Game Theory Fall 2003

Problem Set 6

[1] Let us study first-price auctions with independent private valuations (following F-T). The arguments used in this problem set are applicable to many other instances involving incomplete information (for some other examples, see F-T).

There are two bidders, and a single good is to be auctioned using sealed bid (first price). Each player has a private valuation for the good, θ_i, which we assume comes from the compact interval $\Theta \equiv [\underline{\theta}, \bar{\theta}]$. Each θ_i is drawn iid from the same density g over Θ (call the cdf G). The auctioneer has a reservation price $r \geq \underline{\theta}$, below which the good will not be sold. Player i learns her valuation, then bids b_i. The higher bid wins. Ties are broken equi-probably. Both players are risk-neutral and their utility function is simply given by $\theta - b$ if they win, or 0 if they lose.

(a) Precisely describe the Bayesian game that is involved here. In writing the payoff function, it will be useful to view each player i's bidding strategy as a distribution function G_i over her bid, as far as the other player is concerned.

Note: G_i does *not* have a density at this stage (i.e., we cannot assume that)! Indeed, there may be mass points at particular bids. To proceed further, we will use the notation $g_i(b)$ to refer to probability mass induced by G_i at any bid level b. And we will use the notation $G_i(b_)$ to denote the probability that i's bid is *strictly less* than b. [Later, when we prove that there are no mass points, you can go back to the usual interpretation.]

So the payoff to a type θ_i from making a bid b_i is just

$$F_i(\theta_i, b_i) = (\theta_i - b_i)G_j(b_i) + \frac{\theta_i - b_i}{2} g_j(b_i)$$

(as you should have verified in part (a) already).

(b) Use a revealed preference argument to show that b_i is a nondecreasing function of θ_i. That is, pick two values $\theta < \theta'$ and let b and b' be corresponding bids. By optimality,

$$F_i(\theta, b) \geq F_i(\theta, b')$$

and

$$F_i(\theta', b') \geq F_i(\theta', b)$$

Write these out fully, add both sides, cancel common terms, and examine carefully to obtain the result.

(c) Next, we show that except possibly at r (the reservation value), there are no mass points in the bid function. Put another way, we must show that the bid function $b_i(\theta_i)$ is *strictly* increasing for all $\theta_i > r$. [Note: the argument in part (b) does *not* establish this.] To do this, we must use the full power of an equilibrium argument (not just the optimality argument in part (b)); proceed as follows.

Suppose that $g_i(b_i) > 0$ for some $b_i > r$. Then first prove that there exists $\epsilon > 0$ such that person j will never bid in the interval $[b_i - \epsilon, b_i)$. Then complete the proof by showing that in this case, a bid of b_i is not optimal for player i.

(d) Next, show that there cannot be any "gaps" in the range of i's bids, or equivalently, that the bid function must be *continuous*. This uses the same style of argument as in part (c): if there is a gap in i's bid, first, prove something about j's bid, and then return to i to complete a proof by contradiction.

(e) Prove that the *maximum* bids are the same: $b_i(\bar{\theta}) = b_j(\bar{\theta})$.

(f) To proceed further, first let's invert the bid functions over the range $(r, \bar{b}]$, where $\bar{b}$ is the common maximum bid (the previous steps allow us to do this — why?). Call the inverted functions ϕ_i and ϕ_j. Make sure you've understood what the inverses mean: when player i bids b, this means she has a valuation of $\phi_i(b)$.

(g) Noting that the ϕ's are differentiable almost everywhere because they are monotonic, show that the following differential equation is satisfied for each i:

$$G(\phi_j(b)) = [\phi_i(b) - b]g(\phi_j(b))\phi_j'(b). \tag{1}$$

Idea: let i choose the bid b to maximize expected profits. Write down the first order condition and notice that the appropriate valuation θ_i for i must be $\phi_i(b)$.

Equivalently, note that (1) can be written as

$$G_j(b) = [\phi_i(b) - b]g_j(b), \tag{2}$$

where you will recall that G_j and g_j are the induced distribution and density (defined almost everywhere) over person j's bid.

(h) The rest of the analysis is technical but not difficult. See FT pp. 224–225 for the details.

[2] Here is another example of a Bayesian game, which studies the implications of unanimity rules in juries. The idea of unanimity is, of course, to guard as much as possible against the possibility of convicting an innocent individual. In the model that we construct below, some doubt is thrown on this assertion. In this exercise, you will have not only another application of Bayesian types but will also use the concept of "pivotality", something that is used often in political economy models and elsewhere.

(I follow Osborne's exposition in his new book, *Game Theory*).

