Economics 713, Final Exam
Spring 2006

Answer two of the three equally-weighted questions. Explain your answers. Use a separate blue book for every question you answer.

Question 1:

Consider an education signaling model where a worker's education adds to his productivity. There are two firms in the labor market who compete for the services of a worker. The worker has productivity $\theta \in \{\theta_L, \theta_H\}$ which is unobserved by the firm; he can acquire education $e \geq 0$ to signal his productivity. The prior probability of a high type is $P(\theta = \theta_H) = \lambda$. The firm's profits from hiring a worker of productivity θ and education level e are $\pi(e, w, \theta) = \theta + e - w$, where w is the wage paid to the worker. The worker has utility from w and e given by $u(w, e; \theta_k) = w - c_k e$, $k \in \{H, L\}$, $c_H < c_L$.

After nature draws a type for the worker, the worker acquires an education level. Firms then simultaneously post wage offers to the worker, which the worker will afterwards either accept or reject. (Assume his reservation utility is zero).

a) In a Perfect Bayesian equilibrium, firm's behavior in the wage-setting part of the game tree for any observed education level gives rise to a wage schedule $w(e)$. What must such a wage schedule look like? Draw a diagram and explain.

b) Give necessary and sufficient conditions on the parameters of the game such that a Perfect Bayesian (or sequential) equilibrium exists. (Hint: Use your answer to a).

c) Fully characterize a PBE of this game (including out-of-equilibrium beliefs).

d) Suppose that instead of acquiring education at the start of the game, the worker only sends a message e at the start of the game, signifying his intention to acquire education e (though this signal does not actually commit him to acquiring exactly e). Firms then simultaneously post wage schedules $w : \mathbb{R}_0^+ \to \mathbb{R}_0^+$; then the worker acquires some education e' and gets paid $w(e')$. Show that in this game, no separating equilibrium exists. Does a pooling equilibrium exist? (A separating equilibrium here is one in which workers of different types send different messages; in a pooling equilibrium, all types send the same messages).

Parts a)-c) are straightforward. For d), suppose a separating equilibrium did exist, and consider the optimal wage schedules that firms post after observing e_H^*, the equilibrium message of the high type. As firms now believe that they are facing a high type with probability 1, they must post wages $w(e) = \theta_H$ at least for all education levels $e \in [0, e^{HCS}]$ by the usual Bertrand arguments. But then a low type can mimic a high *type at no cost to himself by choosing the (costless) message e_H^* and afterwards acquire f.ex. no education at all. Hence, a separating equilibrium cannot exist. Pooling equilibria exist whenever the Rothschild-Stiglitz screening equilibrium exists - both types send some identical message at the start of the game, and then the usual screening equilibrium outcome results.*

Question 2:

Consider the following principal-agent problem involving a regulator (the principal) and a monopolist of unobservable type (the agent): the monopolist has constant marginal costs of production $\theta \in [\underline{\theta}, \bar{\theta}]$ which are unknown to the regulator. The monopolist faces a known inverse demand function $P(q)$ with $P'(q) < 0$ and P continuous. θ is distributed according to a distribution function F which is everywhere continuous, with continuous associated density $f > 0$.

A contract is a tuple of functions (p, q, s, r) that depend on θ. The regulator sets prices p as well as quantity q; additionally, the regulator may give the monopolist a subsidy s. The regulator shuts the firm down with probability $1 - r$. The monopolist's utility from such a contract is $V(p, q, s, r; \theta) := \pi(\theta) = r[(p - \theta)q] + s$.

