\mathcal{E}_{\theta_i}$  for each  $\theta_i \in \Theta_i$ ;
- feed  $\pi_i^t(\theta'_i; a'_i) \rho_i(\theta_i) u_i^t(\theta_i, a_i)$  as the reward for decision  $a_i \in A_i$  to subroutine  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$  for each  $\theta_i, \theta'_i \in \Theta_i$  and  $a'_i \in A_i$ ;

**end**

---

### 3.1 Reduction to $\Phi$ -regret minimization

First, we interpret untruthful swap regret minimization as  $\Phi$ -regret minimization with decision space  $\mathcal{X} \subseteq [0, 1]^{\Theta_i \times A_i}$  instead of  $\Delta(A_i)^{\Theta_i}$ , where

$$\mathcal{X} = \left\{ x \in [0, 1]^{\Theta_i \times A_i} \mid \sum_{a_i \in A_i} x(\theta_i, a_i) = 1 \ (\forall \theta_i \in \Theta_i) \right\} \quad (5)$$

is the set of vectors corresponding to  $\Delta(A_i)^{\Theta_i}$ . Each  $\pi_i \in \Delta(A_i)^{\Theta_i}$  is associated to some  $x \in \mathcal{X}$  by  $x(\theta_i, a_i) = \pi_i(\theta_i; a_i)$  for each  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ . Let  $x^t \in \mathcal{X}$  be the output of the algorithm in each round  $t \in [T]$  that corresponds to  $\pi_i^t$ . Let  $\bar{u}^t \in [0, 1]^{\Theta_i \times A_i}$  be the reward vector weighted by prior probabilities  $\rho_i$ , i.e.,  $\bar{u}^t(\theta_i, a_i) = \rho_i(\theta_i) u_i^t(\theta_i, a_i)$  for each  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ .

Now we prove that the competitor is supposed to choose an optimal transformation of all  $(x^t)_{t=1}^T$  from the set of linear transformations defined by

$$\mathcal{Q} = \left\{ Q \in [0, 1]^{(\Theta_i \times A_i) \times (\Theta_i \times A_i)} \mid \begin{array}{l} \text{there exists some } W \in [0, 1]^{\Theta_i \times \Theta_i} \text{ such that} \\ \sum_{\theta'_i \in \Theta_i} W(\theta_i, \theta'_i) = 1 \ (\forall \theta_i \in \Theta_i) \text{ and} \\ \sum_{a_i \in A_i} Q((\theta_i, a_i), (\theta'_i, a'_i)) = W(\theta_i, \theta'_i) \ (\forall \theta_i, \theta'_i \in \Theta_i, a'_i \in A_i) \end{array} \right\}. \quad (6)$$

**Lemma 3.1.** *The untruthful swap regret can be written as*

$$R_{\text{US},i} = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Qx^t, \bar{u}^t \rangle - \sum_{t=1}^T \langle x^t, \bar{u}^t \rangle.$$

*Proof.* The expected reward for the algorithm is

$$\sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, a_i)] \right] = \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{a_i \in A_i} \pi_i^t(\theta_i; a_i) u_i^t(\theta_i, a_i) = \sum_{t=1}^T \langle x^t, \bar{u}^t \rangle.$$

Next, we express the competitor's transformation that uses  $\psi: \Theta_i \rightarrow \Theta_i$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$  as a linear transformation. Given  $\pi_i^t \in \Delta(A_i)^{\Theta_i}$ , the competitor's expected reward in each round  $t \in [T]$  is

$$\begin{aligned} \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta_i, \phi(\theta_i, a_i))] \right] &= \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{a_i \in A_i} \pi_i^t(\psi(\theta_i); a_i) u_i^t(\theta_i, \phi(\theta_i, a_i)) \\ &= \sum_{\theta_i \in \Theta_i} \sum_{a_i \in A_i} x^t(\psi(\theta_i), a_i) \bar{u}^t(\theta_i, \phi(\theta_i, a_i)) \\ &= \langle Q_\psi x^t, Q_\phi \bar{u}^t \rangle, \end{aligned}$$

where  $Q_\psi, Q_\phi \in [0, 1]^{(\Theta_i \times A_i) \times (\Theta_i \times A_i)}$  are the matrices defined as follows.

- The linear transformation  $Q_\psi$  for  $x^t$  is defined as the Kronecker product  $Q_\psi = W_\psi \otimes I$  of the zero-one stochastic matrix  $W_\psi \in \{0, 1\}^{\Theta_i \times \Theta_i}$ , where  $W_\psi(\theta_i, \theta'_i) = 1$  if  $\theta'_i = \psi(\theta_i)$  and 0 otherwise, and the identity matrix  $I \in \{0, 1\}^{A_i \times A_i}$ .
- The linear transformation  $Q_\phi$  for  $\bar{u}^t$  is defined as the block diagonal matrix, where each block  $Q_{\phi(\theta_i, \cdot)} \in \{0, 1\}^{A_i \times A_i}$  corresponding to  $\theta_i \in \Theta_i$  is the zero-one stochastic matrix defined by  $Q_{\phi(\theta_i, \cdot)}(a_i, a'_i) = 1$  if  $a'_i = \phi(\theta_i, a_i)$  and 0 otherwise.

Since  $\langle Q_\psi x^t, Q_\phi \bar{u}^t \rangle = \left\langle \left( Q_\phi^\top Q_\psi \right) x^t, \bar{u}^t \right\rangle$ , the competitor can be interpreted as transforming  $x^t$  by using matrix  $Q_\phi^\top Q_\psi$ , which we denote by  $Q_{\psi, \phi}$ . Hence, the set of all possible transformations for the competitor is

$$\mathcal{Q}' = \{Q_{\psi, \phi} \mid \psi: \Theta_i \rightarrow \Theta_i, \phi: \Theta_i \times A_i \rightarrow A_i\}.$$

Next, we prove that  $\mathcal{Q}$  equals the convex hull of  $\mathcal{Q}'$ . For each  $\psi$  and  $\phi$ , by a simple calculation, we can see that each entry is  $Q_{\psi, \phi}((\theta_i, a_i), (\theta'_i, a'_i)) = 1$  if  $\theta'_i = \psi(\theta_i)$  and  $a_i = \phi(\theta_i, a'_i)$ , and 0 otherwise. All the constraints (6) for  $\mathcal{Q}$  are satisfied by  $Q_{\psi, \phi} \in \mathcal{Q}$  for any  $\psi$  and  $\phi$ . The convex hull is therefore a subset of  $\mathcal{Q}$ .

Conversely, we show that any  $Q \in \mathcal{Q}$  can be expressed as a convex combination of  $Q_{\psi, \phi}$ . Fix any  $Q \in \mathcal{Q}$ . We can regard  $Q$  as a block matrix with block  $Q_{\theta_i, \theta'_i} \in [0, 1]^{A_i \times A_i}$  for each  $\theta_i, \theta'_i \in \Theta_i$ . From the definition of  $\mathcal{Q}$ , there exists some stochastic matrices  $W \in [0, 1]^{\Theta_i \times \Theta_i}$  and  $\tilde{Q}_{\theta_i, \theta'_i} \in [0, 1]^{A_i \times A_i}$  for each  $\theta_i, \theta'_i \in \Theta_i$  such that  $Q_{\theta_i, \theta'_i}^\top = W(\theta_i, \theta'_i) \tilde{Q}_{\theta_i, \theta'_i}$  for each  $\theta_i, \theta'_i \in \Theta_i$ . Since any stochastic matrix is a convex combination of zero-one stochastic matrices, there exists some distribution  $\gamma$  over all

possible  $\psi$  such that  $W = \sum_{\psi} \gamma(\psi) W_{\psi}$ . Similarly, for each  $\theta_i, \theta'_i \in \Theta_i$ , there exists some distribution  $\kappa_{\theta_i, \theta'_i}$  over all possible  $\phi(\theta_i, \cdot)$  such that  $\tilde{Q}_{\theta_i, \theta'_i} = \sum_{\phi(\theta_i, \cdot)} \kappa_{\theta_i, \theta'_i}(\phi(\theta_i, \cdot)) Q_{\phi(\theta_i, \cdot)}$ . Each entry of  $Q$  is

$$\begin{aligned} Q((\theta_i, a_i), (\theta'_i, a'_i)) &= W(\theta_i, \theta'_i) \tilde{Q}_{\theta_i, \theta'_i}(a'_i, a_i) \\ &= \left( \sum_{\psi} \gamma(\psi) W_{\psi}(\theta_i, \theta'_i) \right) \left( \sum_{\phi(\theta_i, \cdot)} \kappa_{\theta_i, \theta'_i}(\phi(\theta_i, \cdot)) Q_{\phi(\theta_i, \cdot)}(a'_i, a_i) \right) \\ &= \left( \sum_{\psi} \gamma(\psi) \mathbf{1}_{\{\theta'_i = \psi(\theta_i)\}} \right) \left( \sum_{\phi(\theta_i, \cdot)} \kappa_{\theta_i, \theta'_i}(\phi(\theta_i, \cdot)) \mathbf{1}_{\{a_i = \phi(\theta_i, a'_i)\}} \right) \\ &= \Pr_{\psi \sim \gamma} (\theta'_i = \psi(\theta_i)) \Pr_{\phi(\theta_i, \cdot) \sim \kappa_{\theta_i, \theta'_i}} (a_i = \phi(\theta_i, a'_i)). \end{aligned}$$

We consider a distribution that generates  $\psi$  and  $\phi$  with probability  $\gamma(\psi) \prod_{\theta_i \in \Theta_i} \kappa_{\theta_i, \psi(\theta_i)}(\phi(\theta_i, \cdot))$ . That is,  $\psi$  is generated from  $\gamma$  and then each  $\phi(\theta_i, \cdot)$  is generated from  $\kappa_{\theta_i, \psi(\theta_i)}$  independently for each  $\theta_i \in \Theta_i$ . The expected value of each entry according to this distribution is

$$\begin{aligned} \mathbb{E}_{\psi, \phi} [Q_{\psi, \phi}((\theta_i, a_i), (\theta'_i, a'_i))] &= \Pr_{\psi, \phi} (\theta'_i = \psi(\theta_i), a_i = \phi(\theta_i, a'_i)) \\ &= \Pr_{\psi \sim \gamma} (\theta'_i = \psi(\theta_i)) \Pr_{\phi(\theta_i, \cdot) \sim \kappa_{\theta_i, \psi(\theta'_i)}} (a_i = \phi(\theta_i, a'_i)) \\ &= Q((\theta_i, a_i), (\theta'_i, a'_i)). \end{aligned}$$

Since every  $Q \in \mathcal{Q}$  can be expressed as a convex combination of some matrices in  $\mathcal{Q}'$ , we can see that  $\mathcal{Q}$  is a subset of the convex hull of  $\mathcal{Q}'$ . We conclude that  $\mathcal{Q}$  equals the convex hull of  $\mathcal{Q}'$ .

The competitor's expected reward is

$$\begin{aligned} &\max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta_i, \phi(\theta_i, a_i))] \right] \\ &= \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \langle Q_{\psi, \phi} x^t, \bar{u}^t \rangle \\ &= \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q x^t, \bar{u}^t \rangle, \end{aligned}$$

where the last equality holds because  $\langle Q x^t, \bar{u}^t \rangle$  is linear in terms of  $Q$ , and  $\mathcal{Q}$  is the convex hull of  $Q_{\psi, \phi}$  for all  $\psi$  and  $\phi$ .  $\square$

### 3.2 Reduction to online linear optimization

We apply the framework developed by Gordon et al. [2008] that reduces  $\Phi$ -regret minimization to fixed point computation and online linear optimization. First, we show that each  $Q \in \mathcal{Q}$  has a fixed point in  $\mathcal{X}$ .

**Lemma 3.2.** *For every  $Q \in \mathcal{Q}$ , there exists  $x \in \mathcal{X}$  such that  $Qx = x$ .*

*Proof.* We prove this lemma by using Brouwer’s fixed point theorem, which claims that any continuous function that maps a compact convex set to itself has a fixed point. Since  $\mathcal{X}$  is a product of probability simplices, it is compact and convex. In addition,  $Q$  is continuous since it is a matrix. Thus it is sufficient to prove that  $Q$  maps every  $x \in \mathcal{X}$  to a vector in  $\mathcal{X}$ . Let  $x \in \mathcal{X}$  be an arbitrary vector. Let  $x_{\theta_i}$  be the block of  $x$  corresponding to  $\theta_i$  and  $Q_{\theta_i, \theta'_i}$  be the block of  $Q$  corresponding to  $\theta_i, \theta'_i$ . Then  $x_{\theta_i} \in \Delta(A_i)$  is a stochastic vector for each  $\theta_i \in \Theta_i$ . From the definition of  $Q$ , there exists some stochastic matrices  $W \in [0, 1]^{\Theta_i \times \Theta_i}$  and  $\tilde{Q}_{\theta_i, \theta'_i} \in [0, 1]^{A_i \times A_i}$  for each  $\theta_i, \theta'_i$  such that  $Q_{\theta_i, \theta'_i}^\top = W(\theta_i, \theta'_i) \tilde{Q}_{\theta_i, \theta'_i}$  for each  $\theta_i, \theta'_i \in \Theta_i$ . Since  $\tilde{Q}_{\theta_i, \theta'_i}^\top x_{\theta'_i} \in \Delta(A_i)$  is also a stochastic vector, each block of  $Qx$  corresponding to  $\theta_i \in \Theta_i$  is

$$\sum_{\theta'_i \in \Theta_i} Q_{\theta_i, \theta'_i} x_{\theta'_i} = \sum_{\theta'_i \in \Theta_i} W(\theta_i, \theta'_i) \left( \tilde{Q}_{\theta_i, \theta'_i}^\top x_{\theta'_i} \right).$$

This is a convex sum of stochastic vectors, hence a stochastic vector in  $\Delta(A_i)$ . This implies  $Qx \in \mathcal{X}$  for any  $x \in \mathcal{X}$ . From Brouwer’s fixed point theorem, there exists a fixed point  $x \in \mathcal{X}$  such that  $Qx = x$ .  $\square$

Next, we show that this fixed point can be obtained by eigenvector computation. As described later, all the subroutines are variants of the multiplicative weights, and therefore, all entries of their outputs are positive. We can assume that each entry of  $Q^t$ , which is computed as a product of the subroutines’ outputs, is positive for every  $t \in [T]$ .

**Lemma 3.3.** *If all entries of  $Q \in \mathcal{Q}$  are positive, then we can compute  $x \in \mathcal{X}$  such that  $Qx = x$  in time polynomial in  $|\Theta_i|$  and  $|A_i|$ .*

*Proof.* We apply the Perron–Frobenius theorem, which claims that any positive matrix has an eigenvalue whose eigenspace is 1-dimensional, and any positive eigenvector is an eigenvector corresponding to this eigenvalue. From Theorem 3.2, there exists some  $x \in \mathcal{X}$  such that  $Qx = x$ , which implies that  $Q$  has eigenvector 1, and  $x$  is its corresponding eigenvector. Since  $Q$  is a positive matrix and  $x$  is a non-negative non-zero vector,  $Qx = x$  is also a positive vector. From the Perron–Frobenius theorem, the eigenspace corresponding to eigenvalue 1 is 1-dimensional and contains only a multiple of  $x \in \mathcal{X}$ . We can obtain  $x \in \mathcal{X}$  by computing an eigendecomposition of  $Q$  and appropriately scaling the eigenvector corresponding to eigenvalue 1. Since  $Q$  is a  $|\Theta_i||A_i| \times |\Theta_i||A_i|$  matrix, its eigendecomposition can be computed in  $O(|\Theta_i|^\omega |A_i|^\omega)$  time, where  $\omega$  is the exponent of matrix multiplication.  $\square$

By using this fixed point computation, we can solve untruthful swap regret minimization with a subroutine for online linear optimization with decision space  $\mathcal{Q}$ . Let  $Q^t \in \mathcal{Q}$  be the subroutine’s output in round  $t \in [T]$ . Then compute a fixed point  $x^t \in \mathcal{Q}$  that satisfies  $Q^t x^t = x^t$ , and let  $x^t$  be the decision for round  $t$ . For the subroutine, feed reward  $U^t \in [0, 1]^{(\Theta_i \times A_i) \times (\Theta_i \times A_i)}$  defined by  $U^t((\theta_i, a_i), (\theta'_i, a'_i)) = \bar{u}^t(\theta_i, a_i) x^t(\theta'_i, a'_i)$  for each  $\theta_i, \theta'_i \in \Theta_i$  and  $a_i, a'_i \in A_i$ . We define external regret for this subroutine as

$$R_{\mathcal{Q}}^T = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q, U^t \rangle - \sum_{t=1}^T \langle Q^t, U^t \rangle.$$

Then we show that untruthful swap regret  $R_{\text{US},i}^T$  equals this subroutine’s external regret  $R_{\mathcal{Q}}^T$ .

**Lemma 3.4.**  $R_{\text{US},i}^T = R_{\mathcal{Q}}^T$ .

*Proof.* From the definition of  $x^t$ , it holds that  $Q^t x^t = x^t$  for each  $t \in [T]$ . From the definition of  $U^t$ , the algorithm's expected reward is

$$\sum_{t=1}^T \langle Q^t, U^t \rangle = \sum_{t=1}^T \langle Q^t, \bar{u}^t (x^t)^\top \rangle = \sum_{t=1}^T \langle Q^t x^t, \bar{u}^t \rangle = \sum_{t=1}^T \langle x^t, \bar{u}^t \rangle,$$

where the last equality is due to  $Q^t x^t = x^t$  for each  $t \in [T]$ . Similarly, the competitor's expected reward is

$$\max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q, U^t \rangle = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q, \bar{u}^t (x^t)^\top \rangle = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q x^t, \bar{u}^t \rangle.$$

Therefore, we obtain

$$R_{\mathcal{Q}}^T = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q, U^t \rangle - \sum_{t=1}^T \langle Q^t, U^t \rangle = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q x^t, \bar{u}^t \rangle - \sum_{t=1}^T \langle x^t, \bar{u}^t \rangle = R_{\text{US},i},$$

where the last equality is due to Theorem 3.1.  $\square$

### 3.3 Decomposition into small external regret minimization problems

Next, we reduce online linear optimization with decision space  $\mathcal{Q}$  to small external regret minimization problems. For each  $\theta_i, \theta'_i \in \Theta_i$  and  $a'_i \in A_i$ , we use a subroutine  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$  with decision space  $A_i$ . Moreover, for each  $\theta_i \in \Theta_i$ , we use a subroutine  $\mathcal{E}_{\theta_i}$  with decision space  $\Theta_i$ . In total, we use  $|A_i| |\Theta_i|^2 + |\Theta_i|$  subroutines.

The reduction proceeds as follows. Let  $y_{\theta_i, \theta'_i, a'_i}^t \in \Delta(A_i)$  be the output of  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$  and  $w_{\theta_i}^t \in \Delta(\Theta_i)$  the output of  $\mathcal{E}_{\theta_i}$  in round  $t \in [T]$ . According to these outputs, we set  $Q^t((\theta_i, a_i), (\theta'_i, a'_i)) = w_{\theta_i}^t(\theta'_i) y_{\theta_i, \theta'_i, a'_i}^t(a_i)$  for each  $\theta_i, \theta'_i \in \Theta_i$  and  $a_i, a'_i \in A_i$ . Note that this  $Q^t$  is contained in  $\mathcal{Q}$ , which can be checked by setting  $W(\theta_i, \theta'_i) = w_{\theta_i}^t(\theta'_i)$  for each  $\theta_i, \theta'_i \in \Theta_i$ .

In each round  $t \in [T]$ , based on the observed reward  $U^t = \bar{u}^t (x^t)^\top$ , we feed the reward for each subroutine as follows. For each  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$ , the reward for decision  $a_i \in A_i$  is  $x^t(\theta'_i, a'_i) \bar{u}_i^t(\theta_i, a_i)$ . The external regret  $R_{\theta_i, \theta'_i, a'_i}^T$  for  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$  is defined as

$$R_{\theta_i, \theta'_i, a'_i}^T = \max_{a_i^* \in A_i} \sum_{t=1}^T x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i^*) - \sum_{t=1}^T \sum_{a_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i).$$

For each  $\mathcal{E}_{\theta_i}$ , the reward for decision  $\theta'_i \in \Theta_i$  is  $\sum_{a_i, a'_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i)$ . The external regret  $R_{\theta_i}^T$  for  $\mathcal{E}_{\theta_i}$  is defined as

$$\begin{aligned} R_{\theta_i}^T &= \max_{\theta_i^* \in \Theta_i} \sum_{t=1}^T \sum_{a_i, a'_i \in A_i} y_{\theta_i, \theta_i^*, a'_i}^t(a_i) x^t(\theta_i^*, a'_i) \bar{u}^t(\theta_i, a_i) \\ &\quad - \sum_{t=1}^T \sum_{\theta'_i \in \Theta_i} w_{\theta_i}^t(\theta'_i) \sum_{a_i, a'_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i). \end{aligned}$$

The external regret for online linear optimization with decision space  $\mathcal{Q}$  is bounded by the sum of the external regrets for these subroutines as follows.

**Lemma 3.5.**

$$R_{\mathcal{Q}}^T \leq \sum_{\theta_i \in \Theta_i} R_{\theta_i}^T + \sum_{\theta_i \in \Theta_i} \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} R_{\theta_i, \theta'_i, a'_i}^T.$$

*Proof.* Recall that  $U^t = \bar{u}^t(x^t)^\top$ . Since each extreme point of  $\mathcal{Q}$  corresponds to a matrix  $(Q_\phi^\top Q_\psi)$  for some  $\psi$  and  $\phi$ , we obtain

$$\begin{aligned} & \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q, U^t \rangle \\ &= \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q x^t, \bar{u}^t \rangle \\ &= \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \langle Q_\psi x^t, Q_\phi \bar{u}^t \rangle \\ &= \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{\theta_i \in \Theta_i} \sum_{a'_i \in A_i} \sum_{t=1}^T x^t(\psi(\theta_i), a'_i) \bar{u}^t(\theta_i, \phi(\theta_i, a'_i)) \\ &= \sum_{\theta_i \in \Theta_i} \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} \max_{a_i \in A_i} \sum_{t=1}^T x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i) \\ &= \sum_{\theta_i \in \Theta_i} \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} \left\{ \sum_{t=1}^T \sum_{a_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i) + R_{\theta_i, \theta'_i, a'_i}^T \right\} \\ & \hspace{15em} \text{(due to the definition of } R_{\theta_i, \theta'_i, a'_i}^T) \\ &\leq \sum_{\theta_i \in \Theta_i} \left\{ \max_{\theta'_i \in \Theta_i} \sum_{t=1}^T \sum_{a'_i \in A_i} \sum_{a_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i) + \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} R_{\theta_i, \theta'_i, a'_i}^T \right\} \\ &= \sum_{\theta_i \in \Theta_i} \left\{ \sum_{t=1}^T \sum_{\theta'_i \in \Theta_i} w_{\theta_i}^t(\theta'_i) \sum_{a'_i \in A_i} \sum_{a_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i) + R_{\theta_i}^T + \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} R_{\theta_i, \theta'_i, a'_i}^T \right\} \\ & \hspace{15em} \text{(due to the definition of } R_{\theta_i}^T) \\ &= \sum_{t=1}^T \langle Q^t, U^t \rangle + \sum_{\theta_i \in \Theta_i} R_{\theta_i}^T + \sum_{\theta_i \in \Theta_i} \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} R_{\theta_i, \theta'_i, a'_i}^T, \quad \text{(due to the definition of } Q^t \text{ and } U^t) \end{aligned}$$

which yields an upper bound on  $R_{\mathcal{Q}}^T = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T \langle Q, U^t \rangle - \sum_{t=1}^T \langle Q^t, U^t \rangle$ .  $\square$

To obtain an upper bound on untruthful swap regret, we use the following upper bounds for external regret minimization. These algorithms are based on the multiplicative weights, and their output is always a positive vector.

