601.436/636 Algorithmic Game Theory **Spring 2022**
Homework #2 **Due:** February 24, 2022, 4:30pm

Remember: you may work in groups of up to three people, but must write up your solution entirely on your own. Collaboration is limited to discussing the problems – you may not look at, compare, reuse, etc. any text from anyone else in the class. Please include your list of collaborators on the first page of your submission. You may use the internet to look up formulas, definitions, etc., but may not simply look up the answers online.

Please include proofs with all of your answers, unless stated otherwise. Your solution must by typeset (*not* handwritten), and must be submitted by gradescope.

---

1 Potential Games (50 points)

A team game is a game in which all players have the same utility function: $u_1(\boldsymbol{s}) = \cdots = u_n(\boldsymbol{s})$ for every outcomes $\boldsymbol{s}$. In a dummy game, the utility of every player i is independent of its strategy: $u_i(\boldsymbol{s}) = u_i(\boldsymbol{s}_{-i}, s'_i)$ for every $s \in S$ and $s'_i \in S_i$.

Prove that a game with utilities $u_1 \ldots, u_n$ is a potential game (i.e., admits a potential function Φ) if and only if it is the sum of a team game $u^t_1, \ldots, u^t_n$ and a dummy game $u^d_1, \ldots, u^d_n$ (i.e., $u_i(\boldsymbol{s}) = u^t_i(\boldsymbol{s}) + u^d_i(\boldsymbol{s})$ for all i and $\boldsymbol{s}$)

Solution: Suppose that our game is a potential game, with potential function Φ. Let $u^t_i(\boldsymbol{s}) = \Phi(\boldsymbol{s})$ for all $i \in [n]$ and $\boldsymbol{s} \in S$. Note that by definition, the u^t_i's form a team game. Let $u^d_i(\boldsymbol{s}) = u_i(\boldsymbol{s}) - \Phi(\boldsymbol{s})$. Then clearly $u_i(\boldsymbol{s}) = u^t_i(\boldsymbol{s}) + u^d_i(\boldsymbol{s})$ for all i and $\boldsymbol{s}$, so we just need to show that the u^d_i's form a dummy game. To prove this, by the definition of a potential function we have that

$$u^d_i(\boldsymbol{s}_{-i}, s'_i) = u_i(\boldsymbol{s}_{-i}, s'_i) - \Phi(\boldsymbol{s}_{-i}, s'_i) = u_i(\boldsymbol{s}) - \Phi(\boldsymbol{s}) = u^d_i(\boldsymbol{s}),$$

and thus the u^d_i's form a dummy game.

For the other direction, assume that our game is the sum of a team game and a dummy game. Let $\Phi(\boldsymbol{s}) = u^t_1(\boldsymbol{s})$ (which by the definition of a team game is also equal to $u^t_i(\boldsymbol{s})$ for all $i \in [n]$). We claim that Φ is a potential function, thus implying that our game is a potential game. A few simple calculations suffices to prove this:

$$\begin{aligned}
u_i(\boldsymbol{s}_{-i}, s'_i) - u_i(\boldsymbol{s}) &= u^t_i(\boldsymbol{s}_{-i}, s'_i) + u^d_i(\boldsymbol{s}_{-i}, s'_i) - u^t_i(\boldsymbol{s}) - u^d_i(\boldsymbol{s}) \\
&= u^t_i(\boldsymbol{s}_{-i}, s'_i) + u^d_i(\boldsymbol{s}) - u^t_i(\boldsymbol{s}) - u^d_i(\boldsymbol{s}) \\
&= u^t_i(\boldsymbol{s}_{-i}, s'_i) - u^t_i(\boldsymbol{s}) \\
&= \Phi(\boldsymbol{s}_{-i}, s'_i) - \Phi(\boldsymbol{s})
\end{aligned}$$

2 No-Regret and Coarse Correlated Equilibria (50 points)

We proved that if every player uses a no-regret algorithm, then the time-averaged distribution converges to a coarse correlated equilibrium. Let's prove *almost* the converse. Consider a

k-player cost-minimization game in which $C_i(\boldsymbol{s}) \neq C_i(\boldsymbol{s}')$ for every agent i and every pair of strategy profiles $\boldsymbol{s}, \boldsymbol{s}' \in S$, (i.e., no agent has the same cost for two different profiles). Let σ be a coarse correlated equilibrium for this game. Prove that there exist no-regret algorithms $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_k$ for the agents so that the time-averaged history of the corresponding no-regret dynamics converges to σ as T tends to infinity.

