Game Theory Fall 2003

Problem Set 6

[1] (a) Standard.

(b) We show that b_i is a nondecreasing function of θ_i. Notice that the payoff of player i when he bids b_i and his valuation is θ_i is

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

Now pick two values $\theta < \theta'$ and let b and b' be corresponding bids. Assume that $b \geq r$ (otherwise there is nothing to prove). Then it is easy to see that b'_i is also no less than r (after all, he can make the bid b and be assured of positive profit now if he was getting nonnegative profit earlier, which he was).

Next notice that if $\theta > b$ (or $\theta' > b'$), neither b nor b' will be placed at a point (say x) such that $b_j = x$ with positive probability: by bidding slightly higher you can get a discontinuously higher return (avoiding ties). And if $\theta = b$ (or $\theta' = b'$) the payoff is zero anyway. So in both cases, we can write

$$F(\theta, b) = (\theta - b)G_j(b) \text{ and } F(\theta', b') = (\theta' - b')G_j(b')$$

Now by the revealed preference argument which you've seen more than once:

$$(\theta - b)G_j(b) \geq (\theta - b')G_j(b')$$

while

$$(\theta' - b')G_j(b') \geq (\theta' - b)G_j(b)$$

Adding and canceling common terms, we have

$$(\theta - \theta')[G_j(b) - G_j(b')] \geq 0,$$

which proves that the bid function is nondecreasing.

(c) Next, we show that there are no mass points in the bid function, except possibly at r. [This in fact shows that the bid function is *strictly* increasing.] Look at the induced distribution of the bid G_i, and suppose on the contrary that G_i has an atom at some value $b_i > r$. Then I claim that there is $\epsilon > 0$ such that person j never bids in the range $[b_i - \epsilon, b_i)$. The reason is that j's bid in this this range would be dominated by j slightly raising the bid above b_i, whereupon j would make a discontinuous gain in the win probability. So there is a (small) blank region to the left of b_i where j never bids. But then all the types of i that bid b_i would be better off by cutting their bid to $b_i - (\epsilon/2)$ (no change in win probability, greater win margin). Contradiction.

(d) Next, we show that there cannot be any "gaps" in the range of i's bids, or equivalently, that the bid function must be *continuous*. Suppose not. Then at some θ the left hand limit is b_- and the right hand limit is b_+, with $b_+ > b_-$ (monotonicity of the bid function, established earlier, guarantees all this). Notice that no type of j should ever bid in the gap (b_-, b_+) (there is no gain in win probability at all and win margins are being given up). But if this is the case, then the i-types with bids slightly above b_+ (or at b_+, if there is such a type) will all gain by discretely lowering their bids to just above b_-. Contradiction.

(e) Prove that the *maximum* bids are the same: $b_i(\bar{\theta}) = b_j(\bar{\theta})$. Same sort of argument. The person with the higher maximum gains nothing in win probabilities by bidding the higher value.

(f) — (h) Explained in the problem itself.

[2] (a) Standard.

(b) Suppose that a juror assesses the defendant to be guilty with probability r. His expected payoff on conviction is then $-(1-r)z$. His expected payoff on acquittal is $-(1-z)r$. So acquittal is weakly better than conviction iff

$$(1-z)r \leq (1-r)z,$$

or iff $r \leq z$.

(c) Suppose that $n = 1$. If juror gets a guilty signal, the probability of guilt is

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

(applying Bayes' Rule) so by part (b) she should (weakly) convict iff

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

On the other hand, if she gets a b-signal, the posterior probability of guilt is

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

so by part (b) she should (weakly) acquit iff

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

So if

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

the juror will make different decisions depending on her signal.

(d) Now suppose there are n jurors. Assume that all other jurors 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 outcome 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!

For me to vote according to my signal, I must prefer to acquit if I get a b-signal. Will do so if

$$\begin{aligned} z &\geq \Pr(G|b, g, \ldots, g) \\ &= \frac{(1-p)p^{n-1}\pi}{(1-p)p^{n-1}\pi + a(1-q)^{n-1}(1-\pi)} \\ &= \frac{1}{1 + \frac{q}{1-p}\frac{1-\pi}{\pi}\left(\frac{1-q}{1-p}\right)^{n-1}}, \end{aligned}$$

which converges to 1 as $n \to \infty$. (In fact, it goes pretty fast: in the numerical example in the problem, you need $z \geq 0.999999$ for acquittal!)

