Seminar in Statistics: Learning Blackjack – Coding task I
ETH Zurich, Spring Semester 2016

Jonas Peters, Christina Heinze, Gian Thanei, Nicolai Meinshausen

March 1, 2016

For this coding task you are asked to fill out some missing code fragments in the classes `RandomPlayer`, `AlwaysStandPlayer` and `EpsGreedyContextualBandit`. The code framework is available from `https://github.com/christinaheinze/learning-blackjack` and the positions of the missing code fragments are marked with the comments `### your code goes here` or `### modify`. You can work in teams of two students.

`RandomPlayer`

The `RandomPlayer` plays randomly as he chooses to split, to double down and to hit with a probability of 0.5, respectively.

`AlwaysStandPlayer`

The `AlwaysStandPlayer` never splits nor doubles down and he always stands.

Hint: For both the `RandomPlayer` and the `AlwaysStandPlayer` you only need to write one line of code.

`EpsGreedyContextualBandit`

This class extends the `Bandit` class which models the first decision between double down, hit or stand as a three-armed bandit. This player never splits and after the first decision he always stands. The `EpsGreedyContextualBandit` uses the history by taking into account the game states and the corresponding decisions taken in each game. Initially, each of the three actions is chosen with a probability of 1/3. After a number of games, the performance is evaluated and updated according to an ϵ-greedy strategy where each game state is treated separately as a single bandit.

Deadline

Please send your solutions to `heinze@stat.math.ethz.ch` until March 20th 2016. If you decide to work in groups of two, please include both of your names in the email.