**Theorem 3.6** ([Cesa-Bianchi and Lugosi, 2006, Theorem 2.2]). *For an online learning problem with rewards in  $[0, 1]$ , there exists an algorithm such that its external regret is bounded above by  $\sqrt{\frac{1}{2}T \log d}$ , where  $d$  is the number of possible decisions.*

**Theorem 3.7** ([Cesa-Bianchi et al., 1997, Theorem 4.6.3]). *For an online learning problem with rewards in  $[0, 1]$ , there exists an algorithm such that its external regret is bounded above by  $6\sqrt{u^* \log d} + 2 \log d$ , where  $d$  is the number of possible decisions and  $u^*$  is the sum of rewards for an optimal fixed decision over  $T$  rounds.<sup>13</sup>*

These three-step reduction yields the algorithm described in Algorithm 3. Finally, we obtain an upper bound on untruthful swap regret as follows.

**Theorem 3.8.** *The untruthful swap regret of Algorithm 3 is bounded as*

$$R_{\text{US},i}^T \leq \sqrt{\frac{1}{2}T \log |\Theta_i|} + 6\sqrt{T|A_i| \log |A_i|} + 2|A_i| \log |A_i|.$$

*Proof.* From Theorems 3.4 and 3.5, the untruthful swap regret of Algorithm 3 is bounded as

$$R_{\text{US},i}^T = R_{\mathcal{Q}}^T \leq \sum_{\theta_i \in \Theta_i} R_{\theta_i}^T + \sum_{\theta_i \in \Theta_i} \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} R_{\theta_i, \theta'_i, a'_i}^T.$$

Since the reward input to  $\mathcal{E}_{\theta_i}$  is bounded above by

$$\max_{t \in [T]} \max_{\theta'_i \in \Theta_i} \sum_{a_i, a'_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \bar{u}^t(\theta_i, a_i) \leq \max_{t \in [T]} \max_{\theta'_i \in \Theta_i} \sum_{a_i, a'_i \in A_i} y_{\theta_i, \theta'_i, a'_i}^t(a_i) x^t(\theta'_i, a'_i) \rho_i(\theta_i) = \rho_i(\theta_i),$$

we can obtain the regret upper bound  $R_{\theta_i}^T \leq \rho_i(\theta_i) \sqrt{\frac{1}{2}T \log |\Theta_i|}$  from Theorem 3.6. Since the reward input to  $\mathcal{E}_{\theta_i, \theta'_i, a'_i}$  is similarly bounded above by  $\rho_i(\theta_i)$  and the total reward for each  $a_i \in A_i$  is bounded above as

$$\sum_{t=1}^T x^t(\theta'_i, a'_i) \bar{u}_i^t(\theta_i, a_i) \leq \sum_{t=1}^T \pi_i^t(\theta'_i; a'_i) \rho_i(\theta_i) = \rho_i(\theta_i) \sum_{t=1}^T \pi_i^t(\theta'_i; a'_i),$$

we can obtain a regret upper bound  $R_{\theta_i, \theta'_i, a'_i}^T = 6\rho_i(\theta_i) \sqrt{\sum_{t=1}^T \pi_i^t(\theta'_i; a'_i) \log |A_i|} + 2\rho_i(\theta_i) \log |A_i|$ . By the Cauchy–Schwarz inequality, we have  $\sum_{a'_i \in A_i} 1 \cdot \sqrt{\sum_{t=1}^T \pi_i^t(\theta'_i; a'_i)} \leq \sqrt{|A_i| \sum_{a'_i \in A_i} \sum_{t=1}^T \pi_i^t(\theta'_i; a'_i)} = \sqrt{|A_i|T}$ . We thus obtain

$$\begin{aligned} & \sum_{\theta_i \in \Theta_i} R_{\theta_i}^T + \sum_{\theta_i \in \Theta_i} \max_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} R_{\theta_i, \theta'_i, a'_i}^T \\ & \leq \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sqrt{\frac{1}{2}T \log |\Theta_i|} + \sum_{\theta_i \in \Theta_i} \max_{\theta'_i \in \Theta_i} \rho_i(\theta_i) \left( 6\sqrt{T|A_i| \log |A_i|} + 2|A_i| \log |A_i| \right) \\ & = \sqrt{\frac{1}{2}T \log |\Theta_i|} + 6\sqrt{T|A_i| \log |A_i|} + 2|A_i| \log |A_i|. \quad \square \end{aligned}$$

<sup>13</sup>For simplicity, the constants are not optimized. We can instead use an algorithm with adaptive step size proposed by Auer et al. [2002].

### 3.4 Application to equilibrium computation

Now we apply the algorithm for untruthful swap regret minimization to equilibrium computation. Theorem 2.17 claims that if each player minimizes the untruthful swap regret, then the dynamics converge to an approximate communication equilibrium with strategy representability. However, exactly simulating the dynamics requires the exact evaluation of rewards  $u_i^t$ , which needs exponential time in general. Instead, we estimate these values by random sampling and bound the error by using Hoeffding's inequality.

**Corollary 3.9.** *For any  $\epsilon > 0$ , there exists an algorithm that outputs a succinct representation of an  $\epsilon$ -approximate communication equilibrium with strategy representability with probability at least  $1 - \delta$  and runs in time polynomial in  $n$ ,  $\max_{i \in N} |\Theta_i|$ ,  $\max_{i \in N} |A_i|$ ,  $1/\epsilon$ , and  $\log(1/\delta)$  with an oracle for utility functions.*

*Proof.* We estimate the value of  $u_i^t(\theta_i, a_i)$  for each  $i \in N$ ,  $t \in [T]$ ,  $\theta_i \in \Theta_i$ , and  $a_i \in A_i$  by generating  $\frac{8}{\epsilon^2} \log \frac{2nT \max_{i \in N} |\Theta_i| |A_i|}{\delta}$  samples of  $\theta_{-i} \sim \rho|\theta_i$  and  $a_{-i} \sim \pi_{-i}^t(\theta_{-i})$  and taking the average of  $v_i(\theta_i, \theta_{-i}; a_i, a_{-i})$ , where  $T$  will be set later. Let  $\tilde{u}_i^t(\theta_i, a_i)$  be this estimation. By using Hoeffding's inequality, we have

$$\begin{aligned} \Pr \left( |u_i^t(\theta_i, a_i) - \tilde{u}_i^t(\theta_i, a_i)| \geq \frac{\epsilon}{4} \right) &\leq 2 \exp \left( -\frac{2\epsilon^2}{16} \cdot \frac{8}{\epsilon^2} \log \frac{2nT \max_{i \in N} |\Theta_i| |A_i|}{\delta} \right) \\ &\leq \frac{\delta}{nT \max_{i \in N} (|\Theta_i| |A_i|)} \end{aligned}$$

for each  $i \in N$ ,  $t \in [T]$ ,  $\theta_i \in \Theta_i$ , and  $a_i \in A_i$ . By the union bound, with probability at least  $1 - \delta$ , the additive errors of all the rewards are at most  $\epsilon/4$ .

We apply Algorithm 3 to these estimated rewards. For each player  $i \in N$ , by using the upper bound on untruthful swap regret in Theorem 3.8, we obtain

$$\begin{aligned} \frac{R_{\text{US},i}}{T} &= \frac{1}{T} \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta_i, \phi(\theta_i, a_i))] - \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, a_i)] \right] \\ &\leq \frac{1}{T} \left\{ \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{a_i \sim \pi_i^t(\psi(\theta_i))} [\tilde{u}_i^t(\theta_i, \phi(\theta_i, a_i))] - \mathbb{E}_{a_i \sim \pi_i^t(\theta_i)} [\tilde{u}_i^t(\theta_i, a_i)] \right] + T \cdot \frac{\epsilon}{2} \right\} \\ &\leq \sqrt{\frac{1}{2T} \log |\Theta_i|} + 6\sqrt{\frac{|A_i| \log |A_i|}{T}} + \frac{2}{T} |A_i| \log |A_i| + \frac{\epsilon}{2} \\ &\leq \epsilon, \end{aligned}$$

where we set  $T = \max\{\frac{18}{\epsilon^2} \log(\max_{i \in N} |\Theta_i|), \frac{1296}{\epsilon^2} \max_{i \in N} |A_i| \log |A_i|\}$ . From Theorem 2.17,  $\pi$  is an  $\epsilon$ -approximate communication equilibrium with strategy representability.  $\square$

## 4 Lower bound for untruthful swap regret

Here we show an  $\Omega(\sqrt{T \log |\Theta_i|})$  lower bound on untruthful swap regret. As in the previous section, this section focuses on an online learning problem with stochastic types for a single player  $i \in N$ . The subscript  $i$  for  $A_i$ ,  $a_i$ ,  $\Theta_i$ ,  $\theta_i$ ,  $\pi_i^t$ , and  $R_{\text{US},i}^T$  is not essential in this section, but we put it for notational consistency.

Since a special case of our setting (the case of  $|\Theta_i| = 1$ ) is swap regret minimization, an  $\Omega(\sqrt{T|A_i|\log|A_i|})$  lower bound given by Ito [2020] can be applied. Combining it with our lower bound, we can argue that the leading term of our upper bound is tight up to a multiplicative constant.

The main theorem of this section is stated as follows.

**Theorem 4.1.** *Assume  $|A_i| = 2$ ,  $|\Theta_i| = 2^{B+1}$  for some  $B \in \mathbb{Z}$ , and  $T$  is a multiple of  $B = \log_2 |\Theta_i| - 1$ . If  $T \geq 2^{-49} |\Theta_i|^2 \log_2 |\Theta_i|$ , then for any randomized algorithm, there exists an oblivious adversary for which the untruthful swap regret of the algorithm is at least  $2^{-28} \sqrt{T \log_2 |\Theta_i|}$ .*

The assumption of  $T = \Omega(|\Theta_i|^2 \log |\Theta_i|)$  is essentially used in the proofs. The other assumptions on  $|A_i|$ ,  $|\Theta_i|$ , and  $T$  are just for simplicity of the proofs and only affect the multiplicative constants.

We want to prove that for any randomized algorithm, there exists a deterministic adversary for which the algorithm's untruthful swap regret is lower bounded. From Yao's minimax principle, it is sufficient to show there exists a randomized adversary for which any deterministic algorithm's expected untruthful swap regret is lower bounded. In the following, we construct such a randomized adversary.

Let  $A_i = \{\alpha_0, \alpha_1\}$  be the set of actions and  $\rho_i \in \Delta(\Theta_i)$  the uniform distribution over  $\Theta_i$ . The set of time rounds  $[T]$  is partitioned into  $B$  blocks of equal length  $L = T/B$ , where  $B = \log_2 |\Theta_i| - 1$ . For each  $b \in [B]$ , we denote the  $b$ th block by  $\mathcal{T}_b = \{t \in [T] \mid (b-1)L < t \leq bL\}$ .

The set of types  $\Theta_i$  is partitioned into  $\Theta'_i$  of size  $2^B$  and  $\Theta''_i$  of size  $2^B$ . The rewards for types  $\Theta'_i$  are determined by a bijection  $\zeta: \Theta'_i \rightarrow \{0, 1\}^B$  that associates each type in  $\theta'_i \in \Theta'_i$  with binary sequence  $\zeta(\theta'_i)$  of length  $B$ . A bijection  $\zeta$  is latently generated from the uniform distribution over all such bijections and unknown to the algorithm in the beginning. For each time round  $t \in \mathcal{T}_b$  in the  $b$ th block, the reward  $u_i^t(\theta_i, \alpha_0)$  is defined by  $\zeta$  such that  $u_i^t(\theta'_i, \alpha_0) = \zeta(\theta'_i)(b)$  for each  $\theta'_i \in \Theta'_i$ . On the other hand, the rewards for types  $\Theta''_i$  are determined by a random map  $\xi: \Theta''_i \rightarrow \{0, 1\}^T$ . For each  $\theta''_i \in \Theta''_i$  and  $t \in [T]$ , the reward  $u_i^t(\theta''_i, \alpha_0) = \xi(\theta''_i)(t)$  independently follows the uniform distribution over  $\{0, 1\}$ . For both  $\Theta'_i$  and  $\Theta''_i$ , the reward for  $\alpha_1$  is defined by flipping the reward for  $\alpha_0$ , i.e.,  $u_i^t(\theta_i, \alpha_1) = 1 - u_i^t(\theta_i, \alpha_0)$  for every type  $\theta_i \in \Theta_i$  and round  $t \in [T]$ . See Figure 2 for an illustration.

Once  $\zeta$  and  $\xi$  are fixed, the problem instance is determined. Then, since the algorithm is deterministic, the algorithm's decisions  $(\pi_i^t)_{t \in [T]}$  are also determined. Furthermore, since in each round  $t \in [T]$ , the algorithm decides  $\pi_i^t$  according to the rewards observed so far,  $\pi_i^t$  is deterministic if  $(\zeta(\theta'_i)(b))_{\theta'_i \in \Theta'_i, b \in [(t-1)/L]}$  and  $(\xi(\theta''_i)(t'))_{\theta''_i \in \Theta''_i, t' \in [t-1]}$  are fixed.

The untruthful swap regret  $R_{\text{US},i}^T$  depends on the problem instance parametrized by  $\zeta$  and  $\xi$ . We will develop an  $\Omega(\sqrt{T \log |\Theta_i|})$  lower bound on the expected value  $\mathbb{E}_{\zeta, \xi} [R_{\text{US},i}^T]$ . This lower bound implies that there exists some  $\zeta$  and  $\xi$  such that  $R_{\text{US},i}^T = \Omega(\sqrt{T \log |\Theta_i|})$ , which leads to the theorem.

#### 4.1 Analysis for randomly branching rewards

The goal of this subsection is to prove that the algorithm with low untruthful swap regret must make significantly different decisions for types  $\Theta'_i$  (Theorem 4.4). Toward this goal, we first show that the algorithm must select  $\alpha_0$  or  $\alpha_1$  for the type with all 0 or all 1 binary sequence, respectively (Theorem 4.2). We then show that the algorithm must make significantly different decisions for "adjacent" types (Theorem 4.3). By taking the expectation over  $\zeta$  and  $\xi$ , we prove Theorem 4.4.

First, we prove that if the untruthful swap regret is small, for the type  $\theta_i^0 \in \Theta'_i$  with  $\zeta(\theta_i^0) = 00 \cdots 00$ , the algorithm must choose  $\alpha_0$  for most of the rounds. Similarly, for the type  $\theta_i^1 \in \Theta'_i$  with  $\zeta(\theta_i^1) = 11 \cdots 11$ , the algorithm must choose  $\alpha_1$  for most of the rounds.

**Lemma 4.2.** *Fix  $\zeta$  and  $\xi$ . For  $\theta_i^0 \in \Theta'_i$  such that  $\zeta(\theta_i^0)(b) = 0$  for all  $b \in [B]$ , it must hold that  $\sum_{t=1}^T \pi_i^t(\theta_i^0; \alpha_0) \geq T - |\Theta_i| R_{\text{US},i}^T$ . For  $\theta_i^1 \in \Theta'_i$  such that  $\zeta(\theta_i^1)(b) = 1$  for all  $b \in [B]$ , it must hold that  $\sum_{t=1}^T \pi_i^t(\theta_i^1; \alpha_1) \geq T - |\Theta_i| R_{\text{US},i}^T$ .*

*Proof.* Let  $\theta_i^0 \in \Theta'_i$  be the type satisfying  $\zeta(\theta_i^0)(b) = 0$  for all  $b \in [B]$ . We consider  $\psi$  and  $\phi$  in the definition of untruthful swap regret such that  $\psi(\theta_i^0) = \theta_i^0$  and  $\phi(\theta_i^0, \alpha_0) = \phi(\theta_i^0, \alpha_1) = \alpha_0$ . Let  $\psi(\theta_i) = \theta_i$  and  $\phi(\theta_i, \alpha) = \alpha$  for all the other types  $\theta_i \in \Theta_i$  and any action  $\alpha \in A_i$ . We thus obtain

$$\begin{aligned} R_{\text{US},i}^T &= \max_{\psi: \Theta_i \rightarrow \Theta_i} \max_{\phi: \Theta_i \times A_i \rightarrow A_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{\alpha \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta, \phi(\theta_i, \alpha))] - \mathbb{E}_{\alpha \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, \alpha)] \right] \\ &\geq \frac{1}{|\Theta_i|} \sum_{t=1}^T \left\{ u_i^t(\theta_i^0, \alpha_0) - \mathbb{E}_{\alpha \sim \pi_i^t(\theta_i^0)} [u_i^t(\theta_i^0, \alpha)] \right\}. \end{aligned}$$

Since  $\zeta(\theta_i^0)(b) = 0$  for all  $b \in [B]$ , we have  $u_i^t(\theta_i^0, \alpha_0) = 1$  for all  $t \in [T]$ , which implies

$$R_{\text{US},i}^T \geq \frac{1}{|\Theta_i|} \left\{ T - \sum_{t=1}^T \pi_i^t(\theta_i^0; \alpha_0) \right\}.$$

Rearranging this inequality yields  $\sum_{t=1}^T \pi_i^t(\theta_i^0; \alpha_0) \geq T - |\Theta_i| R_{\text{US},i}^T$ . By applying the same argument to  $\theta_i^1 \in \Theta'_i$  such that  $\zeta(\theta_i^1)(b) = 1$  for all  $b \in [B]$ , we obtain  $\sum_{t=1}^T \pi_i^t(\theta_i^1; \alpha_1) \geq T - |\Theta_i| R_{\text{US},i}^T$ .  $\square$

The previous lemma shows that for the type with  $\zeta(\theta_i) = 00 \cdots 00$  or  $\zeta(\theta_i) = 11 \cdots 11$ , the algorithm must choose the optimal action for most of the rounds. For the other types, we cannot prove a similar claim since there exists a counterexample (see Section 1.2). However, we can show that the algorithm must make sufficiently different decisions for some pair of ‘‘adjacent’’ types.

Fix any  $\theta_i \in \Theta'_i$ . To define the adjacency for  $\theta_i$ , we define a sequence of types as follows (see Figure 4 for example). Let  $B_0$  and  $B_1$  be the numbers of blocks  $b \in [B]$  such that  $\zeta(\theta_i)(b) = 0$  and  $\zeta(\theta_i)(b) = 1$ , respectively. We define the sorted indices  $b_1 \leq b_2 \leq \cdots \leq b_{B_0}$  of blocks such that  $\zeta(\theta_i)(b_k) = 0$  for each  $k \in [B_0]$ . Similarly, we define the sorted indices  $b'_1 \leq b'_2 \leq \cdots \leq b'_{B_1}$  of blocks such that  $\zeta(\theta_i)(b'_l) = 1$  for each  $l \in [B_1]$ . For notational convenience, we define  $b_0 = b'_0 = 0$  and  $b_{B_0+1} = b'_{B_1+1} = B + 1$ .

For each  $k = 1, 2, \dots, B_0 + 1$ , we define the type  $\nu_\zeta(\theta_i, b_k) \in \Theta'_i$  such that  $\zeta(\nu_\zeta(\theta_i, b_k))$  is the binary sequence whose first  $b_{k-1}$  bits are identical to the prefix of  $\zeta(\theta_i)$  and the remaining bits are all 1, i.e.,  $\zeta(\nu_\zeta(\theta_i, b_k))(b) = \zeta(\theta_i)(b)$  for each  $b \leq b_{k-1}$  and  $\zeta(\nu_\zeta(\theta_i, b_k))(b) = 1$  for each  $b > b_{k-1}$ . Similarly, for each  $l = 1, 2, \dots, B_1 + 1$ , we define the type  $\nu_\zeta(\theta_i, b'_l) \in \Theta'_i$  such that  $\zeta(\nu_\zeta(\theta_i, b'_l))$  is the binary sequence whose first  $b'_{l-1}$  bits are identical to the prefix of  $\zeta(\theta_i)$  and the remaining bits are all 0.

In the following lemma, we show that the algorithm must make significantly different decisions for  $\theta_i$  and  $\nu_\zeta(\theta_i, b)$  in each block  $b \in [B]$  if  $R_{\text{US},i}^T$  is small.

	$\zeta(\cdot)$				
$\nu_\zeta(\theta_i, b_1)$	1	1	1	1	1
$\nu_\zeta(\theta_i, b_2)$	1	0	1	1	1
$\theta_i = \nu_\zeta(\theta_i, b_3) = \nu_\zeta(\theta_i, b'_4)$	1	0	1	1	0
$\nu_\zeta(\theta_i, b'_3)$	1	0	1	0	0
$\nu_\zeta(\theta_i, b'_2)$	1	0	0	0	0
$\nu_\zeta(\theta_i, b'_1)$	0	0	0	0	0
	$b'_1$	$b_1$	$b'_2$	$b'_3$	$b_2$

Figure 4: An example of the definition of  $\nu_\zeta$ . Given  $\theta_i \in \Theta'_i$ , we denote the block indices corresponding to 0 and 1 by  $b_1, b_2, \dots, b_{B_0}$  and  $b'_1, b'_2, \dots, b'_{B_1}$ , respectively. We define the series of types  $\nu_\zeta(\theta_i, b_1), \nu_\zeta(\theta_i, b_2), \dots, \nu_\zeta(\theta_i, b_{B_0}), \nu_\zeta(\theta_i, b_{B_0+1}) = \theta_i$  such that their corresponding binary sequences gradually change from  $11 \cdots 11$  to  $\zeta(\theta_i)$ . Similarly, we define the series of types  $\nu_\zeta(\theta_i, b'_1), \nu_\zeta(\theta_i, b'_2), \dots, \nu_\zeta(\theta_i, b'_{B_1}), \nu_\zeta(\theta_i, b'_{B_1+1}) = \theta_i$  such that their corresponding binary sequences gradually change from  $00 \cdots 00$  to  $\zeta(\theta_i)$ .

