Game Theory Fall 2003

Problem Set 2

[1] The action b is surely not dominated by any pure strategy. To prove that no mixed strategy dominates it either, take a mix of a and c, say, p and $1-p$. So as to dominate the first entry in b, it is obvious that $p > 3/4$. But then the last entry in b cannot be dominated by this mixed strategy.

Indeed, it is easy to see that b is a best response to the correlated belief: "2 plays L and 3 chooses matrix 1 with probability 1/2, while 2 plays R and 3 chooses matrix 2 with probability 1/2".

But b is a never-best strategy when beliefs are independent. To see this, denote by p the probability that 2 plays L, and by q the probability that 3 chooses matrix 1. Let r_i be the expected payoff from choosing action i (for $i =, a, b, c$). Then

$$r_a = 4 - 4(1-p)(1-q),$$

$$r_b = 4 - 4pq,$$

and

$$r_c = 3[pq + (1-p)(1-q)].$$

It will suffice to prove the following assertion: *whenever $r_a \leq r_b$, $r_c > r_b$*. To this end, note that if $r_a \leq r_b$, then bit a little algebra,

$$(1) \qquad 10pq - 7p - 7q + 3 \geq 0.$$

It is under this condition (1) that we are to show that $r_c > r_b$, or equivalently (a bit more algebra) that

$$(2) \qquad 10pq - 3p - 3q < 1.$$

Suppose, on the contrary, that (1) is true but (2) is false, so that for some p and q in $[0, 1]$,

$$(3) \qquad 10pq - 3p - 3q \geq 1.$$

Notice that neither p for q can be equal to one, for then (1) can't hold. Moreover, neither p nor q can equal 0, for then (3) can't be true. Therefore $0 < p, q < 1$. Now consider the problem:

$$\max 10pq - 3p - 3q$$

subject to (1). Set up the Lagrangean and note that because of the previous paragraph the first order conditions must hold with equality. You can use this to show that $p = q \equiv x \in (0, 1)$. So (1) reduces to

$$10x^2 - 14x + 3 \geq 0,$$

while the maximand becomes $10x^2 - 6x$. Now simply graph these functions to show that if (1) holds, (3) can't.

Note: There may be a less tedious way to do all this but I couldn't find it.

[2] (a) No. In general, Σ_i^{k+1} will not be a convex set. For instance, the set may contain two pure strategues that are each a best response to *different* beliefs, while any mixture of them is not a best response to *any* belief. Example (only the row player's payoffs are recorded):

	L	R
a	3	0
b	2	2
c	0	3

In this game a and c lie in Σ^1_{Row} but no mixture of them does.

(b) The idea behind using $\bar{\Sigma}_j^k$ is simply this. Just as in the previous question, there may be two strategies for player j which are each best responses and are hence in Σ_j^k, but no mix of these is a best response for j (in the previous iteration). Should this mean that player i should *not* entertain a belief that she faces, say, a 50-50 mix between a and b? The answer is no. The reason is that i may be uncertain about which of the two pure responses j is going to play. Notice that this interpretation completely divorces one interpretation of mixed stratgies (as beliefs) from the other (such strategies are objects that players choose). Player i *cannot rationally believe* that player j will use a mixed strategy, but she can be uncertain about which pure strategy j is going to play!

(c) Essentially, correlation. Player i would have to believe that players j and k somehow correlate their choice of actions. Or in the interpretation in question (b), the uncertainty in player i's mind about the choices of the others is somehow correlated. But this position is untenable if it is commonly known that each player makes her decision in full independence of the others.

(d) Consider the following recursion:

$$\Omega_i^{k+1} \equiv \{\sigma_i \in \Sigma_i | \text{ there is } \sigma_{-i} \in \prod_{j \neq i} \mathcal{M}(S_j^k) \text{ such that } \sigma_i \text{ is a best response to } \sigma_{-i}\},$$

where for each k and j, S_j^k is just the set of pure strategies in the support of Ω_j^k. We want to prove that this is the same recursion: that $\Omega_j^k = \Sigma_j^k$ for all k and j. Proceed by induction. Suppose true for k and all j. But then Ω_j^k and Σ_j^k have the same pure strategies S_j^k. It is obvious that $\mathcal{M}(S_j^k) = \bar{\Sigma}_j^k$ (you can generate the convex hull of the latter by only using the pure strategies, after all).