Intuition: a bit like herding. In the pivotal event, everyone else is voting guilty and by assumption, they all have a guilty signal. This makes it very hard for me to follow my own signal in the face of so many guilty signals. Note that the other guilty signals don't exist — they only exist in my mind in the pivotal event — but unfortunately, it's only the pivotal event I care about.

(e) Follow the hint.

(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. Let us compute this equilibrium. Assuming all are following this strategy,

$$\begin{aligned} \Pr(G|b \text{ and } n-1 \text{ votes for } C) &= \frac{\Pr(b|G)[\Pr(C|G)]^{n-1}\Pr(G)}{\Pr(b|G)[\Pr(C|G)]^{n-1}\Pr(G) + \Pr(b|B)[\Pr(C|B)]^{n-1}\Pr(B)} \\ &= \frac{(1-p)[p+(1-p)\beta]^{n-1}\pi}{(1-p)[p+(1-p)\beta]^{n-1}\pi + q[1-q+q\beta]^{n-1}(1-\pi)}, \end{aligned}$$

and this must equal z. Now simplify.

[To complete the argument, you should check that a person with a g-signal strictly prefers, in the pivotal case, to vote for C.]

(g) By inspection. The intuition of this comes from trying to maintain indifference in the pivotal state. If n is large, then for sure have both many g and many b signals. If behavior conditional on the two remains markedly different, and yet we have $n-1$ C's, then it means that the true state is G with very high probability. So I will strictly ignore my signal, which contradicts mixing. So the behavior conditional on getting a g signal and a b signal have to converge to each other.

[3] In both the games under consideration, let A stand for the generic strategy that involves play of L (for Column) or U (for Row), and B for the generic strategy that involves play of R (for Column) or D (for Row). In both cases note that playing B is likely to be "better" under low values of the signal, so that is how we will orient the calculations.

Suppose, then, that we imagine that a player will play B if the signal is some value X or less. Let us calculate the recursion value $\psi(X)$ such that *under this assumption*, someone will play B if his signal is $\psi(X)$ or less.

These examples have the same general structure. Suppose that the signal space is located on some interval $[\ell, h]$. For signals very close to ℓ playing B is dominant. For signals very close to h, playing A is dominant. So $\psi(\ell) > \ell$ and $\psi(h) < h$. Finally, we will show that ψ is nondecreasing but has a slope strictly less than one. This yields a unique intersection x^* (which depends on the extent of the noise ϵ). By *exactly* the same arguments as in Morris-Shin (see my notes), there is a unique equilibrium of the imperfect observation game: play B iff the signal falls short of x^*. Finally, we describe x^* as $\epsilon \to 0$.

(a) In the first example, suppose that your opponent plays B if his signal is X or less. Suppose you see a signal x, and play B. if the true state is θ, the chance that your opponent plays B is just the chance that your opponent's signal falls below the threshold X, given θ. This is given by the expression

$$\max\{\frac{X - (\theta - \epsilon)}{2\epsilon}, 0\},$$

and so your expected payoff (now taking expectations over θ conditional on your signal) is

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon} (b - \theta) \max\{\frac{X - (\theta - \epsilon)}{2\epsilon}, 0\} d\theta. \tag{1}$$

Likewise, if you play A, the chance that your opponent also plays A is

$$1 - \max\{\frac{X - (\theta - \epsilon)}{2\epsilon}, 0\},$$

and so your expected payoff conditional on x is

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon} (a + \theta) \left[1 - \max\{\frac{X - (\theta - \epsilon)}{2\epsilon}, 0\}\right] d\theta. \tag{2}$$

[Above, I am integrating from $x - \epsilon$ to $x + \epsilon$. I should be worrying about the lower and upper bounds on θ if I am too close to one edge of the signal space. But we can ignore this, because we know the behavior of Ψ at the edges of the signal space without having to write down the exact expressions.]

The equality of expressions (1) and (2) give you the threshold x for which you are indifferent between A and B, under the presumption that a signal below X results in a play of B for your opponent. In other words, $\psi(X)$ is the solution (in x) to the equation

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon} (b-\theta) \max\{\frac{X - (\theta - \epsilon)}{2\epsilon}, 0\} d\theta = \frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon} (a+\theta) \left[1 - \max\{\frac{X - (\theta - \epsilon)}{2\epsilon}, 0\}\right] d\theta. \tag{3}$$