**Lemma 4.3.** Fix  $\zeta$  and  $\xi$ . Let  $\theta_i \in \Theta'_i$  be an arbitrary type. Then it holds that

$$\sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \{ \pi_i^t(\theta_i; \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) \} + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ \pi_i^t(\theta_i; \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b'_l); \alpha_1) \} \geq \frac{T}{2} - 2|\Theta_i|R_{\text{US},i}^T.$$

*Proof.* We consider  $\psi$  in the definition of untruthful swap regret such that  $\psi(\nu(\theta_i, b_{k+1})) = \nu(\theta_i, b_k)$  for each  $k \in [B_0]$ . Let  $\psi(\theta_i) = \theta_i$  for all the other types  $\theta_i \in \Theta_i$ . We use the identity map for actions, i.e.,  $\phi(\theta_i, \alpha) = \alpha$  for all types  $\theta_i \in \Theta_i$  and actions  $\alpha \in A_i$ . Then we obtain a lower bound on the untruthful swap regret as

$$R_{\text{US},i}^T \geq \frac{1}{|\Theta_i|} \sum_{k=1}^{B_0} \sum_{t=1}^T \left\{ \mathbb{E}_{\alpha \sim \pi_i^t(\nu_\zeta(\theta_i, b_k))} [u_i^t(\nu_\zeta(\theta_i, b_{k+1}), \alpha)] - \mathbb{E}_{\alpha \sim \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}))} [u_i^t(\nu_\zeta(\theta_i, b_{k+1}), \alpha)] \right\}. \quad (7)$$

From the definition of  $\nu_\zeta(\theta_i, b_{k+1})$ , the binary sequence  $\zeta(\nu_\zeta(\theta_i, b_{k+1}))$  is identical to  $\zeta(\theta_i)$  for the first  $b_k$  blocks and all 1 for the remaining  $B - b_k$  blocks. This implies that  $u_i^t(\nu_\zeta(\theta_i, b_{k+1}), \alpha_0) = 1$  for all blocks with indices  $b_1, b_2, \dots, b_k$ , and  $u_i^t(\nu_\zeta(\theta_i, b_{k+1}), \alpha_1) = 1$  for all blocks with indices  $b_{k+1}, b_{k+2}, \dots, b_{B_0}$  and  $b'_1, b'_2, \dots, b'_{B_1}$ . Each term on the right-hand side of (7) for  $k \in [B_0]$  can be decomposed as

$$\begin{aligned} & \sum_{t=1}^T \left\{ \mathbb{E}_{\alpha \sim \pi_i^t(\nu_\zeta(\theta_i, b_k))} [u_i^t(\nu_\zeta(\theta_i, b_{k+1}), \alpha)] - \mathbb{E}_{\alpha \sim \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}))} [u_i^t(\nu_\zeta(\theta_i, b_{k+1}), \alpha)] \right\} \\ &= \underbrace{\sum_{k'=1}^k \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_0) \}}_{(A)} \end{aligned}$$

$$\begin{aligned}
& + \underbrace{\sum_{k'=k+1}^{B_0} \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_1) \}}_{(B)} \\
& + \underbrace{\sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_1) \}}_{(C)}.
\end{aligned}$$

We take the summation of these terms over all  $k = 1, \dots, B_0$ . The sum of (A) for all  $k = 1, \dots, B_0$  is

$$\begin{aligned}
& \sum_{k=1}^{B_0} \sum_{k'=1}^k \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_0) \} \\
& = \sum_{k'=1}^{B_0} \sum_{k=k'}^{B_0} \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_0) \} \\
& = \sum_{k'=1}^{B_0} \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_{k'}); \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b_{B_0+1}); \alpha_0) \},
\end{aligned}$$

where the first equality is obtained by changing the order of summations and the second equality is obtained by the telescoping sum. Similarly, the sum of (B) for all  $k = 1, \dots, B_0$  is

$$\begin{aligned}
& \sum_{k=1}^{B_0} \sum_{k'=k+1}^{B_0} \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_1) \} \\
& = \sum_{k'=1}^{B_0} \sum_{k=1}^{k'-1} \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_1) \} \\
& = \sum_{k'=1}^{B_0} \sum_{t \in \mathcal{T}_{b_{k'}}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_{k'}); \alpha_1) \}.
\end{aligned}$$

By using the telescoping sum, the sum of (C) for all  $k = 1, \dots, B_0$  is

$$\begin{aligned}
& \sum_{k=1}^{B_0} \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_{k+1}); \alpha_1) \} \\
& = \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_{B_0+1}); \alpha_1) \}.
\end{aligned}$$

By substituting these terms into (7), we obtain

$$|\Theta_i| R_{\text{US},i}^T \geq \sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) - \pi_i^t(\theta_i; \alpha_0) + \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) \}$$

$$+ \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1) - \pi_i^t(\theta_i; \alpha_1) \},$$

where we used  $\nu_\zeta(\theta_i, b_{B_0+1}) = \theta_i$ . Recall  $\zeta(\nu_\zeta(\theta_i, b_1))(b) = 1$  for all  $b \in [B]$ . From Theorem 4.2, we have

$$\sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1) + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1) \geq T - |\Theta_i| R_{\text{US},i}^T.$$

By summing these two inequalities and rearranging, we obtain

$$\sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \{ \pi_i^t(\theta_i; \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) + \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) \} + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \pi_i^t(\theta_i; \alpha_1) \geq T - 2|\Theta_i| R_{\text{US},i}^T,$$

where we removed  $\sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1)$  and  $\sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \pi_i^t(\nu_\zeta(\theta_i, b_1); \alpha_1)$  from both sides.

Finally, by using  $\pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) + \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_1) = 1$  for each  $t \in [T]$ , we obtain

$$\sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \{ \pi_i^t(\theta_i; \alpha_0) - 2\pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) + 1 \} + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \pi_i^t(\theta_i; \alpha_1) \geq T - 2|\Theta_i| R_{\text{US},i}^T.$$

By using the symmetric analysis, we obtain

$$\sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \pi_i^t(\theta_i; \alpha_0) + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ \pi_i^t(\theta_i; \alpha_1) - 2\pi_i^t(\nu_\zeta(\theta_i, b'_l); \alpha_1) + 1 \} \geq T - 2|\Theta_i| R_{\text{US},i}^T.$$

By summing these two inequalities, we obtain

$$\begin{aligned} \sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \{ 2\pi_i^t(\theta_i; \alpha_0) - 2\pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) + 1 \} + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ 2\pi_i^t(\theta_i; \alpha_1) - 2\pi_i^t(\nu_\zeta(\theta_i, b'_l); \alpha_1) + 1 \} \\ \geq 2T - 4|\Theta_i| R_{\text{US},i}^T. \end{aligned}$$

Since  $\sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} 1 + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} 1 = \sum_{t \in [T]} 1 = T$ , by dividing both sides by 2, we complete the proof.  $\square$

The previous lemma claims that compared to the decisions for the adjacent type  $\nu_\zeta(\theta_i, b)$ , the algorithm must choose the optimal action for type  $\theta_i$  more frequently in each block  $b \in [B]$ . The next lemma compares the algorithm's decisions in block  $b \in [B]$  for types whose binary sequences first branch in block  $b$ , not only adjacent types.

For each  $b \in [B]$  and a prefix  $p \in \{0, 1\}^{b-1}$ , let  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  be the sets of types whose corresponding binary sequences have the prefix  $p$  and the next bit is 0 or 1, respectively. Formally, we define  $\mathcal{I}_{p,0} = \{ \theta_i \in \Theta'_i \mid \zeta(\theta_i)(b') = p(b') \ (\forall b' < b), \zeta(\theta_i)(b) = 0 \}$  and  $\mathcal{I}_{p,1} = \{ \theta_i \in \Theta'_i \mid \zeta(\theta_i)(b') = p(b') \ (\forall b' < b), \zeta(\theta_i)(b) = 1 \}$ . Note that  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  are random variables determined by  $\zeta_b$ . Let  $\mathcal{I}_p = \mathcal{I}_{p,0} \cup \mathcal{I}_{p,1}$  be their union, which is determined by  $\zeta_{b-1}$ .

In block  $b$ , the optimal action is  $\alpha_0$  for types in  $\mathcal{I}_{p,0}$  and  $\alpha_1$  for types in  $\mathcal{I}_{p,1}$ . In the following lemma, we show that the average decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  are sufficiently different with some constant probability.

**Lemma 4.4.** *Assume*

$$\mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] < 2^{-28} \sqrt{T \log_2 |\Theta_i|}.$$

*Then it holds that*

$$\Pr_{\substack{b \sim [B], \zeta_b, \xi_{b-1}, \\ p \sim \{0, 1\}^{b-1}}} \left( \mathbb{E}_{\xi_b} \left[ \mathbb{E}_{\theta_i \sim \mathcal{I}_{p, 0}} \left[ \sum_{t \in \mathcal{T}_b} \pi_i^t(\theta_i; \alpha_0) \right] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p, 1}} \left[ \sum_{t \in \mathcal{T}_b} \pi_i^t(\theta_i; \alpha_0) \right] \middle| \zeta_b, \xi_{b-1} \right] \geq \frac{L}{8} \right) \geq \frac{1}{8}.$$

*Proof.* Since Theorem 4.3 holds for any  $\theta_i \in \Theta'_i$ , it holds in expectation for any distribution of  $\theta_i$ . In this proof, assuming that  $\theta_i$  follows the uniform distribution over  $\Theta'_i$ , we take the expectation of Theorem 4.3 with respect to  $\zeta$ ,  $\xi$ , and  $\theta_i$ .

The expected value of the right-hand side of Theorem 4.3 is bounded as

$$\mathbb{E}_{\zeta, \xi, \theta_i} \left[ \frac{T}{2} - 2|\Theta_i| R_{\text{US}, i}^T \right] = \frac{T}{2} - 2|\Theta_i| \mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] \geq \frac{T}{4}, \quad (8)$$

where the inequality is due to the assumptions  $\mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] < 2^{-28} \sqrt{T \log_2 |\Theta_i|}$  and  $T \geq 2^{-49} |\Theta_i|^2 \log_2 |\Theta_i|$ .

The expected value of the left-hand side of Theorem 4.3 can be decomposed for each block as

$$\begin{aligned} & \mathbb{E}_{\zeta, \xi, \theta_i} \left[ \sum_{k=1}^{B_0} \sum_{t \in \mathcal{T}_{b_k}} \{ \pi_i^t(\theta_i; \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b_k); \alpha_0) \} + \sum_{l=1}^{B_1} \sum_{t \in \mathcal{T}_{b'_l}} \{ \pi_i^t(\theta_i; \alpha_1) - \pi_i^t(\nu_\zeta(\theta_i, b'_l); \alpha_1) \} \right] \\ &= \sum_{b=1}^B \left\{ \Pr_{\zeta, \xi, \theta_i} (\zeta(\theta_i)(b) = 0) \mathbb{E}_{\zeta, \xi, \theta_i} \left[ \sum_{t \in \mathcal{T}_b} \{ \pi_i^t(\theta_i; \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b); \alpha_0) \} \middle| \zeta(\theta_i)(b) = 0 \right] \right. \\ & \quad \left. + \Pr_{\zeta, \xi, \theta_i} (\zeta(\theta_i)(b) = 1) \mathbb{E}_{\zeta, \xi, \theta_i} \left[ \sum_{t \in \mathcal{T}_b} \{ \pi_i^t(\nu_\zeta(\theta_i, b); \alpha_0) - \pi_i^t(\theta_i; \alpha_0) \} \middle| \zeta(\theta_i)(b) = 1 \right] \right\}, \quad (9) \end{aligned}$$

where we used  $\pi_i^t(\theta_i; \alpha_0) + \pi_i^t(\theta_i; \alpha_1) = 1$  and  $\pi_i^t(\nu_\zeta(\theta_i, b); \alpha_0) + \pi_i^t(\nu_\zeta(\theta_i, b); \alpha_1) = 1$  for each  $t \in [T]$ . We take the expectation separately for each  $b \in [B]$ .

First, we take the expectation conditioned on observations from the first block to the  $b$ th block. Let  $\zeta_b$  and  $\xi_b$  be the part of  $\zeta$  and  $\xi$  that is revealed until the end of block  $b \in [B]$ . Formally, we define  $\zeta_b = (\zeta(\theta_i)(b'))_{\theta_i \in \Theta'_i, b' \in [b]}$  and  $\xi_b = (\xi(\theta_i)(t))_{\theta_i \in \Theta'_i, t \in [bL]}$ . If we condition on  $\zeta_b$  and  $\xi_b$ , the algorithm's outputs  $(\pi_i^t)_{t \in [bL]}$  until the end of block  $b$  are deterministic. Therefore, conditioned on  $\zeta_b$ ,  $\xi_b$ , and  $\theta_i$ , only  $\nu_\zeta(\theta_i, b)$  is stochastic in the term of which we take the expectation in (9). Recall that  $\nu_\zeta(\theta_i, b)$  is defined to be the type with the corresponding binary sequence equal to  $\zeta(\theta_i)$  before block  $b - 1$  and all 1 (or 0) after block  $b$  if  $\zeta(\theta_i)(b)$  is 0 (or 1, respectively). Let  $p \in \{0, 1\}^{b-1}$  be a random variable that represents the prefix of  $\zeta(\theta_i)$ . Since  $\zeta$  revealed after block  $b$  is uniformly distributed, if the  $b$ th bit  $\zeta(\theta_i)(b)$  is 0, then  $\nu_\zeta(\theta_i, b)$  is distributed uniformly over  $\mathcal{I}_{p, 1}$ . We thus obtain

$$\mathbb{E}_{\zeta, \xi} \left[ \sum_{t \in \mathcal{T}_b} \{ \pi_i^t(\theta_i; \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b); \alpha_0) \} \middle| \zeta_b, \xi_b \right] = \sum_{t \in \mathcal{T}_b} \left\{ \pi_i^t(\theta_i; \alpha_0) - \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p, 1}} [\pi_i^t(\theta'_i; \alpha_0)] \right\}. \quad (10)$$

Next, we take the expectation over  $\theta_i$  conditioned on  $\zeta_b$  and  $\xi_b$ . Since  $\theta_i$  is uniformly distributed over  $\Theta'_i$ , the prefix  $p$  follows the uniform distribution over  $\{0, 1\}^{b-1}$ . Then we obtain

$$\begin{aligned}
& \mathbb{E}_{\theta_i} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \pi_i^t(\theta_i; \alpha_0) - \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta'_i; \alpha_0)] \right\} \middle| \zeta(\theta_i)(b) = 0, \zeta_b, \xi_b \right] \\
&= \mathbb{E}_{p \sim \{0,1\}^{b-1}} \left[ \mathbb{E}_{\theta_i} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \pi_i^t(\theta_i; \alpha_0) - \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta'_i; \alpha_0)] \right\} \middle| \zeta(\theta_i)(b) = 0, p \text{ is a prefix of } \zeta(\theta_i), \zeta_b, \xi_b \right] \right] \\
&= \mathbb{E}_{p \sim \{0,1\}^{b-1}} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right\} \right]. \tag{11}
\end{aligned}$$

Finally, we take the expectation over  $\zeta_b$  and  $\xi_b$ , which determine  $\pi_i^t$  for each  $t \in [bL]$  and the partition  $(\mathcal{I}_{p,0}, \mathcal{I}_{p,1})$  for each  $p \in \{0, 1\}^{b-1}$ . By combining it with (10) and (11), we obtain

$$\begin{aligned}
& \mathbb{E}_{\zeta, \xi, \theta_i} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \pi_i^t(\theta_i; \alpha_0) - \pi_i^t(\nu_\zeta(\theta_i, b); \alpha_0) \right\} \middle| \zeta(\theta_i)(b) = 0 \right] \\
&= \mathbb{E}_{\zeta, \xi, p \sim \{0,1\}^{b-1}} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right\} \right].
\end{aligned}$$

By applying the same argument to the case of  $\zeta(\theta_i)(b) = 1$ , we obtain

$$\begin{aligned}
& \mathbb{E}_{\zeta, \xi, \theta_i} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \pi_i^t(\nu_\zeta(\theta_i, b); \alpha_0) - \pi_i^t(\theta_i; \alpha_0) \right\} \middle| \zeta(\theta_i)(b) = 1 \right] \\
&= \mathbb{E}_{\zeta, \xi, p \sim \{0,1\}^{b-1}} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right\} \right].
\end{aligned}$$

By plugging these expected values into (9) and bounding it by the right-hand side (8), we obtain

$$\sum_{b \in [B]} \mathbb{E}_{\zeta, \xi, p \sim \{0,1\}^{b-1}} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right\} \right] \geq \frac{T}{4}.$$

By dividing both sides by  $B$  and taking the conditional expectation, we obtain

$$\mathbb{E}_{b \sim [B], \zeta_b, \xi_b, p \sim \{0,1\}^{b-1}} \left[ \mathbb{E}_{\zeta, \xi} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right\} \middle| \zeta_b, \xi_b \right] \right] \geq \frac{L}{4}.$$

Since  $\mathcal{I}_{p,0}$ ,  $\mathcal{I}_{p,1}$ , and  $(\pi_i^t)_{t \in \mathcal{T}_b}$  are determined only by  $\zeta_b$  and  $\xi_b$ , we can replace the expectation with respect to  $\zeta$  and  $\xi$  by the expectation with respect to  $\xi_b$ . By applying Theorem 4.5 proved below, we complete the proof.  $\square$

In this proof, we used the following lemma.

**Lemma 4.5.** Fix  $x, y \in \mathbb{R}$  such that  $0 < y < x$ . Let  $X$  be a random variable whose value is at most  $x$ . If  $\mathbb{E}[X] \geq y$ , then  $\Pr(X \geq y/2) \geq \frac{y}{2x}$ .

*Proof.* By Markov's inequality to  $x - X$ , for any  $\epsilon > 0$ , we obtain

$$\Pr\left(X \leq \frac{y}{2} + \epsilon\right) = \Pr\left(x - X \geq x - \frac{y}{2} - \epsilon\right) \leq \frac{\mathbb{E}[x - X]}{x - \frac{y}{2} - \epsilon} \leq \frac{x - y}{x - \frac{y}{2} - \epsilon} = 1 - \frac{\frac{y}{2} - \epsilon}{x - \frac{y}{2} - \epsilon}$$

By setting sufficiently small  $\epsilon \in (0, \frac{y^2}{4x-2y})$ , we obtain  $\Pr\left(X \leq \frac{y}{2} + \epsilon\right) \leq 1 - \frac{y}{2x}$ . By considering the complementary event, we obtain  $\Pr(X \geq y/2) \geq \Pr(X > y/2 + \epsilon) \geq \frac{y}{2x}$ .  $\square$

## 4.2 Assignment of types

The goal of this subsection is to define an assignment  $\psi'_{\zeta, \xi}: \Theta''_i \rightarrow \Theta'_i$  such that against the reward sequence for each  $\theta''_i \in \Theta''_i$ , the decision for  $\psi'_{\zeta, \xi}(\theta''_i) \in \Theta'_i$  achieves significantly better than  $T/2$ . In the previous subsection, we showed that if the expected untruthful swap regret is small, the algorithm must take decisions different for  $\mathcal{I}_{p,0} \subseteq \Theta'_i$  and  $\mathcal{I}_{p,1} \subseteq \Theta'_i$  with some constant probability (Theorem 4.4). We consider assigning these decisions to the rewards for each  $\theta''_i \in \Theta''_i$ . Since these rewards are determined by the independent random variables  $\xi$ , the total reward obtained by any decision is  $L/2$  in expectation for every block. Intuitively, if the decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  are sufficiently different, then the better one fits the rewards for  $\theta''_i \in \Theta''_i$ , thus achieving the expected reward  $L/2 + \Omega(\sqrt{L})$  due to anti-concentration.

A simple but failed approach is to decide  $\psi'_{\zeta, \xi}(\theta''_i)$  independently for each  $\theta''_i \in \Theta''_i$ . For each  $\zeta$  and  $\xi$ , this approach iteratively selects  $\mathcal{I}_{p,0}$  or  $\mathcal{I}_{p,1}$  whose average decision achieves better for the rewards of  $\theta''_i$  as follows. First, let  $p$  be the empty sequence. For each  $b = 1, 2, \dots, B$ , we assign  $\theta''_i$  to 0 if the average decision for  $\mathcal{I}_{p,0}$  is better than that for  $\mathcal{I}_{p,1}$  against the reward sequence of  $\theta''_i$ , i.e.,

$$\mathbb{E}_{\theta''_i \sim \mathcal{I}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_t^t(\theta''_i)} [u_i^t(\theta''_i, \alpha)] \right] \geq \mathbb{E}_{\theta''_i \sim \mathcal{I}_{p,1}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_t^t(\theta''_i)} [u_i^t(\theta''_i, \alpha)] \right]$$

and 1 otherwise, and we then add the selected 0 or 1 to the end of  $p$ . Finally, we set  $\psi'_{\zeta, \xi}(\theta''_i) = \zeta^{-1}(p)$ , where  $\zeta^{-1}: \{0, 1\}^B \rightarrow \Theta'_i$  is the inverse map of the bijection  $\zeta$ . This assignment process is done for each  $\theta''_i \in \Theta''_i$  independently. This approach has the following two issues.

- Since the assignment is independent for each  $\theta''_i \in \Theta''_i$ , the assigned types  $(\psi'_{\zeta, \xi}(\theta''_i))_{\theta''_i \in \Theta''_i}$  might be unbalanced, that is, for some  $\zeta_{b-1}$ ,  $\xi_{b-1}$ , and  $p \in \{0, 1\}^{b-1}$ , there is no  $\theta''_i \in \Theta''_i$  that is assigned to  $\mathcal{I}_{p,0}$  or  $\mathcal{I}_{p,1}$  conditioned on  $\zeta_{b-1}$  and  $\xi_{b-1}$ . If the average decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  are sufficiently different for all  $p \in \{0, 1\}^{b-1}$  with any  $b \in [B]$ , then this assignment is enough. However, Theorem 4.4 guarantees this condition only with a constant fraction of  $p$  and  $b$ . If the algorithm is manipulatively designed, this condition may fail to hold with a significant probability for each  $\zeta_{b-1}$ ,  $\xi_{b-1}$ , and  $p$  that  $\theta''_i \in \Theta''_i$  is assigned to.
- The assignment for each block is decided according to the average decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$ . However, to obtain a lower bound on the untruthful swap regret, we need to prove a lower bound for the decision for  $\psi_{\zeta, \xi}(\theta''_i)$ , not the average decision for  $\mathcal{I}_{p,0}$  or  $\mathcal{I}_{p,1}$ . The selection of  $\psi_{\zeta, \xi}(\theta''_i)$  from  $\mathcal{I}_{p,0}$  or  $\mathcal{I}_{p,1}$  depends on the assignment in the future blocks  $b, b+1, \dots, B$ . If the algorithm is manipulatively designed, the distribution of  $\psi_{\zeta, \xi}(\theta''_i)$  can be far from the uniform distribution over  $\mathcal{I}_{p,0}$  or  $\mathcal{I}_{p,1}$ .

To overcome these two obstacles, we use another definition of  $\psi'_{\zeta,\xi}$ . The assignment process for each  $\zeta$  and  $\xi$  proceeds as follows. Let  $\mathcal{J} = \Theta''_i$ , which is supposed to be  $\mathcal{J}_p$  with the empty sequence  $p$ . For each  $b = 1, 2, \dots, B$  and each  $p \in \{0, 1\}^{b-1}$ , we partition  $\mathcal{J}_p$  into two subsets  $\mathcal{J}_{p,0}$  and  $\mathcal{J}_{p,1}$  of equal size. To decide this partition, we add an arbitrary index for each type in  $\mathcal{J}_p$  as  $\mathcal{J}_p = \{j_1, j_2, \dots, j_{2^{B-b+1}}\}$ . For each  $k \in [2^{B-b}]$ , we add one of  $\{j_{2k-1}, j_{2k}\}$  into  $\mathcal{J}_{p,0}$  and the other into  $\mathcal{J}_{p,1}$ . Conditioned on  $\zeta_b$  and  $\xi_{b-1}$ , we consider the distribution of

$$\mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_t^i(\theta'_i)} [u_i^t(j_{2k-1}, \alpha) - u_i^t(j_{2k}, \alpha)] \right] - \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,1}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_t^i(\theta'_i)} [u_i^t(j_{2k-1}, \alpha) - u_i^t(j_{2k}, \alpha)] \right], \quad (12)$$

whose randomness comes from the randomness of  $\xi_b$ . This value compares the average decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  in terms of the performance difference for the rewards of  $j_{2k-1}$  and  $j_{2k}$ . For the upper 50% of this distribution, we set  $j_{2k-1} \in \mathcal{J}_{p,0}$  and  $j_{2k} \in \mathcal{J}_{p,1}$ . For the lower 50% of this distribution, we set  $j_{2k-1} \in \mathcal{J}_{p,1}$  and  $j_{2k} \in \mathcal{J}_{p,0}$ . Since this is a discrete distribution, there might be arbitrariness in how to divide the mass on the median, but any division is allowed. In the end of this process, for each  $p \in \{0, 1\}^B$ , the assigned type  $\mathcal{J}_p \subseteq \Theta''_i$  is a singleton. Recall that  $\mathcal{I}_p$  is also a singleton for each  $p \in \{0, 1\}^B$ . For each  $p \in \{0, 1\}^B$ , we define  $\psi'_{\zeta,\xi}(\theta''_i) = \theta'_i$ , where  $\theta'_i \in \mathcal{I}_p$  and  $\theta''_i \in \mathcal{J}_p$  are the unique element of  $\mathcal{I}_p$  and  $\mathcal{J}_p$ , respectively.

This assignment addresses the two issues described above as follows.

- The first issue is caused when  $(\psi'_{\zeta,\xi}(\theta''_i))_{\theta''_i \in \Theta''_i}$  is unbalanced over  $\Theta'_i$ . To avoid this issue, we consider all types in  $\Theta''_i$  at the same time and partition them into equally sized subsets in each block. This guarantees that  $\psi'_{\zeta,\xi}: \Theta''_i \rightarrow \Theta'_i$  is always a bijection. Therefore, for any  $b \in [B]$ ,  $p \in \{0, 1\}^{b-1}$ ,  $\zeta_{b-1}$ , and  $\xi_{b-1}$  such that the average decisions are sufficiently different for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$ , there is some  $\theta''_i$  assigned to it.
- The second issue is addressed by assigning each type in  $\mathcal{J}_p$  to  $\mathcal{J}_{p,0}$  and  $\mathcal{J}_{p,1}$  with equal probability. Since this partition applies recursively,  $\psi'_{\zeta,\xi}(\theta''_i)$  follows the uniform distribution over  $\mathcal{I}_p$ . This is formally stated in the following lemma.

