Thomas Kesselheim
Anna Heuser

December 10, 2025

Algorithms and Uncertainty

Winter Term 2025/26

Tutorial Session - Week 9

Exercise 1:
Consider the problem of learning for Pandora's box, but this time we know the distribution of $n - f$ boxes and only need to learn the distribution of the other f boxes. Again let $v_i \in [0, 1]$ with probability 1. Show that for all $\epsilon, \delta > 0$, if $T \geq \frac{f^2 \ln(2f/\delta)}{\epsilon^2}$, then the expected reward for the learned policy is at least $V^* - \epsilon$ with probability at least $1 - \delta$.