

## 1 Decisions and MDPs

1. Suppose you are a turtle guru, and you have found your way into a turtle-themed spinoff of "Who Wants To Be A Millionaire". You have gotten to the 50,000 turtle question, and would win 15,500 turtles if you decide to quit before seeing the 50,000 turtle question. You have a 14 percent chance of getting the next question right. What is the expected payoff of going for the next question? What is the optimal decision?
2. Is a sequence of actions sufficient to get an agent in a nondeterministic environment to reach some goal?
3. What are the components of an MDP?
4. In class, we talked about an example environment where taking the "North" action was nondeterministic. Is it possible for all actions in an environment to be nondeterministic?
5. (**No Answer Necessary**) Review any gopollock questions you got wrong at [ter.ps/cmsc389fgp3](http://ter.ps/cmsc389fgp3)

## 2 Exploring OpenAI and GridWorld

Note: This week, we will begin a very cursory exploration of OpenAI Gym and FrozenLake (a version of the GridWorld environment that we discussed in class). We will solve FrozenLake in later lectures and assignments using a variety of RL algorithms that we will learn.

### 1. OpenAI's FrozenLake

- Open the view-only Python 3 Notebook at [ter.ps/cmsc389fhw3](http://ter.ps/cmsc389fhw3)
- Click on "File", "Save a Copy in Drive"
- Go to your copy of the notebook and complete the exercises
- Share your finished notebook with [cmsc389f@gmail.com](mailto:cmsc389f@gmail.com), "HW2 - First Name, Last Name"