**Lemma 4.6.** *Let  $b \in [B]$  be an arbitrary block index. Let  $\zeta_b$  and  $\xi_b$  be any realization of  $\zeta$  and  $\xi$  for the first  $b$  blocks. Let  $p \in \{0, 1\}^b$  be an arbitrary bit sequence of length  $b$ . For any pair of  $\theta'_i \in \mathcal{I}_p$  and  $\theta''_i \in \mathcal{J}_p$ , it holds that*

$$\Pr_{\zeta,\xi}(\psi'_{\zeta,\xi}(\theta''_i) = \theta'_i \mid \zeta_b, \xi_b) = \frac{1}{2^{B-b}}.$$

*Proof.* Fix any  $b \in [B]$ ,  $\zeta_b$ ,  $\xi_b$ ,  $p \in \{0, 1\}^b$ ,  $\theta'_i \in \mathcal{I}_p$ , and  $\theta''_i \in \mathcal{J}_p$ . Let  $\mathcal{F}_{b'}$  be the event that  $\theta'_i$  and  $\psi'_{\zeta,\xi}(\theta''_i)$  are in the same partition in the  $b'$ th block. Formally,  $\mathcal{F}_{b'}$  is the event that  $\theta'_i \in \mathcal{I}_{p'}$  and  $\theta''_i \in \mathcal{J}_{p'}$  hold for some  $p' \in \{0, 1\}^{b'}$ . By induction, we prove

$$\Pr_{\zeta,\xi}(\mathcal{F}_{b'} \mid \zeta_b, \xi_b) = \frac{1}{2^{b'-b}} \quad (13)$$

for each  $b' \in [B]$  such that  $b \leq b' \leq B$ .

For the case of  $b' = b$ , since  $\theta'_i \in \mathcal{I}_p$  and  $\theta''_i \in \mathcal{J}_p$ , we obtain

$$\Pr_{\zeta,\xi}(\mathcal{F}_b \mid \zeta_b, \xi_b) = 1,$$

which proves the base case.

Fix any  $b' \in [B]$  such that  $b \leq b' \leq B - 1$ . Assume (13) holds for  $b'$ . For any  $p' \in \{0, 1\}^{b'}$  that has  $p$  as a prefix, we have

$$\begin{aligned}
& \Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p', 0}, \theta''_i \in \mathcal{J}_{p', 0} \mid \zeta_b, \xi_b) \\
&= \Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \mid \zeta_b, \xi_b) \\
&\times \Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p', 0} \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'}) \Pr_{\zeta, \xi} (\theta''_i \in \mathcal{J}_{p', 0} \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p', 0}, \theta''_i \in \mathcal{J}_{p'}). \tag{14}
\end{aligned}$$

Conditioned on  $\zeta_{b'}$  and  $\xi_{b'}$ , whether  $\theta'_i \in \mathcal{I}_{p', 0}$  or  $\theta'_i \in \mathcal{I}_{p', 1}$  holds is determined only by the randomness of the  $(b' + 1)$ th block of  $\zeta$ . The second factor of (14) is

$$\begin{aligned}
& \Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p', 0} \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'}) \\
&= \mathbb{E}_{\zeta_{b'}, \xi_{b'}} \left[ \Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p', 0} \mid \zeta_{b'}, \xi_{b'}, \theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'}) \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \right] \\
&= \mathbb{E}_{\zeta_{b'}, \xi_{b'}} \left[ \frac{1}{2} \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \right] \\
&= \frac{1}{2}. \tag{15}
\end{aligned}$$

Conditioned on  $\zeta_{b'+1}$  and  $\xi_{b'}$ , whether  $\theta''_i \in \mathcal{J}_{p', 0}$  or  $\theta''_i \in \mathcal{J}_{p', 1}$  is determined only by the randomness of  $\xi_{b'+1}$ . Since we decide it according to whether the realized  $\xi_{b'+1}$  is in the upper 50% or lower 50% of the distribution of some random variable, the third factor of (14) is

$$\begin{aligned}
& \Pr_{\zeta, \xi} (\theta''_i \in \mathcal{J}_{p', 0} \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p', 0}, \theta''_i \in \mathcal{J}_{p'}) \\
&= \mathbb{E}_{\zeta_{b'+1}, \xi_{b'}} \left[ \Pr_{\zeta, \xi} (\theta''_i \in \mathcal{J}_{p', 0} \mid \zeta_{b'+1}, \xi_{b'}, \theta'_i \in \mathcal{I}_{p', 0}, \theta''_i \in \mathcal{J}_{p'}) \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p', 0}, \theta''_i \in \mathcal{J}_{p'} \right] \\
&= \mathbb{E}_{\zeta_{b'+1}, \xi_{b'}} \left[ \frac{1}{2} \mid \zeta_b, \xi_b, \theta'_i \in \mathcal{I}_{p', 0}, \theta''_i \in \mathcal{J}_{p'} \right] \\
&= \frac{1}{2}. \tag{16}
\end{aligned}$$

By plugging (15) and (16) into (14), we obtain

$$\Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p', 0}, \theta''_i \in \mathcal{J}_{p', 0} \mid \zeta_b, \xi_b) = \frac{1}{4} \Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \mid \zeta_b, \xi_b).$$

In the same way, we can obtain

$$\Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p', 1}, \theta''_i \in \mathcal{J}_{p', 1} \mid \zeta_b, \xi_b) = \frac{1}{4} \Pr_{\zeta, \xi} (\theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \mid \zeta_b, \xi_b).$$

Then the probability of  $\mathcal{F}_{b'+1}$  is

$$\Pr_{\zeta, \xi} (\mathcal{F}_{b'+1} \mid \zeta_b, \xi_b)$$

$$\begin{aligned}
&= \sum_{p' \in \{0,1\}^{b'+1}} \Pr_{\zeta, \xi}(\theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \mid \zeta_b, \xi_b) \\
&= \sum_{p' \in \{0,1\}^{b'}} \left\{ \Pr_{\zeta, \xi}(\theta'_i \in \mathcal{I}_{p',0}, \theta''_i \in \mathcal{J}_{p',0} \mid \zeta_b, \xi_b) + \Pr_{\zeta, \xi}(\theta'_i \in \mathcal{I}_{p',1}, \theta''_i \in \mathcal{J}_{p',1} \mid \zeta_b, \xi_b) \right\} \\
&= \sum_{p' \in \{0,1\}^{b'}} \left\{ \frac{1}{4} \Pr_{\zeta, \xi}(\theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \mid \zeta_b, \xi_b) + \frac{1}{4} \Pr_{\zeta, \xi}(\theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \mid \zeta_b, \xi_b) \right\} \\
&= \frac{1}{2} \sum_{p' \in \{0,1\}^{b'}} \Pr_{\zeta, \xi}(\theta'_i \in \mathcal{I}_{p'}, \theta''_i \in \mathcal{J}_{p'} \mid \zeta_b, \xi_b) \\
&= \frac{1}{2} \Pr_{\zeta, \xi}(\mathcal{F}_{b'}) \mid \zeta_b, \xi_b) \\
&= \frac{1}{2} \cdot \frac{1}{2^{b'-b}} \tag{from the induction hypothesis} \\
&= \frac{1}{2^{b'-b+1}},
\end{aligned}$$

which proves the induction step.

By induction, (13) holds for each  $b \leq b' \leq B$ . By considering the case of  $b' = B$ , we obtain

$$\Pr_{\zeta, \xi}(\mathcal{F}_B \mid \zeta_b, \xi_b) = \frac{1}{2^{B-b}}.$$

From the definition of  $\psi'_{\zeta, \xi}$ , the event  $\mathcal{F}_B$  implies  $\psi'_{\zeta, \xi}(\theta''_i) = \theta'_i$ , which completes the proof.  $\square$

### 4.3 Analysis of the expected reward in each block

In Section 4.1, we showed that the average decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  are sufficiently different with a constant probability. In Section 4.2, we defined the assignment  $\psi'_{\zeta, \xi}: \Theta''_i \rightarrow \Theta'_i$  such that  $\mathcal{J}_{p,0}$  and  $\mathcal{J}_{p,1}$  correspond to  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  in each block, respectively. The next lemma claims that if the average decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  are sufficiently different, the average decisions for  $\mathcal{J}_{p,0}$  and  $\mathcal{J}_{p,1}$  achieve  $L/2 + \Omega(\sqrt{L})$  for the rewards of  $\mathcal{J}_{p,0}$  and  $\mathcal{J}_{p,1}$ , respectively.

The proof proceeds in the following steps. First, by introducing variables  $Y_{\theta''_i}^t$  and  $Z_{\theta''_i}^t$  for each  $\theta''_i \in \Theta''_i$  and  $t \in \mathcal{T}_b$ , we express the expected value achieved by the assignment as (18). The first term of (18) is equal to  $L/2$  as shown in (19). The remaining term of (18) can be written as a random walk term  $D_k^{t_{\max}}$  with a sign  $h_k(\xi_b)$  as shown in (21). To bound it by a term without  $h_k(\xi_b)$ , we define a more tractable  $h'_k(\xi_b)$  and obtain (22). We bound (22) by using Doob's optional stopping theorem. Combining all these inequalities, we prove the following lemma.

**Lemma 4.7.** *Let  $b \in [B]$  be an arbitrary block index. Let  $\zeta_b$  be any partial realization of  $\zeta$  revealed until the end of the  $b$ th block. Let  $\xi_{b-1}$  be any partial realization of  $\xi$  revealed until the end of the  $(b-1)$ th block. Let  $p \in \{0,1\}^{b-1}$  be an arbitrary bit sequence of length  $b-1$ . Let  $c \in [0,1]$  be an arbitrary non-negative constant. If*

$$\mathbb{E}_{\xi_b} \left[ \sum_{t \in \mathcal{T}_b} \left| \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right| \mid \zeta_b, \xi_{b-1} \right] \geq cL,$$

then

$$\begin{aligned} \mathbb{E}_{\xi_b} \left[ \frac{1}{2} \mathbb{E}_{\substack{\theta'_i \sim \mathcal{I}_{p,0} \\ \theta''_i \sim \mathcal{J}_{p,0}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] + \frac{1}{2} \mathbb{E}_{\substack{\theta'_i \sim \mathcal{I}_{p,1} \\ \theta''_i \sim \mathcal{J}_{p,1}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] \middle| \zeta_b, \xi_{b-1} \right] \\ \geq \frac{L}{2} + 2^{-10} e^{9/2} \sqrt{L}. \end{aligned} \quad (17)$$

*Proof.* Fix  $b \in [B]$  and  $p \in \{0, 1\}^{b-1}$ . We define random variables

$$X_0^t = \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta'_i; \alpha_0)] \quad \text{and} \quad X_1^t = \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta'_i; \alpha_0)]$$

that represent the average decisions for  $\mathcal{I}_{p,0}$  and  $\mathcal{I}_{p,1}$  for each  $t \in \mathcal{T}_b$ . Note that they are deterministic if  $\zeta_b$  and  $\xi_b$  are fixed. If we apply the average decision for  $\mathcal{I}_{p,0}$  to the rewards for each  $\theta''_i \in \mathcal{J}_p$ , the expected total reward in round  $t \in \mathcal{T}_b$  is

$$\begin{aligned} & \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,0}} \left[ \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] \\ &= \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta'_i; \alpha_0) u_i^t(\theta''_i, \alpha_0) + \pi_i^t(\theta'_i; \alpha_1) u_i^t(\theta''_i, \alpha_1)] \\ &= \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta'_i; \alpha_0)] u_i^t(\theta''_i, \alpha_0) + \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta'_i; \alpha_1)] u_i^t(\theta''_i, \alpha_1) \\ &= X_0^t \xi(\theta''_i)(t) + (1 - X_0^t)(1 - \xi(\theta''_i)(t)). \end{aligned}$$

Similarly, if we apply the average decision for  $\mathcal{I}_{p,1}$ , the expected total reward is

$$\mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,1}} \left[ \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] = X_1^t \xi(\theta''_i)(t) + (1 - X_1^t)(1 - \xi(\theta''_i)(t)).$$

For each  $\theta''_i \in \Theta''_i$  and  $t \in \mathcal{T}_b$ , we define random variables

$$Y_{\theta''_i}^t = \frac{X_0^t + X_1^t}{2} \xi(\theta''_i)(t) + \left( 1 - \frac{X_0^t + X_1^t}{2} \right) (1 - \xi(\theta''_i)(t))$$

and

$$Z_{\theta''_i}^t = \frac{X_0^t - X_1^t}{2} \xi(\theta''_i)(t) - \frac{X_0^t - X_1^t}{2} (1 - \xi(\theta''_i)(t)),$$

which are determined by  $\zeta_b$  and  $\xi_b$ . The total reward of the average decisions for  $\mathcal{I}_{p,0}$  in block  $b$  is

$$\begin{aligned} \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,0}} \left[ \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] &= \sum_{t \in \mathcal{T}_b} \{ X_0^t \xi(\theta''_i)(t) + (1 - X_0^t)(1 - \xi(\theta''_i)(t)) \} \\ &= \sum_{t \in \mathcal{T}_b} (Y_{\theta''_i}^t + Z_{\theta''_i}^t). \end{aligned}$$

Similarly, the total reward of the average decisions for  $\mathcal{I}_{p,1}$  is

$$\sum_{t \in \mathcal{T}_b} \mathbb{E}_{\theta'_i \sim \mathcal{I}_{p,1}} \left[ \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] = \sum_{t \in \mathcal{T}_b} \{ X_1^t \xi(\theta''_i)(t) + (1 - X_1^t)(1 - \xi(\theta''_i)(t)) \}$$

$$= \sum_{t \in \mathcal{T}_b} (Y_{\theta''_i}^t - Z_{\theta''_i}^t).$$

The left-hand side of the inequality (17) is

$$\begin{aligned} & \mathbb{E}_{\xi_b} \left[ \frac{1}{2} \mathbb{E}_{\substack{\theta'_i \sim \mathcal{I}_{p,0} \\ \theta''_i \sim \mathcal{J}_{p,0}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] + \frac{1}{2} \mathbb{E}_{\substack{\theta'_i \sim \mathcal{I}_{p,1} \\ \theta''_i \sim \mathcal{J}_{p,1}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] \middle| \zeta_b, \xi_{b-1} \right] \\ &= \mathbb{E}_{\xi_b} \left[ \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} (Y_{\theta''_i}^t + Z_{\theta''_i}^t) \right] + \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,1}} \left[ \sum_{t \in \mathcal{T}_b} (Y_{\theta''_i}^t - Z_{\theta''_i}^t) \right] \middle| \zeta_b, \xi_{b-1} \right] \\ &= \mathbb{E}_{\xi_b} \left[ \mathbb{E}_{\theta''_i \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} Y_{\theta''_i}^t \right] + \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} Z_{\theta''_i}^t \right] - \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,1}} \left[ \sum_{t \in \mathcal{T}_b} Z_{\theta''_i}^t \right] \middle| \zeta_b, \xi_{b-1} \right], \end{aligned} \quad (18)$$

where the second equality is due to  $\mathcal{J}_p = \mathcal{J}_{p,0} \cup \mathcal{J}_{p,1}$ .

The first term of (18) is

$$\begin{aligned} & \mathbb{E}_{\xi_b} \left[ \mathbb{E}_{\theta''_i \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} Y_{\theta''_i}^t \right] \middle| \zeta_b, \xi_{b-1} \right] \\ &= \mathbb{E}_{\xi_b} \left[ \mathbb{E}_{\theta''_i \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \frac{X_0^t + X_1^t}{2} \xi(\theta''_i)(t) + \left( 1 - \frac{X_0^t + X_1^t}{2} \right) (1 - \xi(\theta''_i)(t)) \right\} \right] \middle| \zeta_b, \xi_{b-1} \right] \\ &= \mathbb{E}_{\theta''_i \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\xi_b} \left[ \frac{X_0^t + X_1^t}{2} \xi(\theta''_i)(t) + \left( 1 - \frac{X_0^t + X_1^t}{2} \right) (1 - \xi(\theta''_i)(t)) \middle| \zeta_b, \xi_{b-1} \right] \right]. \end{aligned}$$

Here we note that  $\mathcal{J}_p$  is determined by  $\zeta_{b-1}$  and  $\xi_{b-1}$ . Since  $X_0^t$  and  $X_1^t$  are determined by the algorithm before  $\xi(\theta''_i)(t)$  is revealed,  $\frac{X_0^t + X_1^t}{2}$  and  $\xi(\theta''_i)(t)$  are independent. Then we can take the expectation of  $\frac{X_0^t + X_1^t}{2}$  and  $\xi(\theta''_i)(t)$  separately. It is also the case for  $1 - \frac{X_0^t + X_1^t}{2}$  and  $1 - \xi(\theta''_i)(t)$ . Since each entry of  $\xi$  follows the uniform distribution over  $\{0, 1\}$  independently,  $\mathbb{E}_{\xi_b} [\xi(\theta''_i)(t) \mid \zeta_b, \xi_{b-1}] = \mathbb{E}_{\xi_b} [1 - \xi(\theta''_i)(t) \mid \zeta_b, \xi_{b-1}] = 1/2$ . We thus obtain

$$\begin{aligned} & \mathbb{E}_{\xi_b} \left[ \mathbb{E}_{\theta''_i \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} Y_{\theta''_i}^t \right] \middle| \zeta_b, \xi_{b-1} \right] \\ &= \mathbb{E}_{\theta''_i \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} \left\{ \frac{1}{2} \mathbb{E}_{\xi_b} \left[ \frac{X_0^t + X_1^t}{2} \middle| \zeta_b, \xi_{b-1} \right] + \frac{1}{2} \mathbb{E}_{\xi_b} \left[ 1 - \frac{X_0^t + X_1^t}{2} \middle| \zeta_b, \xi_{b-1} \right] \right\} \right] \\ &= \frac{|\mathcal{T}_b|}{2} \\ &= \frac{L}{2}. \end{aligned} \quad (19)$$

Next, we give a lower bound on the second and third terms of (18). Since  $Z_{\theta''_i}^t$  can be expressed as

$$Z_{\theta''_i}^t = \frac{X_0^t - X_1^t}{2} \xi(\theta''_i)(t) - \frac{X_0^t - X_1^t}{2} (1 - \xi(\theta''_i)(t)) = \begin{cases} \frac{X_0^t - X_1^t}{2} & \text{if } \xi(\theta''_i)(t) = 1 \\ -\frac{X_0^t - X_1^t}{2} & \text{if } \xi(\theta''_i)(t) = 0, \end{cases}$$

the cumulative sum  $\left(\sum_{t' \in \mathcal{T}_b: t' \leq t} Z_{\theta''_i}^{t'}\right)_{t \in \mathcal{T}_b}$  can be regarded as a random walk with varying step size  $\left|\frac{X_0^t - X_1^t}{2}\right|$  for each  $\theta''_i \in \mathcal{J}_p$ . The step size in round  $t \in \mathcal{T}_b$  is determined by the algorithm that has not observed  $\xi(\theta''_i)(t)$  for any  $\theta''_i \in \Theta''_i$  yet. Since  $X_0^t \in [0, 1]$  and  $X_1^t \in [0, 1]$ , the step size is bounded as  $\left|\frac{X_0^t - X_1^t}{2}\right| \in [0, \frac{1}{2}]$  for each  $t \in \mathcal{T}_b$ .

Recall that how to partition  $\{j_{2k-1}, j_{2k}\}$  into  $\mathcal{J}_{p,0}$  and  $\mathcal{J}_{p,1}$  is determined independently for each  $k \in [2^{B-b}]$ . This is decided by the value of (12), which can be expressed as

$$\begin{aligned} & \mathbb{E}_{\theta''_i \sim \mathcal{I}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta''_i)} [u_i^t(j_{2k-1}, \alpha) - u_i^t(j_{2k}, \alpha)] \right] - \mathbb{E}_{\theta''_i \sim \mathcal{I}_{p,1}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta''_i)} [u_i^t(j_{2k-1}, \alpha) - u_i^t(j_{2k}, \alpha)] \right] \\ &= \left\{ \sum_{t \in \mathcal{T}_b} (Y_{j_{2k-1}}^t + Z_{j_{2k-1}}^t) - \sum_{t \in \mathcal{T}_b} (Y_{j_{2k}}^t + Z_{j_{2k}}^t) \right\} - \left\{ \sum_{t \in \mathcal{T}_b} (Y_{j_{2k-1}}^t - Z_{j_{2k-1}}^t) - \sum_{t \in \mathcal{T}_b} (Y_{j_{2k}}^t - Z_{j_{2k}}^t) \right\} \\ &= 2 \sum_{t \in \mathcal{T}_b} (Z_{j_{2k-1}}^t - Z_{j_{2k}}^t). \end{aligned}$$

By considering the distribution of this value conditioned on  $\zeta_b$  and  $\xi_{b-1}$ , we set  $j_{2k-1} \in \mathcal{J}_{p,0}$  for the upper half of this distribution and  $j_{2k-1} \in \mathcal{J}_{p,1}$  for the lower half. We define the function  $h_k$  such that  $h_k(\xi_b) = 1$  if  $j_{2k-1} \in \mathcal{J}_{p,0}$  and  $h_k(\xi_b) = -1$  if  $j_{2k-1} \in \mathcal{J}_{p,1}$ . Note that  $h_k(\xi_b)$  might be random on the median of the distribution.

Let  $D_k^t = \sum_{t' \in \mathcal{T}_b: t' \leq t} (Z_{j_{2k-1}}^{t'} - Z_{j_{2k}}^{t'})$  be half of the cumulative sum of this value until round  $t$ . From the definition of  $Z_{\theta''_i}^t$ , we have

$$D_k^t = \begin{cases} D_k^{t-1} & \text{with probability } 1/2 \text{ (if } \xi(j_{2k-1})(t) = \xi(j_{2k})(t)) \\ D_k^{t-1} + (X_0^t - X_1^t) & \text{with probability } 1/4 \text{ (if } \xi(j_{2k-1})(t) = 1 \text{ and } \xi(j_{2k})(t) = 0) \\ D_k^{t-1} - (X_0^t - X_1^t) & \text{with probability } 1/4 \text{ (if } \xi(j_{2k-1})(t) = 0 \text{ and } \xi(j_{2k})(t) = 1), \end{cases} \quad (20)$$

where  $X_0^t - X_1^t$  is determined by the algorithm that has not observed  $\xi(j_{2k-1})(t)$  and  $\xi(j_{2k})(t)$  yet.

The second and third terms of (18) are expressed as

$$\begin{aligned} & \mathbb{E}_{\xi_b} \left[ \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} Z_{\theta''_i}^t \right] - \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,1}} \left[ \sum_{t \in \mathcal{T}_b} Z_{\theta''_i}^t \right] \mid \zeta_b, \xi_{b-1} \right] \\ &= \frac{1}{2^{B-b}} \sum_{k=1}^{2^{B-b}} \frac{1}{2} \mathbb{E}_{\xi_b} \left[ h_k(\xi_b) \left\{ \sum_{t \in \mathcal{T}_b} Z_{j_{2k-1}}^t - \sum_{t \in \mathcal{T}_b} Z_{j_{2k}}^t \right\} \mid \zeta_b, \xi_{b-1} \right] \\ &= \frac{1}{2^{B-b+1}} \sum_{k=1}^{2^{B-b}} \mathbb{E}_{\xi_b} [h_k(\xi_b) D_k^{\max} \mid \zeta_b, \xi_{b-1}], \end{aligned} \quad (21)$$

where  $t_{\max} = \max \mathcal{T}_b$  is the index of the last round in block  $b$ .