By inspecting (3) it should be obvious that $\Psi(X)$ is nondecreasing in X. What is a little less obvious is the assertion that for all $X' > X$,

$$\psi(X') - \psi(X) < X' - X. \tag{4}$$

To prove (4), let X increase to $X + \Delta$. We want to show that the required solution to (3) in x increases by strictly less than Δ. Suppose this is false, then it must be that after raising X to $X + \Delta$, a rise from the previous solution x to $x + \Delta$ still does not (weakly) bring the LHS and RHS of (3) into new equality; i.e., we have

$$\frac{1}{2\epsilon}\int_{x+\Delta-\epsilon}^{x+\Delta+\epsilon} (b-\theta) \max\{\frac{X + \Delta - (\theta - \epsilon)}{2\epsilon}, 0\} d\theta \geq \frac{1}{2\epsilon}\int_{x+\Delta-\epsilon}^{x+\Delta+\epsilon} (a+\theta) \left[1 - \max\{\frac{X + \Delta - (\theta - \epsilon)}{2\epsilon}, 0\}\right] d\theta.$$

Now make the change of variables $\theta' \equiv \theta - \Delta$. Then, after all the substitutions, we may conclude that

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon}(b-\theta'-\Delta)\max\{\frac{X-(\theta'-\epsilon)}{2\epsilon},0\}d\theta' \geq \frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon}(a+\theta'+\Delta)\left[1-\max\{\frac{X-(\theta'-\epsilon)}{2\epsilon},0\}\right]d\theta',$$

but this contradicts (3), the original relationship between X and x. So the claim in (4) is established. Now we have a unique equilibrium using exactly the same arguments as Morris and Shin.

Call this unique threshold x^*. Then, using this fixed point in (3) and noting that the "maxes" in that equation may now be dropped (why?), we have

$$\int_{x^*-\epsilon}^{x^*+\epsilon}\frac{(b-\theta)[x^*-(\theta-\epsilon)]}{2\epsilon}d\theta = \int_{x^*-\epsilon}^{x^*+\epsilon}(a+\theta)\left[1-\frac{x^*-(\theta-\epsilon)}{2\epsilon}\right]d\theta.$$

Now pass to the limit as $\epsilon \to 0$ (use L'Hospital's Rule). It is easy to see that at the limit,

$$x^* = \theta^* = \frac{b-a}{2}.$$

[b] In the second example, make the same provisional assumption: your opponent plays B if his signal is X or less. Suppose you see a signal x, and play B. if the true state is θ, the chance that your opponent plays B is just the chance that your opponent's signal falls below the threshold X, given θ. This is given by the expression

$$\max\{\frac{X-(\theta-\epsilon)}{2\epsilon},0\},$$

just as in (a), and so your expected payoff (conditional on your signal) is

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon}4\max\{\frac{X-(\theta-\epsilon)}{2\epsilon},0\}d\theta.$$

[Again, I am integrating from $x-\epsilon$ to $x+\epsilon$ because we can neglect the edges of the state space (see discussion in part (a) above).]

On the other hand, if you play A, you're guaranteed θ (whatever it may turn out to be), so your expected payoff is just x, of course.

The equality of these two expressions give you the indifference threshold x. That is, $\psi(X)$ solves the equation (in x):

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon}4\max\{\frac{X-(\theta-\epsilon)}{2\epsilon},0\}d\theta = x. \tag{5}$$

Again, you can show that $\Psi(X)$ is nondecreasing in X and has slope less than one; i.e., that (4) holds for the ψ-function here as well. [Use the same sort of argument we did above; things here are even simpler.]

Call this unique threshold x^*. Then, using this fixed point in (5) and once again noting that the "maxes" may be dropped (why?), we have

$$\frac{1}{\epsilon}\int_{x^*-\epsilon}^{x^*+\epsilon}\frac{x^*-\theta+\epsilon}{\epsilon}d\theta = x^*.$$

Now pass to the limit as $\epsilon \to 0$. It is easy to see that

$$x^* = \theta^* = 2.$$

Observe the contrast between parts (a) and (b). In (a), equilibrium selection generally tracks the Pareto-dominant equilibrium. When $a = b$, the switch point is 0 (how *could* it be anything else, by symmetry and uniqueness?), and now if a and b depart from each other, the switch point moves in the "correct" direction. For example, when, if $b > a$, B will be played more often, because the switch point is now positive.

In part (b), the switch point is $\theta = 2$ (which is about its midpoint value, given the support of θ). At this point, $(4, 4)$ is still much better than $(\theta, \theta) = (2, 2)$. Why does $(4, 4)$ have so little attractive power? It is because the play of A has "insurance" properties: if your oppoent does not play A, you still get something (in this example, you get full insurance in fact). But you get no insurance if you play B and your opponent does not. Thus the selection device not only looks at payoffs "at the equilibrium", it looks at payoffs "off the equilibrium" as well to make the selection.