Incentives in Computer Science (COMS 4995-6): Exercise Set #6

Due by Noon on Wednesday, March 4, 2020

Instructions:

(1) You can work individually or in a pair. If you work in a pair, the two of you should submit a single write-up.

(2) Submission instructions: We are using Gradescope for the homework submissions. Go to www.gradescope.com to either login or create a new account. Use the course code MKRKK6 to register for COMS 4995-6. Only one person needs to submit the assignment. When submitting, please remember to add your partner's name (if any) in Gradescope.

(3) Please type your solutions if possible. We encourage you to use the LaTeX template provided on the course home page.

(4) Write convincingly but not excessively. You should be able to fit all of your solutions into 2–3 pages, if not less.

(5) Except where otherwise noted, you may refer to the course lecture notes and the specific supplementary readings listed on the course Web page *only*.

(6) You can discuss the exercises verbally at a high level with other groups. And of course, you are encouraged to contact the course staff (via Piazza or office hours) for additional help.

(7) If you discuss solution approaches with anyone outside of your group, you must list their names on the front page of your write-up.

(8) Refer to the course Web site for the late day policy.

Exercise 31

Prove that the linear scoring rule $S(\mathbf{p}, i) = p_i$ is not even weakly proper. (Here i is an outcome in the set X, $\mathbf{p}$ is a probability distribution over X, and p_i denotes the amount of probability that $\mathbf{p}$ assigns to i.)

Exercise 32

The *spherical scoring rule* is defined as

$$S(\mathbf{p}, i) = \frac{p_i}{\|\mathbf{p}\|},$$

where $\|\mathbf{p}\| = \sqrt{\sum_{i \in X} p_i^2}$ is the Euclidean norm of $\mathbf{p}$. Prove that when X consists of two outcomes, this scoring rule is strictly proper.[1]

Exercise 33

We saw in lecture that the worst-case loss of an automated market maker based on the logarithmic scoring rule is $\ln |X|$, where X is the outcome set, assuming that the initial probability distribution is the uniform distribution. What is the worst-case loss of a market based on the quadratic scoring rule (again, initialized with the uniform distribution)?

[1] It is strictly proper for any number of outcomes, but you don't have to prove this.

Exercise 34

Repeat the previous exercise for an automated market maker based on the spherical scoring rule in Exercise 32.

Exercise 35

Suppose we hire an expert to predict whether it will be sunny or rainy tomorrow. Assume that, absent any information (i.e., unconditionally), the weather is equally likely to be sunny or rainy. Assume also that our expert receives one of two signals, S or R, with the following probabilities (conditional on what the true weather will be):

1. If in fact it will be sunny tomorrow, the expert receives the report S with 70% probability and R otherwise.
2. If in fact it will be rainy tomorrow, the expert receives the report R with 60% probability and S otherwise.

The first question is:

(a) How likely is the weather to be sunny when the expert receives the signal S? The signal R?

If we use a strictly proper scoring rule to pay the expert, she has an incentive to accurately reveal her information. But strictly proper scoring rules can also incentivize the expert to become more accurate. Assume that the expert can purchase an additional signal, drawn independently from the same distribution as the first, at a cost of $c > 0$.

(b) If the expert is paid using the logarithmic scoring rule (with $S(\mathbf{p}, i) = \ln p_i$ for a prediction $\mathbf{p}$ and outcome i), when (i.e., for what values of c) will she be incentivized to purchase an additional signal? Your answer might be different for the case in which her first signal is S and in which her first signal is R.

(c) For which values of $a > 0$ does the scaled logarithmic scoring rule $S(\mathbf{p}, i) = a \ln p_i$ incentivize the expert to always purchase an additional signal? (Your answer should be a function of c.)

Exercise 36

Recall that the Peer Prediction mechanism has a truthful equilibrium. In other words, if every other participant reports their signal truthfully, then your unique best response is to also report truthfully. Give an example where the Peer Prediction mechanism also has a non-truthful equilibrium (i.e., at least one player is not truthful, and neither player can increase their expected payoff via a unilateral deviation).[2]

[2] The Peer Prediction mechanism is parameterized by a strictly proper scoring rule S. In your example, feel free to choose whichever such rule is most convenient for you. You can also use any joint distribution on signals that you want.