Microeconomic Theory II - UNI, Lima
Homework #4 - Due date: September 23rd, in class.

1. **Cournot competition with two-sided incomplete information.** Consider two firms competing à la Cournot, facing inverse demand function $p(Q) = a - Q$, where $Q = q_1 + q_2$ denotes aggregate output. Every firm i's marginal cost is privately known, and takes value c_L with probability $\theta \in (0, 1)$ and value $c_H > c_L$ with probability $1 - \theta$. The two draws are independent, the prior θ is common knowledge, and $a > c_H$. Throughout, denote by

$$\bar{c} \equiv \theta c_L + (1 - \theta)c_H \quad \text{and} \quad \Delta \equiv c_H - c_L > 0$$

the expected marginal cost and the cost gap. In contrast to the setting we analyzed in class, where only one firm held private information, here *both* firms are uninformed about their rival's cost.

Answer the following questions.

   (a) Describe this Bayesian game (players, type spaces, prior beliefs, strategies and payoffs). Then find the best response function of every type of firm i, expressed as a function of the expected output of its rival, $\overline{Q} \equiv \theta q_L + (1 - \theta)q_H$, and check its second-order condition.

   (b) Solve the resulting four-equation system for the symmetric Bayesian Nash equilibrium (BNE), q_L^* and q_H^*. Report $\overline{Q}^*$ as well, and show that every type's equilibrium profit equals the square of its equilibrium output.

   (c) Compare your equilibrium against two benchmarks. First, the complete-information Cournot duopoly in which both costs are commonly known, one firm producing at c_L and the other at c_H. Second, the one-sided-uncertainty model of the lecture, in which firm 1's cost is private while firm 2's cost is commonly known and equal to $\bar{c}$. Does two-sided uncertainty change the behavior of the informed firm relative to the one-sided case?

   (d) *Comparative statics.* Find how the spread $q_L^* - q_H^*$ depends on the prior θ and on the cost gap Δ. Compare this spread against the one arising under complete information. Then find how each type's output responds to a marginal increase in θ, and interpret. [*Hint*: the answer to the first question is less obvious than it looks.]

   (e) Find the condition on θ guaranteeing that the equilibrium is interior, so that the high-cost type produces a positive quantity, and state the restriction on the primitives under which the equilibrium is interior for every $\theta \in (0, 1)$.

   (f) *Numerical example.* Evaluate your results at $a = 10$, $c_L = 1$ and $c_H = 4$, first for $\theta = \dfrac{1}{2}$ and then for $\theta = \dfrac{1}{3}$. Report $\bar{c}$, $\overline{Q}^*$, both outputs, both profits, and the spread. Compare against the complete-information duopoly with costs $(c_L, c_H) = (1, 4)$.

2. **First-price auction with a reserve price.** Consider a seller who owns one indivisible unit of a good, which she values at zero, and who faces $n = 2$ risk-neutral bidders. Every bidder i privately observes his own valuation v_i, independently drawn from a uniform distribution on $[0, 1]$. Before bids are submitted, the seller publicly commits to a reserve price $r \in [0, 1)$: bids below r are rejected, and if no bid is accepted the seller keeps the good. The auction is first-price, so the winner pays his own bid.

   Answer the following questions.

   (a) Find the participation cutoff, that is, the valuation below which a bidder submits no acceptable bid. Then, letting $b(v)$ denote the symmetric equilibrium bidding function of a participating bidder, find the probability that a bidder with valuation v wins the auction.

   (b) Find the symmetric equilibrium bidding function $b(v)$ for every participating bidder. Then *verify* that your candidate is indeed a symmetric equilibrium: show that a bidder with valuation v who submits the bid of a rival type z finds $z = v$ optimal, checking the first-order condition, the second-order condition, and global optimality. Confirm, in addition, that participating dominates staying out.

   (c) Show that $b(v)$ is increasing in v on the participation region and that $b(r) = r$. How is a participating bidder's bid affected by a marginal increase in the reserve price? Which bidders raise their bids by the most?

   (d) Find the seller's expected revenue as a function of the reserve price, $R(r)$. Confirm that $R(0)$ coincides with the expected revenue of the standard first-price auction without a reserve price.

   (e) Find the revenue-maximizing reserve price r^* and check its second-order condition. Then repeat the exercise for an arbitrary number of bidders $n \geq 2$, and confirm that r^* is independent of n.

   (f) Explain intuitively why a reserve price raises the seller's revenue even though it excludes some buyers. [*Hint*: compare the order of magnitude of the gain against that of the loss when r is small.] Then compute total surplus as a function of r and decompose, at $r = r^*$, the seller's gain and the bidders' loss.

   (g) *Numerical example.* Set $r = \dfrac{1}{2}$. Report the bid of a bidder with valuation $v = \dfrac{3}{4}$ and of a bidder with valuation $v = 1$, comparing each against the bid the same bidder would submit if $r = 0$. Report the seller's expected revenue, the probability that the good goes unsold, total surplus, and the bidders' expected surplus, both at $r = \dfrac{1}{2}$ and at $r = 0$.

