Algorithmic Game Theory - handout 6 and 7

Uriel Feige

2 and 9 May 2013

We present algorithms with sublinear regret for sequential optimization, both in the expert setting and in the bandit setting. We show that results on external regret can be used in order to prove the minimax theorem, and more generally, can be used in order to design algorithms that approximately solve linear programs. We show that swap regret can be used in order to find a correlated equilibrium.

Homework.

1. Recall that we have seen in class that the randomized weighted majority (RWM) algorithm achieves external regret $O(\sqrt{T \log N})$ when payoffs are either 0 or 1. Prove this statement when payoffs can have arbitrary values in the range $[0, 1]$.

2. Recall that we have shown in class that if two players repeatedly play a finite 0-sum game and each of them happens to have low external regret, then the empirical averages of what each player plays is (nearly) a minimax strategy. Show that a similar statement fails to hold if one considers arbitrary (not 0-sum) finite 2-player games and the concept of a mixed Nash equilibrium. Namely, show a 2-player finite game G and a sequence of actions played by the two players in a T round version of G such that no player has external regret at all (the external regret is either 0 or negative), but the empirical averages of the strategies that the players play do not form a mixed Nash equilibrium.