In the remaining part of the proof, we fix  $k \in [2^{B-b}]$ . Let  $\mathcal{F}$  be the event that  $|D_k^t| > \frac{\sqrt{c^3 L}}{4}$  for some  $t \in \mathcal{T}_b$ . Let  $\mathcal{F}_+$  be the event that  $D_k^t > \frac{\sqrt{c^3 L}}{4}$  for some  $t \in \mathcal{T}_b$  and  $\mathcal{F}_-$  the event that  $D_k^t < -\frac{\sqrt{c^3 L}}{4}$  for some  $t \in \mathcal{T}_b$ . Since  $\Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \leq \Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}) + \Pr_{\xi_b}(\mathcal{F}_- \mid \zeta_b, \xi_{b-1})$ , one of  $\mathcal{F}_+$  or  $\mathcal{F}_-$  happens with probability at least  $\Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1})/2$ . Assume  $\Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}) \geq \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1})/2$  without loss of generality. For the case of  $\Pr_{\xi_b}(\mathcal{F}_- \mid \zeta_b, \xi_{b-1}) \geq \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1})/2$ , we can prove the statement in the same way.

We define  $h'_k$  as follows.

- First, we consider the case of  $\Pr(\mathcal{F}_+) \geq 0.5$ . In this case, we set  $h'_k(\xi_b) = 1$  for the upper part of  $0.5/\Pr(\mathcal{F}_+)$  on the distribution of  $D_k^{t_{\max}}$  conditioned on  $\zeta_b, \xi_{b-1}$ , and  $\mathcal{F}_+$ . We set  $h'_k(\xi_b) = -1$  for the remaining part of  $\mathcal{F}_+$  and the complementary event  $\mathcal{F}_+^c$ .
- Next, we consider the case of  $\Pr(\mathcal{F}_+) < 0.5$ . If  $\mathcal{F}_+$  holds, we set  $h'_k(\xi_k) = 1$ . If  $\mathcal{F}_+$  does not hold, we decide  $h'_k(\xi_k) = 1$  or  $h'_k(\xi_k) = -1$  by considering the distribution of  $D_k^{t_{\max}}$ . For the upper part of probability  $(0.5 - \Pr(\mathcal{F}_+))/(1 - \Pr(\mathcal{F}_+))$  on the distribution of  $D_k^{t_{\max}}$  conditioned on  $\mathcal{F}_+^c$ , we set  $h'_k(\xi_b) = 1$ . For the lower part of probability  $0.5/(1 - \Pr(\mathcal{F}_+))$ , we set  $h'_k(\xi_b) = -1$ .

In both cases, the mass on the boundary can be arbitrarily partitioned so that the above constraint is satisfied. Note that  $h'_k$  satisfies the constraint  $\Pr_{\xi_b}(h'_k(\xi_b) = 1 \mid \zeta_b, \xi_{b-1}) = 0.5$ . Since  $h_k$  maximizes  $\mathbb{E}_{\xi_b}[h_k(\xi_b)D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}]$  under this constraint, we have

$$\mathbb{E}_{\xi_b}[h_k(\xi_b)D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \geq \mathbb{E}_{\xi_b}[h'_k(\xi_b)D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}].$$

In the following, we prove a lower bound

$$\mathbb{E}_{\xi_b}[h_k(\xi_b)D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \geq \frac{1}{2} \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \quad (22)$$

on each term of (21). Since  $D_k^t$  is a random walk specified as (20), it is a martingale. Therefore,  $\mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] = 0$ . From the optimality of  $h_k$ , the left hand side of (22) is non-negative. If  $\mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] < 0$ , (22) immediately follows. In the following, assuming  $\mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \geq 0$ , we prove (22) for each case separately. Let  $D_k^*$  be the maximum value of  $D_k^{t_{\max}}$  when  $h'_k(\xi_b) = -1$  and  $\mathcal{F}_+^c$ .

- First, we consider the case of  $\Pr(\mathcal{F}_+) \geq 0.5$ . It holds that

$$\begin{aligned} & \mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \\ &= \frac{1}{2} \mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1] + \frac{1}{2} \mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = -1], \end{aligned}$$

where we used the fact that the probabilities of  $h'_k(\xi_b) = 1$  and  $h'_k(\xi_b) = -1$  are both  $1/2$ . Combining it with  $\mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] = 0$ , we obtain

$$\frac{1}{2} \mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1] = -\frac{1}{2} \mathbb{E}_{\xi_b}[D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = -1].$$

We thus obtain

$$\begin{aligned}
& \mathbb{E}_{\xi_b} [h_k(\xi_b) D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \\
& \geq \mathbb{E}_{\xi_b} [h'_k(\xi_b) D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \quad (\text{due to the optimality of } h_k) \\
& = \frac{1}{2} \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1] - \frac{1}{2} \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = -1] \\
& = \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1] \\
& \geq \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \quad (\text{since } h'_k(\xi_b) = 1 \text{ holds for the upper part of } \mathcal{F}_+) \\
& \geq \frac{1}{2} \Pr(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+].
\end{aligned}$$

- If  $\Pr(\mathcal{F}_+) < 0.5$  and  $D_k^* \geq 0$ , the former part of the proof in the previous case also applies to this case. The remaining part is

$$\begin{aligned}
& \mathbb{E}_{\xi_b} [h_k(\xi_b) D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \\
& \geq \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1] \\
& = \Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, h'_k(\xi_b) = 1] \\
& \quad + \Pr_{\xi_b}(\mathcal{F}_+^c \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, h'_k(\xi_b) = 1] \\
& \geq \Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}, h'_k(\xi_b) = 1) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, h'_k(\xi_b) = 1] \\
& = \frac{\Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1})}{\Pr_{\xi_b}(h'_k(\xi_b) = 1 \mid \zeta_b, \xi_{b-1})} \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \\
& = 2 \Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \\
& \geq \frac{1}{2} \Pr(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+].
\end{aligned}$$

The first inequality is due to the non-negativity of  $D_k^{t_{\max}}$  conditioned on  $\mathcal{F}_+^c$  and  $h'_k(\xi_b) = 1$ , which is implied by  $D_k^* \geq 0$ . The second equality holds because  $\mathcal{F}_+$  holds only when  $h'_k(\xi_b) = 1$ . The last inequality is due to the assumption  $\Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}) \geq \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1})/2$ .

- If  $\Pr(\mathcal{F}_+) < 0.5$  and  $D_k^* < 0$ , we obtain

$$\begin{aligned}
& \mathbb{E}_{\xi_b} [h_k(\xi_b) D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \\
& \geq \mathbb{E}_{\xi_b} [h'_k(\xi_b) D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \quad (\text{due to the optimality of } h_k) \\
& = \Pr_{\xi_b}(\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \\
& \quad + \Pr_{\xi_b}(\mathcal{F}_+^c \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, h'_k(\xi_b) = 1]
\end{aligned}$$

$$\begin{aligned}
& - \Pr_{\xi_b} (\mathcal{F}_+, h'_k(\xi_b) = -1 \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, h'_k(\xi_b) = -1] \\
& \geq \Pr_{\xi_b} (\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \\
& \quad + \Pr_{\xi_b} (\mathcal{F}_+, h'_k(\xi_b) = 1 \mid \zeta_b, \xi_{b-1}) D_k^* - \Pr_{\xi_b} (\mathcal{F}_+, h'_k(\xi_b) = -1 \mid \zeta_b, \xi_{b-1}) D_k^* \\
& \geq \Pr_{\xi_b} (\mathcal{F}_+ \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \\
& \geq \frac{1}{2} \Pr_{\xi_b} (\mathcal{F} \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+].
\end{aligned}$$

The second inequality is due to the definition of  $D_k^*$ . The third inequality is due to the negativity of  $D_k^*$  and  $\Pr_{\xi_b} (\mathcal{F}_+, h'_k(\xi_b) = 1 \mid \zeta_b, \xi_{b-1}) \leq \Pr_{\xi_b} (\mathcal{F}_+, h'_k(\xi_b) = -1 \mid \zeta_b, \xi_{b-1})$ .

We thus proved (22). In the remaining part of the proof, we give a lower bound on each of  $\Pr_{\xi_b} (\mathcal{F} \mid \zeta_b, \xi_{b-1})$  and  $\mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+]$ .

First, we show that the expected value of  $D_k^{t_{\max}}$  is at least  $\frac{\sqrt{c^3 L}}{4}$  conditioned on  $\mathcal{F}_+$ . Let  $t_+^*$  be the first time round that  $D_k^t \geq \frac{\sqrt{c^3 L}}{4}$  holds, which is well-defined conditioned on  $\mathcal{F}_+$ . For each  $t \in \mathcal{T}_b$ , if we condition on  $\mathcal{F}_+$  and  $t_+^* = t$ , we have

$$\mathbb{E}_{\xi_b} [D_k^{t'+1} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, t_+^* = t] = \mathbb{E}_{\xi_b} [D_k^{t'} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, t_+^* = t]$$

for any  $t' \geq t$  because whether the events  $\mathcal{F}_+$  and  $t_+^* = t$  hold is determined until round  $t'$ . By considering the expected value for different  $t_+^*$  separately,

$$\begin{aligned}
\mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] &= \sum_{t \in \mathcal{T}_b} \Pr_{\xi_b} (t_+^* = t \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, t_+^* = t] \\
&= \sum_{t \in \mathcal{T}_b} \Pr_{\xi_b} (t_+^* = t \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+) \mathbb{E}_{\xi_b} [D_k^t \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+, t_+^* = t] \\
&\geq \sum_{t \in \mathcal{T}_b} \Pr_{\xi_b} (t_+^* = t \mid \mathcal{F}_+) \frac{\sqrt{c^3 L}}{4} \quad (\text{from the definition of } t_+^*) \\
&= \frac{\sqrt{c^3 L}}{4}. \tag{23}
\end{aligned}$$

Next, we provide a lower bound on  $\Pr_{\xi_b} (\mathcal{F} \mid \zeta_b, \xi_{b-1})$  by using the assumption

$$\mathbb{E}_{\xi_b} \left[ \sum_{t \in \mathcal{T}_b} \left| \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right| \mid \zeta_b, \xi_{b-1} \right] \geq cL.$$

Using random variables  $X_0^t$  and  $X_1^t$ , this assumption can be written as

$$\mathbb{E}_{t \sim \mathcal{T}_b, \xi_b} [|X_0^t - X_1^t| \mid \zeta_b, \xi_{b-1}] \geq c.$$

By applying Theorem 4.5, we obtain

$$\Pr_{t \sim \mathcal{T}_b, \xi_b} \left( (X_0^t - X_1^t)^2 \geq \frac{c^2}{4} \mid \zeta_b, \xi_{b-1} \right) = \Pr_{t \sim \mathcal{T}_b, \xi_b} \left( |X_0^t - X_1^t| \geq \frac{c}{2} \mid \zeta_b, \xi_{b-1} \right) \geq \frac{c}{2}.$$

We thus obtain

$$\mathbb{E}_{\xi_b} \left[ \sum_{t \in \mathcal{T}_b} (X_0^t - X_1^t)^2 \middle| \zeta_b, \xi_{b-1} \right] = L \mathbb{E}_{t \sim \mathcal{T}_b, \xi_b} \left[ (X_0^t - X_1^t)^2 \middle| \zeta_b, \xi_{b-1} \right] = L \cdot \frac{c}{2} \cdot \frac{c^2}{4} \geq \frac{c^3 L}{8}. \quad (24)$$

We show that a stochastic process  $(D_k^t)^2 - \frac{1}{2} \sum_{t' \in \mathcal{T}_b: t' \leq t} (X_0^{t'} - X_1^{t'})^2$  for  $t \in \mathcal{T}_b$  is a martingale. We denote by  $\mathcal{F}_t$  all the events that happen until the end of time round  $t$ . Conditioned on  $\mathcal{F}_{t-1}$ , the expected difference is

$$\begin{aligned} & \mathbb{E}_{\xi_b} \left[ (D_k^t)^2 - \frac{1}{2} \sum_{t' \in \mathcal{T}_b: t' \leq t} (X_0^{t'} - X_1^{t'})^2 \middle| \zeta_b, \xi_{b-1}, \mathcal{F}_{t-1} \right] - \left\{ (D_k^{t-1})^2 - \frac{1}{2} \sum_{t' \in \mathcal{T}_b: t' \leq t-1} (X_0^{t'} - X_1^{t'})^2 \right\} \\ &= \mathbb{E}_{\xi_b} [(D_k^t)^2 - (D_k^{t-1})^2 \middle| \zeta_b, \xi_{b-1}, \mathcal{F}_{t-1}] - \frac{1}{2} (X_0^t - X_1^t)^2 \\ &= \mathbb{E}_{\xi_b} [(D_k^t - D_k^{t-1})(D_k^t + D_k^{t-1}) \middle| \zeta_b, \xi_{b-1}, \mathcal{F}_{t-1}] - \frac{1}{2} (X_0^t - X_1^t)^2 \\ &= \mathbb{E}_{\xi_b} [(D_k^t - D_k^{t-1})(2D_k^{t-1} + (D_k^t - D_k^{t-1})) \middle| \zeta_b, \xi_{b-1}, \mathcal{F}_{t-1}] - \frac{1}{2} (X_0^t - X_1^t)^2 \\ &= 2D_k^{t-1} \mathbb{E}_{\xi_b} [D_k^t - D_k^{t-1} \middle| \zeta_b, \xi_{b-1}, \mathcal{F}_{t-1}] + \mathbb{E}_{\xi_b} [(D_k^t - D_k^{t-1})^2 \middle| \zeta_b, \xi_{b-1}, \mathcal{F}_{t-1}] - \frac{1}{2} (X_0^t - X_1^t)^2 \\ &= 0 + \frac{1}{2} (X_0^t - X_1^t)^2 - \frac{1}{2} (X_0^t - X_1^t)^2 \quad (\text{since (20)}) \\ &= 0. \end{aligned}$$

We define the random variable  $t^* \in \mathcal{T}_b$  such that

$$t^* = \begin{cases} \min \left\{ t \in \mathcal{T}_b \mid |D_k^t| \geq \frac{\sqrt{c^3 L}}{4} \right\} & \text{if } \mathcal{F} \text{ holds} \\ t_{\max} & \text{otherwise} \end{cases}$$

is the time round when  $|D_k^t|$  exceeds  $\frac{\sqrt{c^3 L}}{4}$  for the first time if  $\mathcal{F}$  holds. Since the stopping time  $t^*$  is bounded above by  $t_{\max}$ , we can apply Doob's optional stopping theorem to the martingale defined above and obtain

$$\mathbb{E}_{\xi_b} \left[ (D_k^{t^*})^2 - \sum_{t \in \mathcal{T}_b: t \leq t^*} (X_0^t - X_1^t)^2 \middle| \zeta_b, \xi_{b-1} \right] = 0.$$

Since  $\sum_{t \in \mathcal{T}_b: t > t^*} (X_0^t - X_1^t)^2$  is 0 if  $\mathcal{F}$  does not hold and bounded above by  $L$  if  $\mathcal{F}$  holds, we have

$$\begin{aligned} & \mathbb{E}_{\xi_b} \left[ \sum_{t \in \mathcal{T}_b} (X_0^t - X_1^t)^2 \middle| \zeta_b, \xi_{b-1} \right] \\ &= \mathbb{E}_{\xi_b} \left[ \sum_{t \in \mathcal{T}_b: t \leq t^*} (X_0^t - X_1^t)^2 \middle| \zeta_b, \xi_{b-1} \right] + \mathbb{E}_{\xi_b} \left[ \sum_{t \in \mathcal{T}_b: t > t^*} (X_0^t - X_1^t)^2 \middle| \zeta_b, \xi_{b-1} \right] \\ &\leq \mathbb{E}_{\xi_b} \left[ (D_k^{t^*})^2 \middle| \zeta_b, \xi_{b-1} \right] + \Pr(\mathcal{F} \mid \zeta_b, \xi_{b-1}) L \end{aligned}$$

$$\begin{aligned}
&= \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} \left[ (D_k^{t*})^2 \mid \zeta_b, \xi_{b-1}, \mathcal{F} \right] + \Pr_{\xi_b}(\mathcal{F}^c \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} \left[ (D_k^{t*})^2 \mid \zeta_b, \xi_{b-1}, \mathcal{F}^c \right] + \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) L \\
&\leq \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \left( \frac{\sqrt{c^3 L}}{4} + 1 \right)^2 + \Pr_{\xi_b}(\mathcal{F}^c \mid \zeta_b, \xi_{b-1}) \left( \frac{\sqrt{c^3 L}}{4} \right)^2 + \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) L \\
&\leq 4L \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) + \frac{c^3 L}{16},
\end{aligned}$$

where we used  $c \leq 1$ ,  $\sqrt{L} \leq L$ ,  $L \geq 1$ , and  $\Pr_{\xi_b}(\mathcal{F}^c \mid \zeta_b, \xi_{b-1}) \leq 1$  for the last inequality. Combining it with (24), we obtain

$$\Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \geq \frac{c^3 L / 8 - c^3 L / 16}{4L} = \frac{c^3}{64}. \quad (25)$$

Plugging (23) and (25) into (22), we obtain

$$\begin{aligned}
&\mathbb{E}_{\xi_b} \left[ \frac{1}{2} \mathbb{E}_{\substack{\theta'_i \sim \mathcal{I}_{p,0} \\ \theta''_i \sim \mathcal{J}_{p,0}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] + \frac{1}{2} \mathbb{E}_{\substack{\theta'_i \sim \mathcal{I}_{p,1} \\ \theta''_i \sim \mathcal{J}_{p,1}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta'_i)} [u_i^t(\theta''_i, \alpha)] \right] \middle| \zeta_b, \xi_{b-1} \right] \\
&= \mathbb{E}_{\xi_b} \left[ \mathbb{E}_{\theta''_i \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} Y_{\theta''_i}^t \right] + \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} Z_{\theta''_i}^t \right] - \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,1}} \left[ \sum_{t \in \mathcal{T}_b} Z_{\theta''_i}^t \right] \middle| \zeta_b, \xi_{b-1} \right] \quad (\text{since (18)}) \\
&= \frac{L}{2} + \frac{1}{2^{B-b+1}} \sum_{k=1}^{2^{B-b}} \mathbb{E}_{\xi_b} [h_k(\xi_b) D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}] \quad (\text{since (19) and (21)}) \\
&\geq \frac{L}{2} + \frac{1}{2^{B-b+1}} \sum_{k=1}^{2^{B-b}} \frac{1}{2} \Pr_{\xi_b}(\mathcal{F} \mid \zeta_b, \xi_{b-1}) \mathbb{E}_{\xi_b} [D_k^{t_{\max}} \mid \zeta_b, \xi_{b-1}, \mathcal{F}_+] \quad (\text{since (22)}) \\
&\geq \frac{L}{2} + \frac{1}{2^{B-b+1}} \sum_{k=1}^{2^{B-b}} \frac{1}{2} \cdot \frac{c^3}{64} \cdot \frac{\sqrt{c^3 L}}{4} \quad (\text{since (23) and (25)}) \\
&= \frac{L}{2} + 2^{-10} c^{9/2} \sqrt{L},
\end{aligned}$$

which completes the proof.  $\square$

#### 4.4 Proof of lower bound

Finally, we prove the theorem using Theorems 4.4, 4.6 and 4.7.

*Proof of Theorem 4.1.* Since the rewards for type  $\theta''_i \in \Theta''_i$  are determined by independent Bernoulli random variables  $\xi$ , and therefore the expected reward obtained by the algorithm for each round is equal to 1/2. Formally, since the algorithm's decision  $\pi_i^t(\theta''_i)$  is independent from  $u_i^t(\theta''_i, \alpha_0) = 1 - u_i^t(\theta''_i, \alpha_1) = \xi(\theta''_i)(t)$ , the expected reward is

$$\begin{aligned}
\mathbb{E}_{\zeta, \xi} \left[ \mathbb{E}_{\alpha \sim \pi_i^t(\theta''_i)} [u_i^t(\theta''_i, \alpha)] \right] &= \mathbb{E}_{\zeta, \xi} [\pi_i^t(\theta''_i; \alpha_0) \xi(\theta''_i)(t) + (1 - \pi_i^t(\theta''_i; \alpha_0))(1 - \xi(\theta''_i)(t))] \\
&= \pi_i^t(\theta''_i; \alpha_0) \cdot \frac{1}{2} + (1 - \pi_i^t(\theta''_i; \alpha_0)) \cdot \frac{1}{2}
\end{aligned}$$

$$= \frac{1}{2}.$$

We consider  $\phi$  in the definition of untruthful swap regret such that  $\phi(\theta_i, \cdot)$  is the identity map for every  $\theta_i \in \Theta_i$ . By considering only the types in  $\Theta_i''$ , we obtain

$$\begin{aligned} \mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] &\geq \mathbb{E}_{\zeta, \xi} \left[ \max_{\psi: \Theta_i \rightarrow \Theta_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{\alpha \sim \pi_i^t(\psi(\theta_i))} [u_i^t(\theta_i, \alpha)] - \mathbb{E}_{\alpha \sim \pi_i^t(\theta_i)} [u_i^t(\theta_i, \alpha)] \right] \right] \\ &\geq \mathbb{E}_{\zeta, \xi} \left[ \max_{\psi: \Theta_i \rightarrow \Theta_i} \frac{1}{|\Theta_i|} \sum_{\theta_i'' \in \Theta_i''} \sum_{t=1}^T \left\{ \mathbb{E}_{\alpha \sim \pi_i^t(\psi(\theta_i''))} [u_i^t(\theta_i'', \alpha)] - \mathbb{E}_{\alpha \sim \pi_i^t(\theta_i'')} [u_i^t(\theta_i'', \alpha)] \right\} \right] \\ &= \mathbb{E}_{\zeta, \xi} \left[ \max_{\psi: \Theta_i \rightarrow \Theta_i} \frac{1}{|\Theta_i|} \sum_{\theta_i'' \in \Theta_i''} \sum_{t=1}^T \mathbb{E}_{\alpha \sim \pi_i^t(\psi(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] - \frac{1}{|\Theta_i|} \sum_{\theta_i'' \in \Theta_i''} \sum_{t=1}^T \frac{1}{2} \\ &= \frac{1}{2} \mathbb{E}_{\zeta, \xi} \left[ \max_{\psi: \Theta_i \rightarrow \Theta_i} \frac{1}{|\Theta_i''|} \sum_{\theta_i'' \in \Theta_i''} \sum_{t=1}^T \mathbb{E}_{\alpha \sim \pi_i^t(\psi(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] - \frac{T}{4}, \end{aligned}$$

where the last equality is due to  $|\Theta_i''| = |\Theta_i|/2$ . By using  $\psi'_{\zeta, \xi}$  defined in Section 4.2 for each  $\zeta$  and  $\xi$ , we obtain

$$\begin{aligned} \mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] &\geq \frac{1}{2} \mathbb{E}_{\zeta, \xi} \left[ \frac{1}{|\Theta_i''|} \sum_{\theta_i'' \in \Theta_i''} \sum_{t=1}^T \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] - \frac{T}{4} \\ &= \frac{1}{2} \sum_{b \in [B]} \mathbb{E}_{\zeta, \xi} \left[ \frac{1}{|\Theta_i''|} \sum_{\theta_i'' \in \Theta_i''} \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta_i''))} [u_i^t(\theta_i'', \alpha)] - \frac{L}{2} \right] \\ &= \frac{B}{2} \left\{ \mathbb{E}_{b \sim [B]} \left[ \mathbb{E}_{\zeta, \xi} \left[ \mathbb{E}_{\theta_i'' \sim \Theta_i''} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] \right] \right] - \frac{L}{2} \right\}. \quad (26) \end{aligned}$$