There are n individuals in a jury, and a defendant, who is either guilty G or Blameless B (no one knows this, but everybody's common prior is that he is G with probability π). If he is guilty, each juror receives an iid signal, which indicates guilt with probability p (a g-signal) or innocence with probability $1 - p$ (a b-signal). Assume $p > 1/2$. If he is blameless, then the same signals are received, this time with probability $1 - q$ (for g) and q (for b), where $q > 1/2$. Each juror can vote C[onvict] or [A]cquit.

The defendant is convicted *if and only if* everyone votes to convict.

Assume that every juror has the same payoffs, which are:

0 if a truly guilty person is convicted or a truly blameless person is acquitted.

$-z$ if an innocent person is convicted.

$-(1-z)$ if a guilty person is acquitted.

(a) Set this up as a Bayesian game.

(b) Show that z can be viewed as a measure of "resistance to conviction" in the following sense: if any juror assesses the defendant to be guilty with probability r, then acquital is weakly better than conviction iff $r \leq z$.

(c) Suppose that $n = 1$. Prove that she makes different decisions depending on her received signal (i.e., C if she sees g, A if she sees b) if and only if (neglecting weak vs strict inequalities)

$$\frac{(1-p)\pi}{(1-p)\pi + q(1-\pi)} \leq z \leq \frac{p\pi}{p\pi + (1-q)(1-\pi)}.$$

(d) Now let us see under what conditions a juror would use her own signal when there are n jurors. Let us assume for this step that all other jurors except for me are using their signal in the way mentioned in part (c). Now notice that my own vote is only salient or *pivotal* when all other jurors are voting C (otherwise the outcomne is acquittal anyway). This makes me think about the circumstances in which all other jurors are in fact voting C — they must all have a guilty signal! Now suppose I do receive a b-signal. Show that I will use it — i.e., vote to acquit in this case — only if

$$z \geq \frac{1}{1 + \frac{q}{1-p}\frac{1-\pi}{\pi}\left(\frac{1-q}{p}\right)^{n-1}}$$

and show that the RHS of this inequality converges to 1 as $n \to \infty$. Interpret this result. How is pivotality (plus our assumption that other jurors are using their signals) driving this result?

Just for fun, and to see how quickly signals may start getting ignored under these presumptions, geta calculator and calculate the value of the RHS when $n = 12$, $\pi = 0.5$, and $p = q = 0.8$ (pretty reasonable parameters).

(e) Step (d) essentially tells us that for n large, there is no Nash equilibrium in which every juror votes strictly according to her signal. What about mixing? Prove that there is no symmetric Nash equilibrium in which jurors mix if they get a g-signal. [Hint: if they do so, they must be indifferent when they get a g-signal, so must strictly prefer to acquit when they get a b-signal. Now show that same argument as in part (d) applies.]

(f) So the only kind of symmetric Nash equilibrium involves C with probability β when the signal is b and announcing C for sure when the signal is g. Compute this equilibrium and show that it must have the property that

$$\left(\frac{1-q+q\beta}{p+(1-p)\beta}\right)^{n-1} = \frac{(1-p)\pi(1-z)}{q(1-\pi)z}.$$

[To show this, note that we must have indifference when a b-signal is received. Thus in the pivotal case corresponding to that signal we must have indifference as well. That is, by part (b),

$$\Pr(G| \text{ signal is } b \text{ and } n-1 \text{ votes for } C\ ) = z.$$

Now expand this expression.

(g) Prove that β goes to 1 as $n \to \infty$ (even the innocent signals generate a conviction vote with high probability).

[3] Here is some more practice along Carlsson-Van Damme and Morris-Shin lines of reasoning. Consider the following two-player game:

	L	R
U	$a+\theta, a+\theta$	$0, 0$
D	$0, 0$	$b-\theta, b-\theta$

where a and b both lie strictly between 0 and 1, and θ is a random variable distributed uniformly on $[-1, 1]$. Use the Carlsson-van Damme / Morris-Shin construction: θ is observed with some uniform noise on $[\theta - \epsilon, \theta + \epsilon]$, where ϵ is a tiny positive number. The noise is iid across the two players.

[a] Solve the equilibrium strategy of the perturbed game using the techniques studied in class. Find the limit value of the switch point θ^* as $\epsilon \to 0$ and evaluate this limit relative to the values of a and b.

By the way, take special note of this: in the Morris-Shin world, a lot of the argument works because of reasoning like this: player i thinks that player j thinks that player k thinks that ... But in this model there are only two players! Explain why the above sort of reasoning still matters.

[b] Apply the same logic to the game

	L	R
U	θ, θ	$\theta, 0$
D	$0, \theta$	$4, 4$

where θ is a random variable on some interval $[\underline{\theta}, \bar{\theta}]$, with $\underline{\theta} < 0$ and $\bar{\theta} > 4$. Is it true that (as $\epsilon > 0$) the critical switch point involves the Pareto-dominant equilibrium being played? Is this in contrast to the game of part [a], and why or why not?