Now observe that you can omit the restriction $\sigma_i \in \Sigma_i^k$ in the first definition. Any other σ was not a best response to any beliefs over earlier-stage strategies of opponents, say in stage $s < k$, call it $\mathcal{M}(S_j^s)$. So certainly they will not form a best response to to any beliefs in $\mathcal{M}(S_j^k)$, which is a smaller set.

(e) Define rationalizable (mixed) strategies to be the infinite intersection of the Σ_i^k's. It is easy to see that if Σ_j^k is nonempty and compact for all j, then so is Σ_i^{k+1} for all i (just examine the recursive definition). Because $\Sigma_j^0 = \Sigma_j$ for all j, this is certainly so at stage 0 and therefore it is true at all stages. Therefore for each i, Σ_i^k forms a nested sequence of nonempty compact sets. By a standard theorem, $R_i \equiv \cap_k \Sigma_i^k \neq \emptyset$. To show that R_i contains at least one pure strategy, notice that each Σ_i^k contains at least one pure strategy. Thus we can choose an infinite sequence of pure strategies $s^k \in S_i$ such that $s^k \in \Sigma_i^k$ for each k. Because S_i is finite, at least one of the pure strategies of i must repeat itself infinitely often in this sequence. It is easy to see that any such pure strategy must lie in R_i.

[3] (a) Recall the set of rationalizable mixed strategies for each i, and let

$$S_i^* \equiv \{s_i \in S_i | \sigma_i(s_i) > 0 \text{ for some } \sigma_i \in R_i\}.$$

These are the set of rationalizable pure stratgies for i.

We first observe that $\{S_1^*, \ldots, S_n^*\}$ forms a rationalizable family. Recall from an earlier problem that for each k and i,

$$\Sigma_i^{k+1} \equiv \{\sigma_i \in \Sigma_i | \text{ there is } \sigma_{-i} \in \prod_{j \neq i} \mathcal{M}(S_j^k) \text{ such that } \sigma_i \text{ is a best response to } \sigma_{-i}\},$$

where for each k and j, S_j^k is just the set of pure strategies in the support of Σ_j^k. Because the S_j^k's are finite and nested, the recursion must end in a finite number of steps, arriving at precisely S_i^*. We then have for each i:

$$\S_i^* = \{s_i \in S_i | \text{ there is } \sigma_{-i} \in \prod_{j \neq i} \mathcal{M}(S_j^*) \text{ such that } s_i \text{ is a best response to } \sigma_{-i}\},$$

which proves the claim.

The point is, moreover, that there there is no rationalizable family that picks up actions other than the rationalizable actions. To see this, simply prove inductively that $S_i^k \supseteq \tilde{S}_i$ for all k and i.

(b) In class we defined the concept of a *dominance-proof family*, which is the corresponding notion when we are eliminating strictly dominated strategies. This is a family $\{S_1', \ldots, S_n'\}$ such that for every i,

$$S_i' = \{s_i \in S_i | \text{ there is no } \sigma_i \in \Sigma_i \text{ s.t. } f_i(\sigma_i, s_{-i}) > f(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}'\}.$$

We also showed that a strategy survives iterated elimination of strictly dominated strategies if and only if it belongs to a set in a dominance-proof family.

So the two families show very quickly the connection between rationalizability and dominance-proofness. By the supporting hyperplane theorem, the two are equivalent for all two-player games. But more generally, rationalizability *implies* dominance-proofness, not the other way around.

[4] (69.1) P is not partitional in general. For instance, suppose that ω_1 and ω_2 have the same answers to the first two questions in Q. Then it is immediate that $\omega_2 \in P(\omega_1)$, but that $P(\omega_2)$ is a strict subset of $P(\omega_1)$.

