Problem Set 1

CS 286r

beginning of class, Monday 10/1

Preamble

You may work in pairs and not discuss this problem set with anyone other than your (optional) partner. Turn in a single problem set for each pair.

Due by email to the course TF as a PDF (we suggest you write in LaTex) before class begins on Monday 10/1. You are not expected to be immediately familiar with all concepts on this problem set, but should be capable of learning them as necessary. You will also have to review at least one paper not read for class, so start early!

All answers should be formal proofs unless an informal argument is specified. E.g., a question that asks for an informal answer may require a few lines of non-mathematical argument, but otherwise you should derive an answer formally and show your work.

1 Game theory

1.1 Supergames and subgame perfection

A "supergame" is an infinitely repeated "stage game," where a stage game is a "one-shot" simultaneous action single round game (finite action spaces), like the Prisoner's Dilemma. The stage game is played over an infinite number of periods, 0, 1, 2, . . . , on to infinity, and after each period all actions are revealed and players receive the stage game's payoff in that round. Assume the stage game payoffs are bounded and that players discount future returns by a discount factor $\delta \in (0, 1)$.

1. Prove that playing a Nash equilibrium of the stage game at each round of the supergame is a subgame perfect equilibrium of the supergame. What can some off-equilibrium beliefs be in this case? **Hint:** what if we "commit" to this Nash equilibrium?

2. Are these necessarily the only subgame perfect equilibrium of a supergame? (Informally)

1.2 Exit and perfect Bayesian equilibrium

A game is said to allow "exit" if a player can guarantee a payoff of zero.

1. What is the expected score of players playing in perfect Bayesian equilibrium of any zero-sum game?
2. What might players expect if the game is not zero-sum but positive-constant-sum (players can receive a net positive but bounded payment)? Why might market makers subsidize prediction markets? (Informally)

2 Building a Scoring Rule

2.1 Scoring Rule from a Convex Function

Let Δ_n be the set of all probability distributions over n outcomes, and consider the function

$$g(p) = \sum_{i=0}^{n-1} p_i \log p_i \tag{1}$$

defined over vectors from Δ_n for any finite n.

1. Is this function convex? Is it strictly convex?
2. Using the Savage representation from Gneiting and Raftery's survey, derive a scoring rule from this function.
3. Which popular scoring rule does this convex function specify?

2.2 Properties of the Logarithmic Scoring Rule

The logarithmic scoring rule is any function of the form $s(p, \omega) = a_\omega + b \log p_\omega$, although we will usually consider it with $a_\omega = 0, \forall\omega$ and $b = 1$.

1. Why is the logarithmic scoring rule strictly proper? **Hint:** It is derived from a strictly —.
2. Write the logarithmic scoring rule's expected score function. What function is this?

2.3 A Market Scoring Rule

Now consider running a prediction market that uses the logarithmic scoring rule as a market scoring rule. For clarity, we refer to an expert's payment for each prediction as a *net score* since it is the difference of that prediction and the immediately preceeding one, and still refer to the *score* of each prediction separately.

1. Why does a MSR market open with an initial prediction? (Informally)
2. What is a market maker's worst-case loss (abstractly)? What is the worst-case loss when the market maker uses the logarithmic scoring rule? When is this worst-case loss realized? What initial prediction, if any, minimizes the worst-case loss?
3. Write the expected *net* score function for a prediction in this market where an expert has beliefs $q \in \Delta_n$, predicts $p \in \Delta_n$ and the previous expert predicted $p' \in \Delta_n$.
4. Why does a strictly proper scoring rule, which always uniquely maximizes the score for a prediction when it's "honest," also cause the unique maximum of a prediction's net score to be when the trader is "honest?" Remember the expert's beliefs and the prior prediction can be arbitrary (but still in Δ_n).

2.4 A Cost Function

Now consider we are running a logarithmic cost function market with only two outcomes.

1. If the outstanding shares vector is $X \in \mathrm{R}^2$, what is the cost of changing the shares by $Y \in \mathrm{R}^2$?
2. Given outstanding quanities X, what is the implied likelihood of each event?
3. What does the b parameter in the cost function do? (Informally)

Intrade.com offers derivatives on future events like the U.S. presidential election. Instead of using a cost function market maker they allow their users to trade shares using a double auction, just like how stocks are traded.

4. How are prices related to the likelihood the event will happen on Intrade? (Informally)

3 Scoring Experts

Suppose we hire an expert to predict whether it will be **fair** or **cloudy** tomorrow. Assume our expert receives one of two signals, F or C, according to the following distributions

	F	C
fair	70%	30%
cloudy	40%	60%

that is, if tomorrow will be fair then 70% of the time the expert receives the signal F, and etc..

1. How likely is the weather to be fair when the expert receives the signal F? The signal C?
2. If we pay the expert a flat fee for its prediction, what incentive does it have to be accurate? (Informally)

If we use a strictly proper scoring rule to pay the expert it has an incentive to accurately reveal its information. But strictly proper scoring rules may also incentive the expert to become more accurate. Assume the expert can purchase an additional signal, drawn independently from the same distribution as the first, at a cost of c. When we refer to the log scoring rule, we mean $s(p, \omega) = \log p_\omega$.

3. If the expert is paid using the logarithmic scoring rule, when will it be incentivized to purchase additional information? (i.e. after receiving the signal F? What about after receiving the signal C?)
4. What value for b on the logarithmic scoring rule $s(p, \omega) = b \log p_\omega$ incentivizes the expert to always purchase additional information?
5. Would changing the a value on the logarithmic scoring rule $s(p, \omega) = a + b \log p_\omega$ affect the expert's incentives? Why might you want this to not be zero, even if incentives are not affected? (Informally)

4 Some Research using Scoring Rules

These questions are based on Yiling and Pennock's paper, "A Utility Framework for Bounded-Loss Market Makers."[1]

Consider, for the moment, a market maker with linear utility $u(m) = m$ and a uniform distribution $\vec{\pi}$ over $n \in \mathrm{N}^+$ outcomes.

1. What are the market maker's initial risk-neutral probabilities?
2. What are they after a trader buys 10 shares of the first outcome?
3. What undesirable property(ies?) does this market maker have? Was this a poor choice of utility function? (Informally)

Assume now that we bound the domain of u below by -5.

4. In your own words, what does this mean? (Informally)
5. Is this market maker better or worse than the previous one? Why or why not? (Informally)
6. Suppose a trader wants to buy 20 shares in outcome x. When will this trade be possible?

[1] Available at `http://yiling.seas.harvard.edu/wp-content/uploads/MM_long.pdf`