For each  $b \in [B]$ , since  $(\mathcal{J}_p)_{p \in \{0,1\}^{b-1}}$  is the equally sized partition of  $\Theta_i''$ , each term can be expressed as

$$\begin{aligned} &\mathbb{E}_{\zeta, \xi} \left[ \mathbb{E}_{\theta_i'' \sim \Theta_i''} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] \right] \\ &= \mathbb{E}_{\zeta, \xi} \left[ \mathbb{E}_{p \sim \{0,1\}^{b-1}} \left[ \mathbb{E}_{\theta_i'' \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] \right] \right] \\ &= \mathbb{E}_{\substack{\zeta_b, \xi_{b-1} \\ p \sim \{0,1\}^{b-1}}} \left[ \mathbb{E}_{\zeta, \xi} \left[ \mathbb{E}_{\theta_i'' \sim \mathcal{J}_p} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] \middle| \zeta_b, \xi_{b-1} \right] \right] \\ &= \mathbb{E}_{\substack{\zeta_b, \xi_{b-1} \\ p \sim \{0,1\}^{b-1}}} \left[ \mathbb{E}_{\zeta, \xi} \left[ \frac{1}{2} \mathbb{E}_{\theta_i'' \sim \mathcal{J}_{p,0}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta_i''))} [u_i^t(\theta_i'', \alpha)] \right] \right] \right] \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,1}} \left[ \mathbb{E}_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta''_i))} [u_i^t(\theta''_i, \alpha)] \middle| \zeta_b, \xi_{b-1} \right] \\
= & \mathbb{E}_{\substack{\zeta_b, \xi_{b-1} \\ p \sim \{0,1\}^{b-1}}} \left[ \mathbb{E}_{\xi_b} \left[ \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,0}} \left[ \mathbb{E}_{\zeta, \xi} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta''_i))} [u_i^t(\theta''_i, \alpha)] \middle| \zeta_b, \xi_b \right] \right. \right. \right. \\
& \left. \left. \left. + \frac{1}{2} \mathbb{E}_{\theta''_i \sim \mathcal{J}_{p,1}} \left[ \mathbb{E}_{\zeta, \xi} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta''_i))} [u_i^t(\theta''_i, \alpha)] \middle| \zeta_b, \xi_b \right] \right] \middle| \zeta_b, \xi_{b-1} \right] \right].
\end{aligned}$$

Conditioned on  $\zeta_b$  and  $\xi_b$ , all of  $\mathcal{J}_{p,0}$ ,  $\mathcal{J}_{p,1}$ ,  $(\pi_i^t)_{t \in \mathcal{T}_b}$ , and  $(u_i^t)_{t \in \mathcal{T}_b}$  are deterministic. For taking the expectation with respect to  $\zeta$  and  $\xi$  conditioned on  $\zeta_b$  and  $\xi_b$ , we need to consider only the randomness of  $\psi'_{\zeta, \xi}(\theta''_i)$ . From Theorem 4.6,  $\psi'_{\zeta, \xi}(\theta''_i)$  is distributed uniformly over  $\mathcal{I}_{p,0}$  for each  $\theta''_i \in \mathcal{J}_{p,0}$ . Similarly,  $\psi'_{\zeta, \xi}(\theta''_i)$  is distributed uniformly over  $\mathcal{I}_{p,1}$  for each  $\theta''_i \in \mathcal{J}_{p,1}$ . Therefore, the above term can be expressed as

$$\begin{aligned}
& \mathbb{E}_{\zeta, \xi} \left[ \mathbb{E}_{\theta''_i \sim \Theta''_i} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta''_i))} [u_i^t(\theta''_i, \alpha)] \right] \right] \\
= & \mathbb{E}_{\substack{\zeta_b, \xi_{b-1} \\ p \sim \{0,1\}^{b-1}}} \left[ \mathbb{E}_{\xi_b} \left[ \frac{1}{2} \mathbb{E}_{\substack{\theta''_i \sim \mathcal{I}_{p,0} \\ \theta''_i \sim \mathcal{J}_{p,0}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta''_i)} [u_i^t(\theta''_i, \alpha)] \right] + \frac{1}{2} \mathbb{E}_{\substack{\theta''_i \sim \mathcal{I}_{p,0} \\ \theta''_i \sim \mathcal{J}_{p,1}}} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\theta''_i)} [u_i^t(\theta''_i, \alpha)] \right] \middle| \zeta_b, \xi_{b-1} \right] \right].
\end{aligned}$$

For each  $b \in [B]$ ,  $\zeta_b$ ,  $\xi_{b-1}$ , and  $p \in \{0,1\}^{b-1}$ , let  $\mathcal{F}_{b, \zeta_b, \xi_{b-1}, p}$  be the event that

$$\mathbb{E}_{\xi_b} \left[ \sum_{t \in \mathcal{T}_b} \left| \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,0}} [\pi_i^t(\theta_i; \alpha_0)] - \mathbb{E}_{\theta_i \sim \mathcal{I}_{p,1}} [\pi_i^t(\theta_i; \alpha_0)] \right| \middle| \zeta_b, \xi_{b-1} \right] \geq \frac{L}{8}.$$

Toward a contradiction, assume  $\mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] < 2^{-28} \sqrt{T \log_2 |\Theta_i|}$ . From Theorem 4.4, we have  $\Pr_{b, \zeta_b, \xi_{b-1}, p}(\mathcal{F}_{b, \zeta_b, \xi_{b-1}, p}) \geq 1/8$ . Next, we apply Theorem 4.7 to each  $b$ ,  $\zeta_b$ ,  $\xi_{b-1}$  and  $p$ . If  $\mathcal{F}_{b, \zeta_b, \xi_{b-1}, p}$  holds, we apply Theorem 4.7 with  $c = 1/8$ . If  $\mathcal{F}_{b, \zeta_b, \xi_{b-1}, p}$  does not hold, since the assumption of Theorem 4.7 always holds for  $c = 0$ , we apply Theorem 4.7 with  $c = 0$ . Then we obtain a lower bound on the above term as

$$\begin{aligned}
& \mathbb{E}_{b \sim [B]} \left[ \mathbb{E}_{\zeta, \xi} \left[ \mathbb{E}_{\theta''_i \sim \Theta''_i} \left[ \sum_{t \in \mathcal{T}_b} \mathbb{E}_{\alpha \sim \pi_i^t(\psi'_{\zeta, \xi}(\theta''_i))} [u_i^t(\theta''_i, \alpha)] \right] \right] \right] \\
& \geq \Pr_{b \sim [B], \zeta_b, \xi_{b-1}, p \sim \{0,1\}^{b-1}} (\mathcal{F}_{b, \zeta_b, \xi_{b-1}, p}) \left( \frac{L}{2} + 2^{-24} \sqrt{L} \right) + \Pr_{b \sim [B], \zeta_b, \xi_{b-1}, p \sim \{0,1\}^{b-1}} (\mathcal{F}_{b, \zeta_b, \xi_{b-1}, p}^c) \frac{L}{2} \\
& = \frac{L}{2} + \Pr_{b \sim [B], \zeta_b, \xi_{b-1}, p \sim \{0,1\}^{b-1}} (\mathcal{F}_{b, \zeta_b, \xi_{b-1}, p}) \cdot 2^{-24} \sqrt{L} \\
& \geq \frac{L}{2} + 2^{-27} \sqrt{L}.
\end{aligned}$$

Finally, substituting this into (26), we obtain

$$\begin{aligned}
\mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] &\geq \frac{B}{2} \cdot 2^{-27} \sqrt{L} \\
&= 2^{-28} \cdot B \cdot \sqrt{T/B} && \text{(since } L = T/B\text{)} \\
&= 2^{-28} \sqrt{TB} \\
&\geq 2^{-28} \sqrt{T \log_2 |\Theta_i|},
\end{aligned}$$

which contradicts the assumption  $\mathbb{E}_{\zeta, \xi} [R_{\text{US}, i}^T] < 2^{-28} \sqrt{T \log_2 |\Theta_i|}$ . Then there exist some  $\zeta$  and  $\xi$  such that for the problem instance with parameter  $\zeta$  and  $\xi$ , it holds  $R_{\text{US}, i}^T \geq 2^{-28} \sqrt{T \log_2 |\Theta_i|}$ .  $\square$

## 5 Smoothness and price of anarchy

In this section, we show lower bounds on the price of anarchy (POA) in Bayesian games via the smoothness assumption. As in existing studies [Roughgarden, 2015b, Syrgkanis, 2012, Syrgkanis and Tardos, 2013], we make the following two assumptions throughout this section.

- Assumption 1: the prior distribution  $\rho \in \Delta(\Theta)$  is a product distribution, that is, there exists some  $\rho_i \in \Delta(\Theta_i)$  for each  $i \in N$  such that  $\rho(\theta) = \prod_{i \in N} \rho_i(\theta_i)$ .
- Assumption 2: the value of  $v_i(\theta; a)$  does not depend on  $\theta_{-i}$ . Under this assumption, we write  $v_i(\theta_i; a)$  in place of  $v_i(\theta; a)$  by abuse of notation.

These assumptions naturally hold when type  $\theta_i$  represents a preference or attribute of each player  $i \in N$ . We define the price of anarchy as follows.

**Definition 5.1** (Price of anarchy of Bayesian games). Given an equilibrium class  $\Pi \subseteq \Delta(A)^\Theta$ , its price of anarchy is defined by

$$\text{POA}_\Pi = \frac{\inf_{\pi \in \Pi} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right]}{\mathbb{E}_{\theta \sim \rho} \left[ \max_{a \in A} v_{\text{SW}}(\theta; a) \right]},$$

where  $v_{\text{SW}}: \Theta \times A \rightarrow [0, 1]$  is the social welfare function.

Most of the existing studies considered the case where  $\Pi$  is Bayes Nash equilibria. There are several exceptions [Caragiannis et al., 2015, Hartline et al., 2015] in which  $\Pi$  is some class of Bayes coarse correlated equilibria, but their applications are limited (see Section 1.3). Throughout this section,  $\Pi$  is the set of (non-approximate) communication equilibria with strategy representability. As described in Sections 2 and 3, an equilibrium in this class can be efficiently computed by simulating no-untruthful-swap-regret dynamics.

We prepare a lemma that is commonly used in all three cases. The proof is based on the ones provided by Roughgarden [2015b], Syrgkanis [2012], Syrgkanis and Tardos [2013], but we here extend them from Bayes Nash equilibria to communication equilibria with strategy representability. We use the characterization of this class given in Theorem 2.18.

**Lemma 5.2.** *Suppose that a Bayesian game satisfies Assumptions 1 and 2. Fix any  $i \in N$ . Given any  $a_{i,\theta,a_i}^* \in A_i$  for each  $\theta \in \Theta$  and  $a_i \in A_i$ , it holds that*

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta_i; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} [v_i(\theta_i; a_{i,\theta,a_i}^*, a_{-i})] \right] \right],$$

where  $\pi \in \Delta(A)^\Theta$  is any communication equilibrium with strategy representability.

*Proof.* Let  $\pi \in \Delta(A)^\Theta$  be any communication equilibrium with strategy representability. From Theorem 2.18, there exists  $\sigma \in \Delta(S)$  that satisfies  $\eta(\sigma) = \pi$  and (IC<sub>ComSR</sub>) for each  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . Recall that  $\eta: \Delta(S) \rightarrow \Delta(A)^\Theta$  is defined as  $(\eta(\sigma))(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a)$  for each  $\theta \in \Theta$  and  $a \in A$  (see Section 2.1). Since  $\eta(\sigma) = \pi$ , the expected payoff for player  $i$  in this equilibrium is

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta_i; a)] \right] = \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right]. \quad (27)$$

Next, we apply (IC<sub>ComSR</sub>) for each  $i \in N$ . Fix any  $\theta' \in \Theta$ . If we set  $\psi(\theta_i) = \theta'_i$  for each  $\theta_i \in \Theta_i$  and  $\phi(\theta_i, a_i) = a_{i,(\theta_i, \theta'_{-i}),a_i}^*$  for each  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ , then (IC<sub>ComSR</sub>) implies

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; a_{i,(\theta_i, \theta'_{-i}),s_i(\theta'_i)}, s_{-i}(\theta'_{-i}))] \right].$$

By taking the expectation over  $\theta' \sim \rho$ , we obtain

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; a_{i,(\theta_i, \theta'_{-i}),s_i(\theta'_i)}, s_{-i}(\theta'_{-i}))] \right] \right].$$

Since  $\rho$  is a product distribution, we can swap  $\theta_{-i}$  and  $\theta'_{-i}$  on the right-hand side and obtain

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; a_{i,\theta,s_i(\theta'_i)}, s_{-i}(\theta'_{-i}))] \right] \right].$$

By plugging this inequality into (27), we obtain

$$\begin{aligned} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta_i; a)] \right] &\geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; a_{i,\theta,s_i(\theta'_i)}, s_{-i}(\theta'_{-i}))] \right] \right] \\ &= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} [v_i(\theta_i; a_{i,\theta,a_i}^*, a_{-i})] \right] \right], \end{aligned}$$

where the equality holds since the distribution of  $s(\theta')$  equals that of  $a \sim \pi(\theta')$  due to  $\pi = \eta(\sigma)$ .  $\square$

## 5.1 POA bounds for the sum of payoffs

First, we assume that the social welfare is the sum of all players' payoffs, i.e.,  $v_{\text{SW}}(\theta; a) = \sum_{i \in N} v_i(\theta_i; a)$ . Under this assumption, Syrgkanis [2012] proved a POA lower bound with (unconditional) smoothness, and Syrgkanis and Tardos [2013] then extended it to a weaker version of smoothness, called *conditional smoothness*.

**Definition 5.3** ( $(\lambda, \mu)$ -Conditional smoothness of Bayesian games [Syrngkanis and Tardos, 2013]). A Bayesian game is  $(\lambda, \mu)$ -conditionally smooth if there exists some  $a_{i,\theta,a_i}^* \in A_i$  for each  $i \in N$ ,  $\theta \in \Theta$ , and  $a_i \in A_i$  such that

$$\sum_{i \in N} v_i(\theta_i; a_{i,\theta,a_i}^*, a_{-i}) \geq \lambda \max_{a' \in A} v_{\text{SW}}(\theta; a') - \mu v_{\text{SW}}(\theta; a)$$

holds for any  $\theta \in \Theta$  and  $a \in A$ .

Syrngkanis and Tardos [2013] introduced conditional smoothness for analyzing the POA of Bayes Nash equilibria. We extend their POA bounds to communication equilibria with strategy representability. Note that if (unconditional) smoothness is assumed, Theorem 5.2 can be proved for a broader class of Bayes coarse correlated equilibria. However, even under the (unconditional) smoothness assumption, the following proof requires both the incentive constraints for communication equilibria and strategy representability.

**Theorem 5.4.** *Assume that the social welfare is the sum of all players' payoffs. For any  $(\mu, \lambda)$ -conditionally smooth Bayesian game satisfying Assumptions 1 and 2, the price of anarchy of communication equilibria with strategy representability is at least  $\lambda/(1 + \mu)$ .*

*Proof.* Let  $\pi \in \Delta(A)^\Theta$  be any communication equilibrium with strategy representability. By taking the summation of Theorem 5.2 over  $i \in N$ , we obtain

$$\begin{aligned} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right] &= \sum_{i \in N} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta_i; a)] \right] && \text{(due to the assumption)} \\ &\geq \sum_{i \in N} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} [v_i(\theta_i; a_{i,\theta,a_i}^*, a_{-i})] \right] \right] \\ &\geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} \left[ \lambda \max_{a' \in A} v_{\text{SW}}(\theta; a') - \mu v_{\text{SW}}(\theta; a) \right] \right] \right] \\ &&& \text{(due to conditional smoothness)} \\ &= \lambda \mathbb{E}_{\theta \sim \rho} \left[ \max_{a' \in A} v_{\text{SW}}(\theta; a') \right] - \mu \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} [v_{\text{SW}}(\theta; a)] \right] \right]. \end{aligned} \quad (28)$$

As in the proof of Theorem 5.2, we use  $\sigma \in \Delta(S)$  defined in Theorem 2.18. By using the assumption of  $v_{\text{SW}}$ , we obtain

$$\begin{aligned} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} [v_{\text{SW}}(\theta; a)] \right] \right] &= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_{\text{SW}}(\theta; s(\theta'))] \right] \right] \\ &= \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} \left[ \sum_{i \in N} v_i(\theta_i; s(\theta')) \right] \right] \right] \\ &= \sum_{i \in N} \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta'))] \right] \right] \\ &= \sum_{i \in N} \mathbb{E}_{\theta'_i \sim \rho_i} \left[ \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s_i(\theta'_i), s_{-i}(\theta_{-i}))] \right] \right], \end{aligned} \quad (29)$$

where in the last equality,  $\theta'_{-i}$  is replaced with  $\theta_{-i}$  since  $\rho$  is a product distribution. Next, we apply (IC<sub>ComSR</sub>) for each  $i \in N$ . Fix any  $\theta'_i \in \Theta_i$ . If we set  $\psi(\theta_i) = \theta'_i$  for any  $\theta_i \in \Theta_i$  and  $\phi(\theta_i, a_i) = a_i$  for any  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ , then (IC<sub>ComSR</sub>) implies

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s_i(\theta'_i), s_{-i}(\theta_{-i}))] \right].$$

By taking the expectation over  $\theta'_i \sim \rho_i$ , we obtain

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right] \geq \mathbb{E}_{\theta'_i \sim \rho_i} \left[ \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s_i(\theta'_i), s_{-i}(\theta_{-i}))] \right] \right].$$

By plugging this inequality into (29), we obtain

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} [v_{\text{SW}}(\theta; a)] \right] \right] \leq \sum_{i \in N} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta_i; s(\theta))] \right] = \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right].$$

From (28), we obtain

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right] \geq \lambda \mathbb{E}_{\theta \sim \rho} \left[ \max_{a' \in A} v_{\text{SW}}(\theta; a') \right] - \mu \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right],$$

and then rearranging terms leads to a desired lower bound on the price of anarchy.  $\square$

**Remark 5.5.** In the proof of Theorem 5.4 (and also in the proof of Theorem 5.2), we use (IC<sub>ComSR</sub>) only for constant map  $\psi$ . We can prove the same POA bound for the equilibrium concept with  $\psi$  restricted to constant maps, which is broader than communication equilibria with strategy representability. However, since each player  $i \in N$  knows their type  $\theta_i$  by definition, the equilibrium concept in which each player ignores their type is artificial.

## 5.2 POA bounds for conditionally smooth mechanisms

Next, we consider applications to mechanism design. Since it is not the focus of this paper, we do not introduce mechanism design in details. See, e.g., Nisan et al. [2007, Chapter 9] or Hartline [2013] for basics of mechanism design and Syrgkanis and Tardos [2013] for smooth mechanisms.

A mechanism (for auctions) is regarded as a function that maps type reports (bids)  $a \in A$  from players to an allocation and payment. Let  $X_i$  be the set of allocation for player  $i \in N$  and  $X \subseteq \prod_{i \in N} X_i$  the set of all possible allocations. An allocation function  $f_i: A \rightarrow X_i$  for each  $i \in N$  determines the allocation for player  $i$  based on the bids  $a \in A$ . We can consider the process of mechanisms as a game in which each player decides a bid  $a_i \in A_i$  and then obtains a payoff based on the allocation and the payment. We assume that each player's utility function  $v_i$  is *quasilinear*, i.e., there exist some valuation function  $v_i^+: \Theta_i \times X_i \rightarrow [0, 1]$  and payment function  $v_i^-: A \rightarrow [0, 1]$  such that

$$v_i(\theta_i; a) = v_i^+(\theta_i; f_i(a)) - v_i^-(a).$$

An important assumption here is the payment function depends only on  $a$ , not on  $\theta_i$ . This assumption naturally holds because the payment is determined by the mechanism that does not know each player's type. We also assume that the payoff of each player is always non-negative, i.e.,  $v_i(\theta_i; a) \geq 0$

for any  $a \in A$ . This should hold because each player can withdraw from the mechanism if the payoff is negative. The social welfare function is defined as  $v_{\text{SW}}(\theta; a) = \sum_{i \in N} v_i^+(\theta_i; f_i(a))$  for each  $\theta \in \Theta$  and  $a \in A$ .

Syrngkanis and Tardos [2013] proved that various simple Bayesian mechanisms and their simultaneous/sequential compositions are conditionally smooth.

**Definition 5.6** ( $(\lambda, \mu)$ -Conditional smoothness of Bayesian mechanisms [Syrngkanis and Tardos, 2013]). A Bayesian mechanism is  $(\lambda, \mu)$ -conditionally smooth if there exists some  $a_{i,\theta,a_i}^* \in A_i$  for each  $i \in N$ ,  $\theta \in \Theta$ , and  $a_i \in A_i$  such that

$$\sum_{i \in N} v_i(\theta_i; a_{i,\theta,a_i}^*, a_{-i}) \geq \lambda \max_{x \in X} \sum_{i \in N} v_i^+(\theta_i; x_i) - \mu \sum_{i \in N} v_i^-(a).$$

holds for any  $\theta \in \Theta$  and  $a \in A$ .

Syrngkanis and Tardos [2013] proved a  $\lambda/\max\{1, \mu\}$  lower bound on the POA of Bayes Nash equilibria with this assumption. We extend it to communication equilibria with strategy representability.

**Theorem 5.7.** *For any  $(\mu, \lambda)$ -conditionally smooth mechanism satisfying Assumptions 1 and 2, the price of anarchy of communication equilibria with strategy representability is at least  $\lambda/\max\{1, \mu\}$ .*

*Proof.* Let  $\pi \in \Delta(A)^\Theta$  be any communication equilibrium with strategy representability. By summing Theorem 5.2 for each  $i \in N$ , we obtain

$$\begin{aligned} \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} \left[ \sum_{i \in N} v_i(\theta_i; a) \right] \right] &\geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} \left[ \sum_{i \in N} v_i(\theta_i; a_{i,\theta,a_i}^*, a_{-i}) \right] \right] \right] \\ &\geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} \left[ \lambda \max_{x \in X} \sum_{i \in N} v_i^+(\theta_i; x_i) - \mu \sum_{i \in N} v_i^-(a) \right] \right] \right] \\ &\hspace{15em} \text{(due to conditional smoothness)} \\ &= \lambda \mathbb{E}_{\theta \sim \rho} \left[ \max_{x \in X} \sum_{i \in N} v_i^+(\theta_i; x_i) \right] - \mu \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} \left[ \sum_{i \in N} v_i^-(a) \right] \right] \\ &= \lambda \mathbb{E}_{\theta \sim \rho} \left[ \max_{a \in A} v_{\text{SW}}(\theta; a) \right] - \mu \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} \left[ \sum_{i \in N} v_i^-(a) \right] \right]. \end{aligned}$$

Since  $v_i(\theta_i; a) = v_i^+(\theta_i; f_i(a)) - v_i^-(a)$  for each  $i \in N$ , we obtain

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right] \geq \lambda \mathbb{E}_{\theta \sim \rho} \left[ \max_{a \in A} v_{\text{SW}}(\theta; a) \right] + (1 - \mu) \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} \left[ \sum_{i \in N} v_i^-(a) \right] \right].$$

If  $\mu \leq 1$ , the second term on the right-hand side is non-negative, and therefore, the POA is at least  $\lambda$ . If  $\mu > 1$ , since  $v_i(\theta_i; a) = v_i^+(\theta_i; f_i(a)) - v_i^-(a) \geq 0$ , the second term on the right-hand side is bounded below by  $(1 - \mu)\mathbb{E}_{\theta \sim \rho} [\mathbb{E}_{s \sim \sigma} [v_{\text{SW}}(\theta; s(\theta))]]$ , which implies a  $\lambda/\mu$  lower bound on the POA.  $\square$

Syrkkanis and Tardos [2013] also proved a similar result for weak smoothness and budget constraints. We can generalize our proof to these settings and obtain POA bounds for more various mechanisms. These results yield lower bounds on the POA of communication equilibria with strategy representability for the simultaneous and sequential composition of first-price, second-price, all-pay, and many other auctions.

**Remark 5.8.** In applications to auctions, a type  $\theta_i$  represents a valuation of an item for player  $i$ , and an action  $a_i \in A_i$  represents a bid for this good. The sets of possible values  $\Theta_i$  and  $A_i$  are usually continuous intervals. Our proposed algorithm for computing an equilibrium does not assume any structure for  $\Theta_i$  and  $A_i$  and requires polynomial time in  $|\Theta_i|$  and  $|A_i|$ , which might be prohibitive. A naive solution to this issue is to discretize  $\Theta_i$  and  $A_i$ . Another possible solution is to design an efficient algorithm using the continuous structure of  $\Theta_i$  and  $A_i$ , which is left for future work.