Define a *feasible contract* as a tuple of functions (p, q, s, r) such that, for all $\theta \in [\underline{\theta}, \bar{\theta}]$:

$$r(\theta) \in [0, 1] \qquad (PR)$$

$$p(\theta) = P(q(\theta)) \qquad (D)$$

$$\pi(\theta) \geq 0 \qquad (P_\theta)$$

$$\text{The monopolist } \theta \text{ finds it optimal to reveal his type correctly} \qquad (R_\theta)$$

a) Provide an explicit expression for R_θ (type $\theta's$ revelation constraint).

b) Show that a mechanism is feasible only if:

$$\begin{aligned} &(PR),\ (D) \text{ and } (P_{\bar{\theta}}) \text{ hold,} \\ &r(\theta)q(\theta) \text{ is decreasing in } \theta, \text{ and} \\ \pi(\theta) \quad = \quad &\pi(\bar{\theta}) + \int_{\theta}^{\bar{\theta}} r(\tilde{\theta})q(\tilde{\theta})d\tilde{\theta}. \end{aligned}$$

c) Define $U(q) = \int_0^q P(\hat{q})d\hat{q}$. Then the consumer surplus in this market is $U(q) - qP(q)$. The regulator wishes to maximize a weighted average of expected consumer and producer surplus, i.e. the principal's objective function can be written as

$$\int_{\underline{\theta}}^{\bar{\theta}} \{[U(q(\theta)) - q(\theta)P(q(\theta))]\, r(\theta) - s(\theta)\}\, f(\theta)d\theta + \alpha \int_{\underline{\theta}}^{\bar{\theta}} \pi(\theta)f(\theta)d\theta,$$

where we take $\alpha \in (0,1)$. It can be shown that for any *feasible* contract, the regulator's objective function equals

$$\int_{\underline{\theta}}^{\bar{\theta}} \left\{ U(q(\theta)) - \left[\theta + (1-\alpha)\frac{F(\theta)}{f(\theta)} \right] q(\theta) \right\} r(\theta) f(\theta) d\theta - (1-\alpha)\pi(\bar{\theta}).$$

Use these facts to characterize the optimal contract as completely as you can; discuss your results. In your characterization, make sure to answer the following questions:

(i) When is it optimal to set $r(\theta)$ to 1 or 0? Is it ever optimal to set $r(\theta) \in (0,1)$?

(ii) What characterizes the optimal $q(\theta)$? Compare with the first-best.

a) and b) are straightforward. The first order condition w.r.t $q(\theta)$ is

$$U'(q^*(\theta)) = P(q^*(\theta)) = \theta + (1-\alpha)\frac{F(\theta)}{f(\theta)}, \forall \theta.$$

First note $q^*(\theta)$ is indeed decreasing in θ whenever $\frac{F(\theta)}{f(\theta)}$ is increasing in θ. As we must have $P(q^{CI}(\theta)) = \theta$ in the first-best (complete information) allocation, the equation above implies $q^{CI}(\theta) < q^*(\theta)$ for all $\theta > \bar{\theta}$, the usual 'no distortion at the top'-result. All other quantities are distorted downwards, so that a high-cost producer will prefer to produce his allocated (very low) quantity instead of the disproportionally larger quantity he would have to produce were he to pretend to be a good type (low-cost producer).

The regulator will set $r^*(\theta) = 1$ whenever $U(q^*(\theta)) - \left[\theta + (1-\alpha)\frac{F(\theta)}{f(\theta)}\right] q^*(\theta) > 0$; if there exist θ such that $U(q^*(\theta)) - \left[\theta + (1-\alpha)\frac{F(\theta)}{f(\theta)}\right] q^*(\theta) < 0$, he will set $r^*(\theta) = 0$. (This might happen if the consumer surplus goes to 0 at very low relevant quantities, for example because there is a prohibitive price - in this case, it is not worthwhile to let the firm produce). $r^*(\theta) \in (0,1)$ can be weakly optimal only for θ such that the regulator's utility from that type is 0 already.

Finally, the regulator will set $s^*(\bar{\theta})$ so that $\pi(\bar{\theta}) = \bar{u}$, the monopolist's reservation utility, and $s^*(\theta)$ satisfies the revelation constraints for all types everywhere else:

$$s^*(\theta) = \pi(\bar{\theta}) - r^*(\theta)[(P(q^*(\theta)) - \theta)q^*(\theta)] + \int_{\theta}^{\tilde{\theta}} r^*(\tilde{\theta}) q^*(\tilde{\theta}) d\tilde{\theta}. \qquad (R_\theta)$$

Question 3:

Suppose two sellers $i = 1, 2$ possess one unit of an indivisible good each. They have valuations $\theta_i^s \in [0, 1]$ for the good, and face a buyer of unknown type θ^b, where θ^b is the buyer's valuation for the good. θ^b is uniformly distributed over $[0, 1]$. The sellers post prices p_i for the good simultaneously, after which the buyer decides which offer (if any) to accept. If the buyer is indifferent between two offers, he randomizes between them with equal probability. Let $q(p)$ be the prior probability that the buyer will accept price p. In the absence of competition, seller i's utility from posting price p can be written as $q(p)[p - \theta_i^s]$. The buyer's utility is $\theta^b - p$ if he accepts an offer of p and 0 otherwise.

a) Assume that the sellers can observe each other's valuations, and that $\frac{1+\theta_i}{2} > \theta_j$, where $\theta_j = \max\{\theta_i, \theta_j\}$. What condition on the parameters of the game is necessary and sufficient for the existence of a Bayes Nash equilibrium of this game? Describe the equilibrium when it exists.

b) Now suppose that buyers are constrained to post prices p_i within a finite set $P = \{p_1, ..., p_K\}$ such that $p_k - p_{k-1} = \Delta > 0$, $\forall k > 1$, and $\theta_i^s \notin P$. Assume Δ is small. If you impose the condition you derived in a), what is the unique equilibrium of this game now? Explain. Then show that this game always has a unique Bayes Nash equilibrium. (To prove uniqueness, you may take Δ to be as small as you like, but not zero).

c) Retain the assumptions of b), including identical valuations, but now suppose that the sellers post prices sequentially: Seller 1 first posts a price p_1, which is observed by Seller 2 before he posts p_2. The buyer then observes both prices and chooses one offer (or none). What is the unique SPNE of this game? As you let Δ go to 0, what happens to the equilibrium outcome? Compare with the outcome of the classic Stackelberg game and discuss.

d) Now suppose that the buyer's valuation for the good is known to be 1, but that the sellers' valuations are unobservable. They each take the other's valuation to be uniformly distributed over $[0, 1]$. Sellers again post prices simultaneously, and they are no longer constrained to post prices in P_i. The buyer then chooses which offer to accept. Assume that sellers' strategies are linear functions of their valuations. What is the symmetric Bayes-Nash equilibrium? (In a symmetric BNE, sellers' strategies have the same functional form).

e) Compare the payoffs to sellers of type θ_i^s and a buyer of type 1 from the games you analyzed in a) and d). Explain.

a) is just a standard Bertrand game, and b) is the discrete variation of it. This part is straightforward. For c), note that if player 1 posts any price but $\bar{p} = \min p : p \in P$ *and* $p \geq \theta_i^s$*, he will end up with zero profit; but his profit is positive if he posts* $\bar{p}$*. Hence,* *that is the unique equilibrium outcome for both, which as* $\Delta \to 0$ *gets as close as you like to the outcome of the Nash Bertrand. No first-mover advantage in this game, unlike the Cournot Stackelberg. Finally note that d) is basically just an auction, with the only difference to the standard auction being that the buyer accepts the lowest possible offer, i.e. a seller wins if his bid is lower than the other's. In a symmetric linear NE, this leads to bidding functions* $b(\theta_i) = \frac{1}{2} + \frac{1}{2}\theta_i$*. For e), for simplicity fix* $\theta_i^s = \theta$, $i = 1, 2$*. In the classic Bertrand game in a), sellers' payoff is 0, the buyer of type 1 buys for sure and has payoff of* $1 - \theta$*. In the auction, the buyer pays* $\frac{1}{2} + \frac{1}{2}\theta$ *and achieves a payoff of* $\frac{1}{2}(1-\theta)$*, while the seller's payoff has increased to* $\frac{1}{2} - \frac{1}{2}\theta$*, which is strictly positive for all* $\theta < 1$*. This comes about as a result of informational rents: The uncertainty regarding sellers' valuations means rents have to be given in order for sellers to truthfully reveal their valuation. Thus, the uncertainty takes the edge off the Bertrand competition, and the buyer can no longer extract all rents from the market.*