3. **Labor-market signaling with three worker types.** Consider a worker whose productivity takes one of three values, $t_1 < t_2 < t_3$, with prior probabilities p_1, p_2 and p_3, which add up to one. The worker privately observes his productivity and then chooses an education level $e \geq 0$, at a cost $c(e, t) = \dfrac{e}{t}$, so that more productive workers acquire education more cheaply. After observing e, but not t, two identical firms compete à la Bertrand in wages, so that the worker is paid a wage equal to his expected productivity given the firms' beliefs after observing e. The worker's payoff is

$w - \frac{e}{t}$, and education is purely a signal, since it does not raise productivity. The lecture version of this model had only two types, so the incentive constraints formed a single pair. With three types the constraints form a chain, and *non-adjacent* constraints must be checked separately.

Answer the following questions.

(a) Define the firms' beliefs and the wage schedule, and define a fully separating perfect Bayesian equilibrium (PBE) of this game. Show that, in every separating PBE, the lowest type acquires no education, $e_1 = 0$.

(b) Write down the six incentive-compatibility constraints of a fully separating PBE, and characterize the entire set of separating education profiles (e_2, e_3). Then find the least-costly separating PBE, along with the off-equilibrium beliefs that support it.

(c) Consider the two *non-adjacent* constraints, namely the one preventing type t_1 from mimicking type t_3 and the one preventing type t_3 from mimicking type t_1. Evaluate both at the least-costly separating profile and report the slack of each. Are they implied by the adjacent constraints?

(d) Consider now a partially pooling PBE in which type t_1 separates at $e_1 = 0$ while types t_2 and t_3 pool at a common education level $e_P > 0$. Find the wage that firms offer at e_P, the range of education levels e_P sustaining this equilibrium, and the off-equilibrium beliefs it requires.

(e) Apply the Cho and Kreps' (1987) Intuitive Criterion. Does the partially pooling PBE of part (d) survive? Which separating PBEs survive? Does the criterion select the least-costly separating PBE, as it does in the two-type model of the lecture? Interpret. [*Hint*: for an off-equilibrium education e', first identify the types for which e' is equilibrium dominated, and then check whether the remaining types gain when firms concentrate their beliefs on the surviving set.]

(f) *Numerical example.* Evaluate your results at $t_1 = 1$, $t_2 = 2$ and $t_3 = 4$, with $p_1 = p_2 = p_3 = \frac{1}{3}$. Report the least-costly separating profile, the entire separating set, the slack of every incentive constraint, the payoffs of every type, the pooling equilibrium of part (d) at $e_P = 3$, and an explicit deviation that destroys it.

4. **Entry deterrence against an incumbent of unknown type, with a noisy signal.** Consider an incumbent firm that is either *tough* with prior probability $\mu \in (0, 1)$ or *weak* with probability $1 - \mu$, and a potential entrant. In contrast to the entry game of the lecture, the entrant does not move on the prior alone: before choosing whether to enter, he observes a public signal $s \in \{t, w\}$ about the incumbent's type, which is correct with probability

$$\Pr(s = t \mid \text{tough}) = \Pr(s = w \mid \text{weak}) = \rho \in \left(\frac{1}{2}, 1\right],$$

so that ρ measures the accuracy of the signal. The timing is as follows.

- Nature draws the incumbent's type, which only the incumbent observes, and then draws the public signal s, which both firms observe.
- The entrant chooses whether to enter (In) or to stay out (Out).
- If the entrant stays out, he earns 0 and the incumbent earns its monopoly profit. If the entrant enters, the incumbent chooses whether to fight or to accommodate. Fighting is a dominant action for the tough incumbent, whereas accommodating is a dominant action for the weak incumbent. The entrant earns $-F < 0$ if he is fought and $A > 0$ if he is accommodated.

Answer the following questions.

(a) Find the entrant's posterior belief that the incumbent is tough after each signal realization, μ_t and μ_w, using Bayes' rule. Show that $\mu_w < \mu < \mu_t$ for every $\rho > \frac{1}{2}$, and confirm that the two posteriors average back to the prior.

(b) Find the cutoff posterior $\overline{\mu}$ such that the entrant enters if and only if his posterior lies below it. How is $\overline{\mu}$ affected by A and by F?

(c) Find the Bayesian Nash equilibrium of this game as a function of μ and ρ. In particular, find the cutoff accuracy $\overline{\rho}$ above which the entrant is deterred after the tough signal, and the second cutoff that applies when the prior alone already deters entry. How is $\overline{\rho}$ affected by the prior μ?

(d) *Comparative statics.* Is entry deterrence more effective as the signal becomes more accurate? Answer this question in three ways: using the probability that a tough incumbent avoids entry, the probability that a weak incumbent avoids entry, and the unconditional probability that entry does not occur. Do the three measures agree? Then find how the entrant's expected payoff responds to ρ, and interpret.

(e) Evaluate the two limiting cases $\rho = \frac{1}{2}$ and $\rho = 1$, and relate each of them to a model we studied in class.

(f) *Numerical example.* Evaluate your results at $A = 4$, $F = 6$ and $\mu = \frac{1}{4}$. Report $\overline{\mu}$ and $\overline{\rho}$, both posteriors and the equilibrium for $\rho = \frac{3}{5}$, $\rho = \frac{2}{3}$ and $\rho = \frac{3}{4}$, the unconditional probability of no entry, and the value of the signal to the entrant.

Additional practice (not graded). For further practice on this material, see the textbook, Chapters 7-8.