### 5.3 POA bounds for type-dependent action sets

We consider the case where the set of actions for each player changes depending on their type. Here we briefly describe how we can extend the concepts introduced so far and then provide a lower bound on the POA.

The set of actions for type  $\theta_i$  is restricted as  $A_{i,\theta_i} \subseteq A_i$  for each  $i \in N$  and  $\theta_i \in \Theta_i$ . The set of strategies  $S_i \subseteq A_i^{\Theta_i}$  is also restricted so that  $s_i(\theta_i) \in A_{i,\theta_i}$  holds for each  $i \in N$ ,  $\theta_i \in \Theta_i$ , and  $s_i \in S_i$ . The utility function is defined as  $v_i(\theta; \cdot): \prod_{i \in N} A_{i,\theta_i} \rightarrow [0, 1]$  for each  $i \in N$  and  $\theta \in \Theta$ , but here we still assume that  $v_i(\theta; a) = v_i(\theta'; a)$  for any  $a \in \prod_{i \in N} A_{i,\theta_i} \cap \prod_{i \in N} A_{i,\theta'_i}$  if  $\theta_i = \theta'_i$ . The definition of approximate communication equilibria is generalized as follows.

**Definition 5.9** ( $\epsilon$ -Approximate communication equilibria for type-dependent action sets). For any  $\epsilon \geq 0$ , a type-wise distribution  $\pi \in \Delta(A)^\Theta$  such that the support of  $\pi(\theta)$  is a subset of  $\prod_{i \in N} A_{i,\theta_i}$  for each  $\theta \in \Theta$  is an  $\epsilon$ -approximate communication equilibrium if for any  $i \in N$ ,  $\psi: \Theta_i \rightarrow \Theta_i$ , and  $\phi(\theta_i, \cdot): A_{i,\psi(\theta_i)} \rightarrow A_{i,\theta_i}$  for each  $\theta_i \in \Theta_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\psi(\theta_i), \theta_{-i})} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \epsilon.$$

The definition of strategy representability is also generalized with restricted strategy profiles  $S = \prod_{i \in N} S_i$ . Our proposed dynamics and algorithm for minimizing untruthful swap regret are naturally extended to this setting.

Roughgarden [2015b] and Syrgkanis [2012] assumed *universal smoothness* of Bayesian games with type-dependent action sets. This condition holds in routing games<sup>14</sup> [Roughgarden, 2015b, Syrgkanis, 2012] and Bayesian effort games [Syrgkanis, 2012]. Here we consider its conditional version, which is an assumption weaker than universal smoothness.

**Definition 5.10** ( $(\lambda, \mu)$ -Universal conditional smoothness of Bayesian games). A Bayesian game is  $(\lambda, \mu)$ -universally conditionally smooth if there exists some  $a_{i,\theta,a_i}^* \in A_{i,\theta_i}$  for each  $i \in N$ ,  $\theta \in \Theta$ , and  $a_i \in A_i$  such that

$$\sum_{i \in N} v_i(\theta; a_{i,\theta,a_i}^*, a_{-i}) \geq \lambda \max_{a' \in \prod_{i \in N} A_{i,\theta_i}} v_{\text{SW}}(\theta; a') - \mu v_{\text{SW}}(\theta'; a).$$

<sup>14</sup>Precisely speaking, we need to consider the *minimization* version in which each player minimizes their cost instead of maximizing their payoff. A similar result can be obtained by the same proof.

holds for any  $\theta, \theta' \in \Theta$  and  $a \in A$  with  $a_i \in A_{i, \theta'_i}$  for each  $i \in N$ .

We can generalize Theorem 5.2 to this setting and apply this smoothness inequality to its summation over  $i \in N$ . We also assume that the social welfare is *payoff-dominating*, i.e.,  $v_{\text{SW}}(\theta; a) \geq \sum_{i \in N} v_i(\theta; a)$  for each  $\theta \in \Theta$  and  $a \in A$  such that  $a_i \in A_{i, \theta_i}$  for each  $i \in N$ . Then we obtain

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta)} [v_{\text{SW}}(\theta; a)] \right] \geq \lambda \mathbb{E}_{\theta \sim \rho} \left[ \max_{a' \in \prod_{i \in N} A_{i, \theta'_i}} v_{\text{SW}}(\theta; a') \right] - \mu \mathbb{E}_{\theta' \sim \rho} \left[ \mathbb{E}_{a \sim \pi(\theta')} [v_{\text{SW}}(\theta'; a)] \right],$$

which leads a  $\lambda/(1 + \mu)$  lower bound on the POA of communication equilibria with strategy representability.

## 6 Concluding remarks

We have studied the classes of Bayes correlated equilibria from the viewpoint of efficient computability and the price of anarchy. We proposed dynamics converging to communication equilibria with strategy representability and showed how to simulate the dynamics efficiently by designing an algorithm for minimizing untruthful swap regret. We also proved a lower bound on untruthful swap regret that matches the upper bound up to a multiplicative constant. The price of anarchy of communication equilibria with strategy representability was analyzed via the smoothness argument that has been applied to Bayes Nash equilibria.

The problem of efficiently computing an approximate strategic-form correlated equilibrium is left unsolved. An approach to this problem is to design an efficient algorithm for minimizing strategy swap regret. Although the number of possible swaps  $\phi_{\text{SF}}: S_i \rightarrow S_i$  is doubly exponential, we can reduce this number by considering the simplicity of the intrinsic decision space  $\mathcal{X}$ . In the opposite direction, it might be possible to prove a superpolynomial lower bound on the query complexity. It is still difficult for us to judge which direction is more promising.

Another possible direction for future work is to investigate relations between communication equilibria and Bayesian mechanism design. The only difference between them is whether a mediator sends recommendations to players or determines an outcome in a top-down fashion. This combination may bring new possibilities for Bayesian mechanism design. For example, we proved a POA lower bound for smooth Bayesian mechanisms, but its consequence is still unexplored. A communication equilibrium of a mechanism provides another mechanism with a truth-telling equilibrium, but we do not know its appropriate interpretation.

## Acknowledgements

This work was supported by JSPS KAKENHI Grant Number 22K17857.

## References

- Peter Auer, Nicolò Cesa-Bianchi, and Claudio Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and Systems Sciences*, 64(1):48–75, 2002.
- Robert Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

- Yakov Babichenko. Query complexity of approximate Nash equilibria. *Journal of the ACM*, 63(4), 2016.
- Yakov Babichenko and Siddharth Barman. Algorithmic aspects of private Bayesian persuasion. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, pages 34:1–34:16, 2017.
- Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.
- Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, 2007.
- Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC 2008)*, pages 373–382, 2008.
- Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos. Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory*, 156:343–388, 2015.
- Andrea Celli and Nicola Gatti. Computational results for extensive-form adversarial team games. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2018)*, volume 32, 2018.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P. Helmbold, Robert E. Schapire, and Manfred K. Warmuth. How to use expert advice. *Journal of the ACM*, 44(3):427–485, 1997.
- Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):14:1–14:57, 2009.
- George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. *Journal of the ACM*, 63(2):11:1–11:19, 2016.
- 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.
- Shaddin Dughmi and Haifeng Xu. Algorithmic Bayesian persuasion. *SIAM Journal on Computing*, 50(3), 2021.
- Gabriele Farina, Tommaso Bianchi, and Tuomas Sandholm. Coarse correlation in extensive-form games. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, pages 1934–1941, 2020.
- 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):41:1–41:41, 2022.

- Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions without complements are (almost) efficient. *Games and Economic Behavior*, 123:327–341, 2020.
- Françoise Forges. An approach to communication equilibria. *Econometrica*, pages 1375–1385, 1986.
- Françoise Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35:277–310, 1993.
- Françoise Forges. Correlated equilibrium in games with incomplete information revisited. *Theory and Decision*, 61:329–344, 2006.
- Françoise Forges. Correlated equilibria and communication in games. In Robert A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 1–13. Springer Berlin, Heidelberg, 2014.
- Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, 1997.
- Anat Ganor and Karthik C. S. Communication complexity of correlated equilibrium with small support. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*, pages 12:1–12:16, 2018.
- Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- Geoffrey J. Gordon, Amy Greenwald, and Casey Marks. No-regret learning in convex games. In *Proceedings of the Twenty-Fifth International Conference on Machine Learning (ICML 2008)*, pages 360–367, 2008.
- Amy Greenwald and Amir Jafari. A general class of no-regret learning algorithms and game-theoretic equilibria. In *Proceedings of the 16th Annual Conference on Computational Learning Theory and 7th Kernel Workshop (COLT/Kernel 2003)*, pages 2–12, 2003.
- John C. Harsanyi. Games with incomplete information played by “Bayesian” players, I–III part I. The basic model. *Management Science*, 14(3):159–182, 1967.
- John C. Harsanyi. Games with incomplete information played by “Bayesian” players part II. Bayesian equilibrium points. *Management Science*, 14(5):320–334, 1968a.
- John C. Harsanyi. Games with incomplete information played by ‘Bayesian’ players, part III. The basic probability distribution of the game. *Management Science*, 14(7):486–502, 1968b.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Sergiu Hart and Noam Nisan. The query complexity of correlated equilibria. *Games and Economic Behavior*, 108:401–410, 2018.
- Jason D. Hartline. Bayesian mechanism design. *Foundations and Trends<sup>®</sup> in Theoretical Computer Science*, 8(3):143–263, 2013.

- Jason D. Hartline, Darrell Hoy, and Sam Taggart. Price of anarchy for auction revenue. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC 2014)*, pages 693–710, 2014.
- Jason D. Hartline, Vasilis Syrgkanis, and Éva Tardos. No-regret learning in Bayesian games. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 3061–3069, 2015.
- 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 (NIPS 2007)*, pages 625–632, 2007.
- Shinji Ito. A tight lower bound and efficient reduction for swap regret. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, 2020.
- Albert Xin Jiang and Kevin Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior*, 91:347–359, 2015.
- Yaonan Jin and Pinyan Lu. First price auction is  $1 - 1/e^2$  efficient. In *Proceedings of the 63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS 2022)*, pages 179–187, 2022.
- Yaonan Jin and Pinyan Lu. The price of stability for first price auction. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA 2023)*, pages 332–352, 2023.
- Adam Tauman Kalai and Santosh S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Wouter M. Koolen, Manfred K. Warmuth, and Jyrki Kivinen. Hedging structured concepts. In *Proceedings of the 23rd Conference on Learning Theory (COLT 2010)*, pages 93–105, 2010.
- Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS 1999)*, pages 404–413, 1999.
- Ehud Lehrer, Dinah Rosenberg, and Eran Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2):670–682, 2010.
- Renato Paes Leme and Éva Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2010)*, pages 735–744, 2010.
- Yishay Mansour, Mehryar Mohri, Jon Schneider, and Balasubramanian Sivan. Strategizing against learners in Bayesian games. In *Proceedings of the 35th Conference on Learning Theory (COLT 2022)*, pages 5221–5252, 2022.
- Roger B Myerson. Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics*, 10(1):67–81, 1982.

- Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3):14:1–14:29, 2008.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Beyond regret. In *Proceedings of the 24th Annual Conference on Learning Theory (COLT 2011)*, volume 19, pages 559–594, 2011.
- Tim Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):32:1–32:42, 2015a.
- Tim Roughgarden. The price of anarchy in games of incomplete information. *ACM Transactions on Economics and Computation*, 3(1):6:1–6:20, 2015b.
- Aviad Rubinstein. Inapproximability of Nash equilibrium. *SIAM Journal on Computing*, 47(3):917–959, 2018.
- 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. Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155, 2012.
- Vasilis Syrgkanis and Éva Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC 2013)*, pages 211–220, 2013.
- Bernhard von Stengel and Françoise Forges. Extensive-form correlated equilibrium: Definition and computational complexity. *Mathematics of Operations Research*, 33(4):1002–1022, 2008.
- Brian Hu Zhang and Tuomas Sandholm. Polynomial-time optimal equilibria with a mediator in extensive-form games. *CoRR*, abs/2206.15395, 2022.
- Chenghan Zhou, Thanh H. Nguyen, and Haifeng Xu. Algorithmic information design in multi-player games: Possibilities and limits in singleton congestion. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC 2022)*, page 869, 2022.

## A On conditional independence property

In this section, we show the difference between the strategy representability and the conditional independence property. The conditional independence property was mentioned by Forges [1993] as an important property that holds for SFCEs and ANFCEs.

**Definition A.1** (Conditional independence property). Let  $\pi \in \Delta(A)^\Theta$ . We consider a joint distribution in  $\Delta(\Theta \times A)$  that first generates  $\theta \in \Theta$  according to  $\rho$  and then generates  $a \in A$  according to  $\pi(\theta)$ . We say  $\pi$  satisfies the *conditional independence property*<sup>15</sup> if  $a_i$  is conditionally independent of  $\theta_{-i}$  given  $\theta_i$  for each  $i \in N$ .

---

<sup>15</sup>It is originally defined as a property of a joint distribution in  $\Delta(\Theta \times A)$ , but we can consider a joint distribution in  $\Delta(\Theta \times A)$  and a type-wise distribution in  $\Delta(A)^\Theta$  interchangeably. For each  $\pi \in \Delta(A)^\Theta$ , the corresponding joint distribution in  $\Delta(\Theta \times A)$  is uniquely determined. Conversely, for any joint distribution whose marginal for  $\theta$  coincides with  $\rho$ , we can uniquely determine the corresponding distribution in  $\Delta(A)^\Theta$ .

As Forges [1993] claimed, this property holds for any strategy-representable distribution. Although it was not claimed explicitly, this property is not equivalent to the strategy representability.

**Proposition A.2.** *If  $\pi \in \Delta(A)^\Theta$  is strategy-representable, then  $\pi$  satisfies the conditional independence property. On the other hand, there exists  $\pi \in \Delta(A)^\Theta$  that satisfies the conditional independence property but not the strategy representability.*

*Proof.* First, we show that the strategy representability implies the conditional independence property. Assume there exists  $\sigma \in \Delta(S)$  such that  $\pi(\theta; a) = \Pr_{s \sim \sigma}(s(\theta) = a)$  for each  $\theta \in \Theta$  and  $a \in A$ . To prove the conditional independence property, it is sufficient to prove the conditional independence of  $a_i$  and  $\theta_{-i}$  given  $\theta_i$  for any  $i \in N$ ,  $\theta \in \Theta$ , and  $a_i \in A_i$ . Using the strategy representability of  $\pi$ , we obtain

$$\begin{aligned}
& \Pr_{\substack{\theta' \sim \rho \\ a' \sim \pi(\theta')}} (\theta' = \theta, a'_i = a_i) \Pr_{\theta' \sim \rho} (\theta'_i = \theta_i) \\
&= \left( \rho(\theta) \Pr_{a' \sim \pi(\theta)} (a'_i = a_i) \right) \Pr_{\theta' \sim \rho} (\theta'_i = \theta_i) \\
&= \rho(\theta) \Pr_{s' \sim \sigma} (s'_i(\theta_i) = a_i) \Pr_{\theta' \sim \rho} (\theta'_i = \theta_i) && \text{(from the definition of } \sigma) \\
&= \rho(\theta) \Pr_{\substack{\theta' \sim \rho \\ s' \sim \sigma}} (s'_i(\theta_i) = a_i, \theta'_i = \theta_i) \\
&= \rho(\theta) \Pr_{\substack{\theta' \sim \rho \\ s' \sim \sigma}} (s'_i(\theta'_i) = a_i, \theta'_i = \theta_i) && \text{(since } s'_i(\theta'_i) \text{ can be replaced by } s'_i(\theta_i) \text{ when } \theta'_i = \theta_i) \\
&= \rho(\theta) \Pr_{\substack{\theta' \sim \rho \\ a' \sim \pi(\theta')}} (a'_i = a_i, \theta'_i = \theta_i). && \text{(from the definition of } \sigma)
\end{aligned}$$

Dividing both sides by  $(\Pr_{\theta' \sim \rho}(\theta'_i = \theta_i))^2$ , we obtain the conditional independence of  $a_i$  and  $\theta_{-i}$  given  $\theta_i$ :

$$\Pr_{\substack{\theta' \sim \rho \\ a' \sim \pi(\theta')}} (\theta'_{-i} = \theta_{-i}, a'_i = a_i \mid \theta'_i = \theta_i) = \Pr_{\substack{\theta' \sim \rho \\ a' \sim \pi(\theta')}} (\theta'_{-i} = \theta_{-i} \mid \theta'_i = \theta_i) \Pr_{\substack{\theta' \sim \rho \\ a' \sim \pi(\theta')}} (a_i = a_i \mid \theta'_i = \theta_i).$$

Next, we show the existence of an example that satisfies the conditional independence property but not the strategy representability. We consider  $\pi \in \Delta(A)^\Theta$  described in Theorem 2.2 again (see Figure 3). Since we already proved that this  $\pi$  is not strategy-representable, it is sufficient to show that  $\pi$  satisfies the conditional independence property. It can be verified by considering the distribution conditional on each player's type. If the first player's type is  $\theta_1$ , the conditional distribution of the first player's action and the second player's type is the uniform distribution on  $A_1 \times \Theta_2$ , hence conditionally independent. This is also the case for  $\theta'_1$ . Since the definition of  $\pi$  is symmetric for two players, we can conclude that  $\pi$  satisfies the conditional independence property.  $\square$

## B Algorithm for minimizing strategy swap regret

Here we propose an algorithm for minimizing strategy swap regret  $R_{SS,i}$  for an online learning problem with stochastic types.

---

**Algorithm 4:** Algorithm for minimizing strategy swap regret
 

---

**Input:** The set of types  $\Theta_i$  and the set of actions  $A_i$  are specified in advance. The reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  is given at the end of each round  $t \in [T]$ .

Let  $\mathcal{E}_{s_i, \theta_i}$  be a multiplicative weights algorithm with the doubling trick [Cesa-Bianchi et al., 1997, Section 4.6] with decision space  $S_i$  for each  $\theta_i \in \Theta_i$  and  $s_i \in S_i$ ;

**for** each round  $t = 1, \dots, T$  **do**

Let  $z_{s_i, \theta_i}^t \in \Delta(A_i)$  be the output of  $\mathcal{E}_{s_i, \theta_i}$  in round  $t$  for each  $s_i \in S_i$  and  $\theta_i \in \Theta_i$ ;

Define  $P^t \in [0, 1]^{S_i \times S_i}$  by  $P^t(s_i, s'_i) = \prod_{\theta_i \in \Theta_i} z_{s'_i, \theta_i}^t(s_i(\theta_i))$  for each  $s_i, s'_i \in S_i$ ;

Compute  $\sigma_i^t \in \Delta(S_i)$  such that  $P^t \sigma_i^t = \sigma_i^t$  by eigenvector computation;

Let  $\sigma_i^t$  be the output for round  $t$ ;

Observe reward vector  $u_i^t \in [0, 1]^{\Theta_i \times A_i}$  and feed the reward vector  $\tilde{u}_{s_i, \theta_i}^t \in [0, 1]^{A_i}$  defined by  $\tilde{u}_{s_i, \theta_i}^t(a_i) = \sigma_i^t(s_i) u_i^t(\theta_i, a_i)$  for each  $a_i \in A_i$  to subroutine  $\mathcal{E}_{s_i, \theta_i}$  for each  $s_i \in S_i, \theta_i \in \Theta_i$ ;

**end**

---

Since strategy swap regret  $R_{\text{SS}, i}$  is swap regret when  $S_i$  is regarded as the decision space, we can apply the swap regret minimization algorithm proposed by Blum and Mansour [2007]. The direct application leads to an upper bound of  $O(\sqrt{T|S_i| \log |S_i|} + |S_i| \log |S_i|)$  on the strategy swap regret.

Their algorithm reduces swap regret minimization to  $|S_i|$  external regret minimization problems with decision space  $S_i$ . In the case of strategy swap regret, we can further reduce them to  $|S_i| |\Theta_i|$  external regret minimization problems with decision space  $A_i$ . As a result, we obtain an upper bound of  $O(\sqrt{T|S_i| \log |A_i|} + |S_i| \log |A_i|)$  on the strategy swap regret, which is slightly better than the above bound. Since  $|S_i| = |A_i|^{|\Theta_i|}$ , this bound is still exponentially large in the number of types  $|\Theta_i|$ . Moreover, the algorithm uses exponentially many subroutines, which require exponential time computation for each round. It is open whether there is an efficient algorithm with strategy swap regret sublinear in  $T$  and polynomial in  $n, |A_i|$ , and  $|\Theta_i|$ . If it exists, we can compute an  $\epsilon$ -SFCE in polynomial time by simulating the dynamics.

The reduction proceeds as follows. Let  $\mathcal{E}_{s_i, \theta_i}$  be the multiplicative weights algorithm with the doubling trick [Cesa-Bianchi et al., 1997, Section 4.6] with decision space  $A_i$  for each  $\theta_i \in \Theta_i$  and  $s_i \in S_i$ . In each round  $t \in [T]$ , each subroutine  $\mathcal{E}_{s_i, \theta_i}$  outputs  $z_{s_i, \theta_i}^t \in \Delta(A_i)$ . From these outputs, we define the stochastic matrix  $P^t \in [0, 1]^{S_i \times S_i}$  by  $P^t(s_i, s'_i) = \prod_{\theta_i \in \Theta_i} z_{s'_i, \theta_i}^t(s_i(\theta_i))$  for each  $s_i, s'_i \in S_i$ . Since each  $P^t \in \mathcal{P}$  is a stochastic matrix, we can compute its stationary distribution  $\sigma_i^t \in \Delta(S_i)$  that satisfies  $P^t \sigma_i^t = \sigma_i^t$ , and let  $\sigma_i^t$  be the decision for round  $t$ . Then feed the reward vector  $\tilde{u}_{s_i, \theta_i}^t \in [0, 1]^{A_i}$  defined by  $\tilde{u}_{s_i, \theta_i}^t(a_i) = \sigma_i^t(s_i) u_i^t(\theta_i, a_i)$  for each  $a_i \in A_i$  to subroutine  $\mathcal{E}_{s_i, \theta_i}$  for each  $s_i \in S_i, \theta_i \in \Theta_i$ . Let

$$R_{s_i, \theta_i}^T = \max_{a_i \in A_i} \sum_{t=1}^T \tilde{u}_{s_i, \theta_i}^t(a_i) - \sum_{t=1}^T \sum_{a_i \in A_i} z_{s_i, \theta_i}^t(a_i) \tilde{u}_{s_i, \theta_i}^t(a_i)$$

be the external regret of  $\mathcal{E}_{s_i, \theta_i}$  for each  $s_i \in S_i$  and  $\theta_i \in \Theta_i$ . Then strategy swap regret  $R_{\text{SS}, i}^T$  equals the sum of the external regrets as follows.

**Lemma B.1.**

$$R_{\text{SS}, i}^T = \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in S_i} R_{s_i, \theta_i}^T.$$

*Proof.* Recall that the strategy swap regret is defined as

$$R_{\text{SS},i}^T = \max_{\phi_{\text{SF}}: S_i \rightarrow S_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, (\phi_{\text{SF}}(s_i))(\theta_i))] \right] - \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, s_i(\theta_i))] \right].$$

We analyze the first term and the second term separately. The first term is

$$\begin{aligned} & \max_{\phi_{\text{SF}}: S_i \rightarrow S_i} \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, (\phi_{\text{SF}}(s_i))(\theta_i))] \right] \\ &= \max_{\phi_{\text{SF}}: S_i \rightarrow S_i} \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in S_i} \sigma_i^t(s_i) u_i^t(\theta_i, (\phi_{\text{SF}}(s_i))(\theta_i)) \\ &= \max_{\phi_{\text{SF}}: S_i \rightarrow S_i} \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in S_i} \tilde{u}_{s_i, \theta_i}^t((\phi_{\text{SF}}(s_i))(\theta_i)) \quad (\text{from the definition of } \tilde{u}_{s_i, \theta_i}^t) \\ &= \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in S_i} \max_{a_i \in A_i} \sum_{t=1}^T \tilde{u}_{s_i, \theta_i}^t(a_i), \end{aligned}$$

where the last equality holds because  $(\phi_{\text{SF}}(s_i))(\theta_i)$  can be optimized separately for each  $s_i \in S_i$  and  $\theta_i \in \Theta_i$ .

