Thomas Kesselheim   October 21, 2024
Rebecca Lehming   Due: October 28, 2024 at 10:00 p.m.

Algorithms and Uncertainty
Winter Term 2024/25
Exercise Set 3

If you want to hand in your solutions for this problem set, please send them via email to
rlehming@uni-bonn.de *by Monday evening. Of course, submitting solutions in groups is also*
possible.

If you would like to present one of the solutions in class, please also send an email to
rlehming@uni-bonn.de *stating* **which task** *you would like to present in* **which of the tu-**
torials*. Deadline for the email is Monday, 10:00 pm. Please note that the tasks will be*
allocated on a first-come-first-serve basis, so sending this email earlier than Monday evening
is highly recommended.

Exercise 1:   (5 Points)
Consider the following randomized rounding for fractional Ski Rental. In step t, given that
you have only rented skis before, buy the skis with probability $\frac{x_{\text{buy}}^{(t)} - x_{\text{buy}}^{(t-1)}}{1 - x_{\text{buy}}^{(t-1)}}$, otherwise rent them
again. Show that if the underlying fractional algorithm is α-competitive so is the randomized
integral algorithm.

Exercise 2:   (4 Points)
We would like to show that optimal online algorithms that know the input distribution in
advance are always deterministic.
To this end, we use the notation of Lecture 6. We have a finite class of deterministic algo-
rithms $\mathcal{A}$ and a finite class of instances $\mathcal{X}$. Let X be a random variable with values in $\mathcal{X}$ and
A be an *independent* random variable with values in $\mathcal{A}$. Show that there is an $a \in \mathcal{A}$ such
that $\mathbf{E}[c(a, X)] \leq \mathbf{E}[c(A, X)]$.

Hint: Use the same techniques as in the proof of Yao's principle.

Exercise 3:   (1+5 Points)
Consider the following online problem: We would like to buy a house. Each time we view a
house, we get to know how good it is. Immediately after, we have to decide whether to buy
the current house or not. If we do not buy it now, there will be no other opportunity later.
Formally, this means that a sequence of values is revealed one after the other. Note that
there are no assumptions concerning the order of the sequence. In particular, the order could
be chosen by an adversary. The algorithm can always decide if it stops the sequence at its
current value or if it waits. We say that an algorithm wins if it stops the sequence at its
highest value.

(a) Assume that the length of the sequence n is known. Give a randomized online algorithm
    that wins with probability $\frac{1}{n}$.

(b) Use Yao's principle to show that no algorithm wins with higher probability.

Hint: Use the profit-maximization variant of Yao's principle. The profit is 1 if the algorithm wins, 0 otherwise. The offline optimum is clearly always 1. Construct a probability distribution over sequences of the form $1, 2, \ldots, t, 0, \ldots, 0$.