Hint: Pre-program σ into $\mathcal{A}_1, \ldots, \mathcal{A}_k$. To make sure that each algorithm $\mathcal{A}_i$ is a no-regret algorithm, switch to your favorite no-regret algorithm (e.g., multiplicative weights) if some other agent j fails to use the agreed-upon algorithm $\mathcal{A}_j$.

Solution: Let's assume for now that the nonzero probabilities in σ are rational. Let M be their GCD. So there exists a function $f : \{0, 1, \ldots, M-1\} \to S$ so that $\frac{1}{M}|\{t \in \{0, 1, \ldots, M-1\} : f(t) = \boldsymbol{s}\}| = \sigma(\boldsymbol{s})$. We define the algorithm $\mathcal{A}_i$ as follows:

(a) At time t, play action $f(t \pmod M)_i$.

(b) If the cost at time t is not $C_i(f(t \pmod M))$, then switch the algorithm and for the rest of time use multiplicative weights.

We need to prove that if every player i uses $\mathcal{A}_i$ then the time-averaged distribution converges to σ as T tends to infinity, and that each $\mathcal{A}_i$ is a no-regret algorithm.

Let's prove first that if every player uses its assigned algorithm, then the second step of the algorithm is never executed by any player (i.e., no player switches to multiplicative weights). We prove this by induction. At time 0 it is clearly true since every player j will play action $f(0)_j$ and thus the realized outcome will be $f(0)$ and so the cost to player i will be $C_i(f(0))$, and thus no player will execute the second step and switch to multiplicative weights. If no player has switched to multiplicative weights by the beginning of time t, then in time t player j will play $f(t \pmod M)_j$ and thus the realized outcome will be $f(t \pmod M)$ and so the cost to player i will be $C_i(f(t \pmod M))$, and thus the player will not switch to multiplicative weights.

Now we know that if every player uses its assigned algorithm then they never switch to multiplicative weights. Let k be the largest integer such that $kM \leq T$. Then for every $\boldsymbol{s} \in S$, the time-averaged fraction in which it appeared is

$$\frac{kM}{T}\sigma(\boldsymbol{s}) \leq |\{t \in \{0, 1, \ldots, T-1\} : f(t \pmod M) = \boldsymbol{s}\}| < \frac{(k+1)M}{T}\sigma(\boldsymbol{s}),$$

and thus as T goes to infinity this converges to $\sigma(\boldsymbol{s})$

Now we need to prove that each $\mathcal{A}_i$ is a no-regret algorithm. There are two cases: if i eventually switches to multiplicative weights, or if doesn't. If it does switch, then after it switches it is (by definition) a no-regret algorithm, and thus the regret tends to 0 as T tends towards infinity as required. On the other hand, if it never switches, then since $C_i(\boldsymbol{s}) \neq C_i(\boldsymbol{s}')$ for all $i, \boldsymbol{s}, \boldsymbol{s}'$ we know that the time-averaged distribution of play converges to σ which is a CCE, and thus by the definition of a CCE player i would not be able to do better by switching to any fixed action and thus the algorithm is no-regret.

If the probabilities are irrational, then instead of a cycle of time M we just need a function $f : \mathbb{N} \to S$ so that as T tends towards infinity the fraction of time in which $f(t) = \boldsymbol{s}$ approaches $\sigma(\boldsymbol{s})$ for all $\boldsymbol{s} \in S$. Such functions clearly exist, so we can just use one of them.