The second term is

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{\theta_i \sim \rho_i} \left[ \mathbb{E}_{s_i \sim \sigma_i^t} [u_i^t(\theta_i, s_i(\theta_i))] \right] \\ &= \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in S_i} \sigma_i^t(s_i) u_i^t(\theta_i, s_i(\theta_i)) \\ &= \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i, s'_i \in S_i} P^t(s_i, s'_i) \sigma_i^t(s'_i) u_i^t(\theta_i, s_i(\theta_i)) \quad (\text{since } P^t \sigma_i^t = \sigma_i^t) \\ &= \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i, s'_i \in S_i} \sigma_i^t(s'_i) \prod_{\theta'_i \in \Theta_i} z_{s'_i, \theta'_i}^t(s_i(\theta'_i)) u_i^t(\theta_i, s_i(\theta_i)) \quad (\text{from the definition of } P^t) \\ &= \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i, s'_i \in S_i} \prod_{\theta'_i \in \Theta_i} z_{s'_i, \theta'_i}^t(s_i(\theta'_i)) \tilde{u}_{s'_i, \theta'_i}^t(s_i(\theta_i)) \quad (\text{from the definition of } \tilde{u}_{s_i, \theta_i}^t) \\ &= \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s'_i \in S_i} \sum_{a_i \in A_i} \left( \sum_{\substack{s_i \in S_i: \\ s_i(\theta_i) = a_i}} \prod_{\theta'_i \in \Theta_i} z_{s'_i, \theta'_i}^t(s_i(\theta'_i)) \right) \tilde{u}_{s'_i, \theta'_i}^t(a_i) \\ &= \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s'_i \in S_i} \sum_{a_i \in A_i} z_{s'_i, \theta'_i}^t(a_i) \tilde{u}_{s'_i, \theta'_i}^t(a_i). \end{aligned}$$

Finally, we obtain

$$\begin{aligned}
R_{\text{SS},i}^T &= \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in \mathcal{S}_i} \max_{a_i \in A_i} \sum_{t=1}^T \tilde{u}_{s_i, \theta_i}^t(a_i) - \sum_{t=1}^T \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in \mathcal{S}_i} \sum_{a_i \in A_i} z_{s_i, \theta_i}^t(a_i) \tilde{u}_{s_i, \theta_i}^t(a_i) \\
&= \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in \mathcal{S}_i} \left\{ \max_{a_i \in A_i} \sum_{t=1}^T \tilde{u}_{s_i, \theta_i}^t(a_i) - \sum_{t=1}^T \sum_{a_i \in A_i} z_{s_i, \theta_i}^t(a_i) \tilde{u}_{s_i, \theta_i}^t(a_i) \right\} \\
&= \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in \mathcal{S}_i} R_{s_i, \theta_i}^T. \quad \square
\end{aligned}$$

By applying Theorem 3.7 (an upper bound for the multiplicative weights with the doubling trick), we obtain an upper bound on strategy swap regret.

**Theorem B.2.**

$$R_{\text{SS},i} \leq 6\sqrt{T|\mathcal{S}_i| \log |A_i|} + 2|\mathcal{S}_i| \log |A_i|.$$

*Proof.* Since the total reward of each action  $a_i \in A_i$  for  $\mathcal{E}_{s_i, \theta_i}$  is bounded above as  $\sum_{t=1}^T \tilde{u}_{s_i, \theta_i}^t(a_i) = \sum_{t=1}^T \sigma_i^t(s_i) u_i^t(\theta_i, a_i) \leq \sum_{t=1}^T \sigma_i^t(s_i)$ , we can obtain an upper bound on the external regret

$$R_{s_i, \theta_i}^T \leq 6\sqrt{\sum_{t=1}^T \sigma_i^t(s_i) \log |A_i|} + 2 \log |A_i|.$$

By summing this upper bound for all  $\theta_i \in \Theta_i$  and  $s_i \in \mathcal{S}_i$ , we obtain

$$\begin{aligned}
R_{\text{SS},i}^T &= \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in \mathcal{S}_i} R_{s_i, \theta_i}^T \\
&\leq \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \sum_{s_i \in \mathcal{S}_i} \left\{ 6\sqrt{\sum_{t=1}^T \sigma_i^t(s_i) \log |A_i|} + 2 \log |A_i| \right\} \\
&\leq \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \left\{ 6\sqrt{|\mathcal{S}_i| \sum_{s_i \in \mathcal{S}_i} \sum_{t=1}^T \sigma_i^t(s_i) \log |A_i|} + 2|\mathcal{S}_i| \log |A_i| \right\} \\
&\hspace{15em} \text{(by the Cauchy-Schwarz inequality)} \\
&= \sum_{\theta_i \in \Theta_i} \rho_i(\theta_i) \left\{ 6\sqrt{T|\mathcal{S}_i| \log |A_i|} + 2|\mathcal{S}_i| \log |A_i| \right\} \quad \text{(since } \sum_{s_i \in \mathcal{S}_i} \sigma_i^t(s_i) = 1 \text{ for each } t \in [T]) \\
&= 6\sqrt{T|\mathcal{S}_i| \log |A_i|} + 2|\mathcal{S}_i| \log |A_i|. \quad \square
\end{aligned}$$

## C On linear swap regret minimization

Mansour et al. [2022] defined linear swap regret for online linear optimization with a polytope constraint as follows.

**Definition C.1** (Linear swap regret [Mansour et al., 2022]). Let  $\mathcal{P} \subseteq \mathbb{R}^d$  be a polytope and  $\mathcal{M}(\mathcal{P})$  be the set of all valid linear transformations, where a linear transformation  $M: \mathbb{R}^d \rightarrow \mathbb{R}^d$  is defined

to be valid when  $Mx \in \mathcal{P}$  holds for all  $x \in \mathcal{P}$ . For online linear optimization with reward vector  $\bar{u}^t \in [0, 1]^d$  for each  $t \in [T]$  and feasible region  $\mathcal{P}$ , linear swap regret is defined as

$$R_{\text{LS}}^T = \max_{M \in \mathcal{M}(\mathcal{P})} \sum_{t=1}^T (\langle Mx^t, \bar{u}^t \rangle - \langle x^t, \bar{u}^t \rangle),$$

where  $x^t \in \mathcal{P}$  is the algorithm's output in each round  $t \in [T]$ .

Mansour et al. [2022] defined linear swap regret for a general polytope but focused on a special case of  $\mathcal{P} = \mathcal{X}$  with applications to two-player Bayesian games. Recall that  $\mathcal{X} = \{x \in [0, 1]^{\Theta_i \times A_i} \mid \sum_{a_i \in A_i} x(\theta_i, a_i) = 1 \ (\forall \theta_i \in \Theta_i)\}$  is the set of vectors that represent each  $\pi_i \in \Delta(A_i)^{\Theta_i}$ .

Here we show that linear swap regret minimization with  $\mathcal{P} = \mathcal{X}$  can be reduced to untruthful swap regret minimization. Note that they propose an algorithm for minimizing linear swap regret in this special case (Algorithm 2 in their paper), but focused on guarantees on the Stackelberg value and did not provide any rigorous upper bound on linear swap regret.

The following proposition claims that all valid linear transformations can be expressed by some  $Q \in \mathcal{Q}$ , which represents a linear transformation for untruthful swap regret. Recall that  $\mathcal{Q}$  is defined by (6).

**Proposition C.2.** *For any  $M \in \mathcal{M}(\mathcal{X})$ , there exists some  $Q \in \mathcal{Q}$  such that  $Mx = Qx$  holds for any  $x \in \mathcal{X}$ .*

*Proof.* Fix any  $M \in \mathcal{M}(\mathcal{X})$ . We construct  $Q \in \mathcal{Q}$  such that  $Mx = Qx$  holds for any  $x \in \mathcal{X}$ .

Since  $M$  is a valid linear transformation,  $Mx \in \mathcal{X}$  holds for any  $x \in \mathcal{X}$ . We use this fact for  $x^1, x^2 \in \mathcal{X}$  defined as follows. Fix any  $\theta'_i \in \Theta_i$  and  $a_i^1, a_i^2 \in A_i$ . Define  $x^1(\theta'_i, a_i^1) = 1$  and  $x(\theta'_i, a_i) = 0$  for any other  $a_i \in A_i \setminus \{a_i^1\}$ . Similarly, let  $x^2(\theta'_i, a_i^2) = 1$  and  $x^2(\theta'_i, a_i) = 0$  for any other  $a_i \in A_i \setminus \{a_i^2\}$ . The other entries of  $x^1$  can be defined arbitrarily, and the other entries of  $x^2$  are set to be equal to  $x^1$ . That is, for any  $\theta_i \in \Theta_i \setminus \{\theta'_i\}$ , we must have  $x^1(\theta_i, a_i) = x^2(\theta_i, a_i)$  for any  $a_i \in A_i$ .

Now we use  $Mx^1 \in \mathcal{X}$  and  $Mx^2 \in \mathcal{X}$ . For each  $\theta_i \in \Theta_i$ , we have

$$\sum_{a_i \in A_i} (Mx^1)(\theta_i, a_i) = 1 \quad \text{and} \quad \sum_{a_i \in A_i} (Mx^2)(\theta_i, a_i) = 1,$$

which implies

$$\sum_{a_i \in A_i} (M(x^1 - x^2))(\theta_i, a_i) = 0.$$

Since  $(x^1 - x^2)(\theta'_i, a_i^1) = 1$ ,  $(x^1 - x^2)(\theta'_i, a_i^2) = -1$ , and the other entries of  $x^1 - x^2$  are 0, we obtain

$$\sum_{a_i \in A_i} M((\theta_i, a_i), (\theta'_i, a_i^1)) = \sum_{a_i \in A_i} M((\theta_i, a_i), (\theta'_i, a_i^2)).$$

Since this holds for any pair of  $a_i^1, a_i^2 \in A_i$ , there exists some value  $W(\theta_i, \theta'_i) \in \mathbb{R}$  for each  $\theta_i, \theta'_i \in \Theta_i$  such that

$$\sum_{a_i \in A_i} M((\theta_i, a_i), (\theta'_i, a_i^1)) = W(\theta_i, \theta'_i)$$

for any  $a_i \in A_i$ . Moreover, by using  $Mx \in \mathcal{X}$  for each  $x \in \mathcal{X}$ , it holds that for any  $\theta_i \in \Theta_i$ ,

$$\begin{aligned}
1 &= \sum_{a_i \in A_i} (Mx)(\theta_i, a_i) \\
&= \sum_{a_i \in A_i} \sum_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} x(\theta'_i, a'_i) M((\theta_i, a_i), (\theta'_i, a'_i)) \\
&= \sum_{\theta'_i \in \Theta_i} \sum_{a'_i \in A_i} x(\theta'_i, a'_i) \left\{ \sum_{a_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i)) \right\} \\
&= \sum_{\theta'_i \in \Theta_i} \left\{ \sum_{a'_i \in A_i} x(\theta'_i, a'_i) \right\} W(\theta_i, \theta'_i) && \text{(from the definition of } W(\theta_i, \theta'_i)) \\
&= \sum_{\theta'_i \in \Theta_i} W(\theta_i, \theta'_i). && \text{(since } x \in \mathcal{X})
\end{aligned}$$

If every entry of  $M$  is included in  $[0, 1]$ , then the proof is finished, but this does not necessarily hold. We show that we can achieve  $M \in \mathcal{Q}$  by shifting each entry of  $M$  without changing  $Mx$  for every  $x \in \mathcal{X}$ . Fix any  $\theta_i \in \Theta_i$  and  $a_i \in A_i$ , and we focus on the row  $M((\theta_i, a_i), (\cdot, \cdot))$ . For any  $\theta'_i \in \Theta_i$ , we define  $a_i^1 \in \operatorname{argmax}_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i))$  and  $a_i^2 \in \operatorname{argmin}_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i))$ . We reuse the above definition of  $x^1$  and  $x^2$  with these new  $a_i^1$  and  $a_i^2$ . Since  $(Mx^1)(\theta_i, a_i)$  and  $(Mx^2)(\theta_i, a_i)$  are included in  $[0, 1]$ , we have

$$\max_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i)) - \min_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i)) \in [0, 1]. \quad (30)$$

If we set  $x(\theta'_i, a_i) = 1$  for some  $a_i \in \operatorname{argmax}_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i))$  for each  $\theta_i$  and set 0 otherwise,  $(Mx)(\theta_i, a_i) \in [0, 1]$  implies

$$\sum_{\theta'_i \in \Theta_i} \max_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i)) \in [0, 1]. \quad (31)$$

Similarly, we can obtain

$$\sum_{\theta'_i \in \Theta_i} \min_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i)) \in [0, 1]. \quad (32)$$

For any  $\theta'_i, \theta''_i \in \Theta_i$ , even if we add an arbitrary value  $C \in \mathbb{R}$  to all entries  $M((\theta_i, a_i), (\theta'_i, \cdot))$  and subtract  $C$  from all entries  $M((\theta_i, a_i), (\theta''_i, \cdot))$ , the value of  $(Mx)(\theta_i, a_i)$  does not change. If we denote this shifted matrix by  $M'$ , this fact can be checked as

$$(M'x)(\theta_i, a_i) = (Mx)(\theta_i, a_i) + \sum_{a'_i \in A_i} Cx(\theta'_i, a'_i) - \sum_{a'_i \in A_i} Cx(\theta''_i, a'_i) = (Mx)(\theta_i, a_i).$$

Now we consider shifting the entries of  $M$  and obtain  $M'$ . Fix any  $\theta_i^* \in \Theta_i$ . For  $\theta'_i \in \Theta_i \setminus \{\theta_i^*\}$ , we define all entries  $M'((\theta_i, a_i), (\theta'_i, \cdot))$  by

$$M'((\theta_i, a_i), (\theta'_i, a''_i)) = M((\theta_i, a_i), (\theta'_i, a''_i)) - \min_{a'_i \in A_i} M((\theta_i, a_i), (\theta'_i, a'_i))$$

for each  $a_i'' \in A_i$ . For  $\theta_i^*$ , we define

$$M'((\theta_i, a_i), (\theta_i^*, a_i'')) = M((\theta_i, a_i), (\theta_i^*, a_i'')) + \sum_{\theta_i' \in \Theta_i \setminus \{\theta_i^*\}} \min_{a_i' \in A_i} M((\theta_i, a_i), (\theta_i', a_i'))$$

for each  $a_i'' \in A_i$ . Note that the sum of shifts is 0, and therefore,  $Mx = M'x$  for every  $x \in \mathcal{X}$ . We check that all the entries of  $M'$  on this row are included in  $[0, 1]$ . For  $\theta_i' \in \Theta_i \setminus \{\theta_i^*\}$ , all entries  $M'((\theta_i, a_i), (\theta_i', \cdot))$  are included in  $[0, 1]$  due to (30). For  $\theta_i^*$ , the largest entry is

$$\begin{aligned} \max_{a_i'' \in A_i} M'((\theta_i, a_i), (\theta_i^*, a_i'')) &= \max_{a_i' \in A_i} M((\theta_i, a_i), (\theta_i^*, a_i')) + \sum_{\theta_i' \in \Theta_i \setminus \{\theta_i^*\}} \min_{a_i' \in A_i} M((\theta_i, a_i), (\theta_i', a_i')) \\ &\leq \sum_{\theta_i' \in \Theta_i \setminus \{\theta_i^*\}} \max_{a_i' \in A_i} M((\theta_i, a_i), (\theta_i', a_i')) \in [0, 1] \end{aligned}$$

due to (31). The smallest entry is

$$\min_{a_i'' \in A_i} M'((\theta_i, a_i), (\theta_i^*, a_i'')) = \sum_{\theta_i' \in \Theta_i} \min_{a_i' \in A_i} M((\theta_i, a_i), (\theta_i', a_i')) \in [0, 1]$$

due to (32). Therefore, all entries of  $M'$  on this row are included in  $[0, 1]$ .

We can apply the same shifting operation to all rows of  $M$  and denote the obtained matrix by  $Q$ . All the entries of  $Q$  are included in  $[0, 1]$  with  $Qx = Mx$  for every  $x \in \mathcal{X}$ . From its definition,  $W(\theta_i, \theta_i')$  is the sum of non-negative values, hence non-negative for each  $\theta_i, \theta_i' \in \Theta_i$ . Since  $\sum_{\theta_i' \in \Theta_i} W(\theta_i, \theta_i') = 1$ , each  $W(\theta_i, \theta_i')$  is at most 1. Therefore,  $Q$  satisfies all the constraints for  $\mathcal{Q}$ , and we have  $Q \in \mathcal{Q}$ .  $\square$

By using this proposition, we can reduce linear swap regret minimization to untruthful swap regret minimization as

$$R_{\text{LS}}^T = \max_{M \in \mathcal{M}(\mathcal{P})} \sum_{t=1}^T (\langle Mx^t, \bar{u}^t \rangle - \langle x^t, \bar{u}^t \rangle) = \max_{Q \in \mathcal{Q}} \sum_{t=1}^T (\langle Qx^t, \bar{u}^t \rangle - \langle x^t, \bar{u}^t \rangle),$$

which equals  $R_{\text{US},i}^T$  from Theorem 3.1.

## D On Bayes coarse correlated equilibria

Here we briefly present definitions of Bayes coarse correlated equilibria for comparison with Bayes correlated equilibria. As with SFCEs, we can define coarse correlated equilibria of the strategic form. In an equilibrium of this class, each player  $i \in N$  does not have incentive to ignore recommendation  $s_i \in S_i$  and stick to any strategy  $s_i' \in S_i$ .

**Definition D.1** (Strategic-form coarse correlated equilibria (SFCCEs)). A distribution  $\sigma \in \Delta(S)$  is a strategic-form coarse correlated equilibrium if for any  $i \in N$  and any  $s_i' \in S_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbb{E}_{s \sim \sigma} [v_i(\theta; s_i'(\theta_i), s_{-i}(\theta_{-i}))] \right].$$

Table 2: A Bayesian game for which the set of SFCCEs and the set of ANFCCEs are different. The left table is the payoffs for player 1 with type  $\theta_1$ , and the right table is the payoffs for player 1 with type  $\theta'_1$ . The payoff for player 2 is always 0.

	$a_2$	$a'_2$
$a_1$	0	0
$a'_1$	0.5	0

	$a_2$	$a'_2$
$a_1$	0	1
$a'_1$	1	0

Similarly, we can define coarse correlated equilibria of the agent normal form. In an equilibrium of this class, each hypothetical player  $(i, \theta_i) \in N'$  does not have incentive to stick to any action  $a'_i \in A_i$ . Hartline et al. [2015] proposed dynamics converging to this class of equilibria.

**Definition D.2** (Agent-normal-form coarse correlated equilibria (ANFCCEs)). A distribution  $\sigma \in \Delta(S)$  is an agent-normal-form coarse correlated equilibrium if for any  $i \in N$ ,  $\theta'_i \in \Theta_i$ , and  $a'_i \in A_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbf{1}_{\{\theta_i = \theta'_i\}} \mathbb{E}_{s \sim \sigma} [v_i(\theta; s(\theta))] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbf{1}_{\{\theta_i = \theta'_i\}} \mathbb{E}_{s \sim \sigma} [v_i(\theta; a'_i, s_{-i}(\theta_{-i}))] \right].$$

The incentive constraints for ANFCCEs are imposed separately for each type  $\theta'_i$ . For (non-coarse) ANFCCEs, the separate constraints are equivalent to the total constraints ( $\text{IC}_{\text{ANF}}$ ) (for the non-approximate version, i.e.,  $\epsilon = 0$ ), and the difference between SFCEs and ANFCCEs comes from the difference between deviation  $\phi_{\text{SF}}: S_i \rightarrow S_i$  and  $\phi: \Theta_i \times A_i \rightarrow A_i$ . For SFCCEs and ANFCCEs, there is no difference in deviations, but the total constraints for SFCEs and the separate constraints for ANFCCEs are not equivalent. If we take the summation of the incentive constraint for ANFCCEs over type  $\theta'_i \in \Theta_i$ , we obtain those for SFCCEs by setting  $s'_i(\theta'_i) = a'_i$  for each  $\theta'_i \in \Theta_i$ . On the other hand, the incentive constraints for SFCCEs do not imply those for ANFCCEs. Thus, the set of ANFCCEs is a subset of the set of SFCCEs, which makes a striking contrast to the relation of SFCEs and ANFCCEs. A generalization of this relation for extensive-form games with imperfect information was proved by Farina et al. [2020]. Here we show this relation is strict even for Bayesian games by presenting a simple example.

**Example D.3.** We consider a Bayesian game with two players. There are two possible types for each player, and their types are completely correlated. More precisely, let  $\Theta_1 = \{\theta_1, \theta'_1\}$  and  $\Theta_2 = \{\theta_2, \theta'_2\}$ , and  $(\theta_1, \theta_2)$  realizes with probability 0.5 and  $(\theta'_1, \theta'_2)$  with probability 0.5. Each player has two possible actions  $A_1 = \{a_1, a'_1\}$  and  $A_2 = \{a_2, a'_2\}$ , respectively. The payoffs for player 1 with each type are presented in Table 2. The payoffs for player 2 are always defined to be 0.

We consider a distribution  $\sigma \in \Delta(S)$  that recommends two possible strategies  $s, s'$  each with probability 0.5. For  $(\theta_1, \theta_2)$ , these two strategies are the same: both of  $s_1$  and  $s_2$  recommend  $a_1$  to player 1 and  $a_2$  to player 2. For  $(\theta'_1, \theta'_2)$ , these two strategies are different:  $s_1(\theta'_1) = a'_1$  and  $s_2(\theta'_2) = a_2$ , while  $s'_1(\theta'_1) = a_1$  and  $s'_2(\theta'_2) = a'_2$ . The expected payoff for player 1 is 0 for  $(\theta_1, \theta_2)$  and 1 for  $(\theta'_1, \theta'_2)$ . The total expected payoff is 0.5.

First, we show that  $\sigma$  is an SFCCE by considering an optimal deviation for player 1. For  $\theta_1$ , selecting  $a'_1$  increases payoff to 0.5. On the other hand, for  $\theta'_1$ , selecting either of  $a_1$  or  $a'_1$  decreases payoff to 0.5. In total, player 1 cannot increase the expected payoff. Since the payoff for player 2 is always 0, this player never has incentive to deviate. Therefore,  $\sigma$  is an SFCCE. On the other hand,  $\sigma$  is not an ANFCCE since player 1 with  $\theta_1$  can gain by choosing  $a'_1$ .

Caragiannis et al. [2015] and Jin and Lu [2023] used the following definition of Bayes coarse correlated equilibria. This is similar to that of ANFCCEs, but defined on  $\Delta(A)^\Theta$ , not on  $\Delta(S)$ . To our knowledge, this class has not yet been given any specific name. A possible candidate is “coarse Bayesian solutions” since this notion is a coarse variant of Bayesian solutions.

**Definition D.4** (coarse Bayesian solutions). A type-wise distribution  $\pi \in \Delta(A)^\Theta$  is a coarse Bayesian solution if for any  $i \in N$ ,  $\theta'_i \in \Theta_i$ , and  $a'_i \in A_i$ , it holds that

$$\mathbb{E}_{\theta \sim \rho} \left[ \mathbf{1}_{\{\theta_i = \theta'_i\}} \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a)] \right] \geq \mathbb{E}_{\theta \sim \rho} \left[ \mathbf{1}_{\{\theta_i = \theta'_i\}} \mathbb{E}_{a \sim \pi(\theta)} [v_i(\theta; a'_i, a_{-i})] \right].$$