(69.2) If the true state is 20, the decision-maker knows the true number lies in the set $\{19, 20, 21\}$. If the true state is 21, the decision-maker knows the true number lies in the set $\{20, 21, 22\}$. Note that $21 \in P(20)$ but that $P(21) \neq P(20)$.

(71.1) [a] To show that $P' = P$ we must show that for every $\omega \in \Omega$, $P(\omega) = P'(\omega)$. Let $x \in P(\omega)$. Now pick any E such that $\omega \in K(E)$. This means that $P(\omega) \subseteq E$. So $x \in E$ for all such events. But

$$P'(\omega) \equiv \cap\{E \subseteq \Omega | \omega \in K(E)\}$$

by construction, which shows that $x \in P'(\omega)$.

Conversely, suppose that $x \in P'(\omega)$. Then by definition, x lies in every E for which $\omega \in K(E)$; i.e., for which $P(\omega) \subseteq E$. But $P(\omega)$ is one such set!

[b] Now begin with $K(E)$, then construct the information function $P(\omega)$, and then define a knowledge function $K'(E)$ from it. To show: $K(E) = K'(E)$ for every E.

Suppose that $\omega \in K(E)$. Then by construction, $P(\omega) \subseteq E$, because

$$P(\omega) \equiv \cap\{E \subseteq \Omega | \omega \in K(E)\}.$$

But this means that $\omega \in K'(E)$, because the latter is the collection of all ω for which $P(\omega) \subseteq E$.

Conversely, suppose that $\omega \in K'(E)$. Then $P(\omega) \subseteq E$. But then (and again by the definition of P from K) $\omega \in K(E)$.

[The assumptions (K.1)–(K.3) are needed to ensure that P is a well-defined information function starting from K.]

(71.2) Trivial. Any act feasible under (technically, measurable with respect to) a coarser information function must also be feasible when the information function is finer. Feasibility requires that $a(\omega) = a(\omega')$ whenever ω and ω' belong to the same element of the information partition P. If P' is finer (in the sense defined in the problem), then the previous sentence also guarantees that $a(\omega) = a(\omega')$ whenever ω and ω' belong to the same element of the information partition P'. So the agent can do no worse.

In Exercise 28.2, we are in a game-theoretic situation where more information not only has a direct (and positive) result, as in this problem, but also affects the actions of the opponent, which may have a negative "indirect" effect. This opens up the usual possibility of the game-theoretic paradoxes with respect to more information. In contrast, what we have here is a simple decision problem with no player interaction.

(76.1) It can certainly be common knowledge (with the same priors) that two people assign different probabilities to the same event. This is to be carefully distinguished from the statement that they assign different *but known* posterior probabilities, which we've seen *cannot* be common knowledge.

To see this, suppose that $\Omega = \{a, b\}$, and assume that $P_1(\omega) = \{\omega\}$ for each ω, while $P_2(\omega) = \Omega$ for each ω. Now suppose that each player has a prior of 50-50 on the two outcomes a and b. The event that the two players have different posteriors is the entire space. Why? The reason is that player 2 will always use a posterior of 50-50, while player 1's posterior will always be degenerate. Therefore, the event that the two have different beliefs is the whole space, which is always common knowledge.

On the other hand, it *cannot* be common knowledge at any ω that player 1's posterior probability regarding some event E is lower than that of 2's. Define this event: it is the set

$$E^* \equiv \{\omega' \in \Omega | \rho(E/P_1(\omega)) < \rho(E/P_2(\omega))\}.$$

If E^* is commonly known at some ω, then there is a self-evident set F such that $\omega \in F \subseteq E^*$. As discussed in class (and in O-R), there will be collections $\{P_1^i\}$ from 1's partition and $\{P_2^i\}$ from 2's partition such that

$$F = \cup_i P_1^i = \cup_j P_2^j.$$

Without loss of generality use indices i and j such that (a) they run over the same index sets, and (b) for each i, $P_1^i \cap P_2^i \neq \emptyset$. [This may require some eleements being repeated.] Then for each i, because $P_1^i \cap P_2^i \neq \emptyset$, and $P_1^i = P_1(\omega)$ and $P_2^i = P_2(\omega)$ for some ω in the intersection, we must have

$$\rho(E/P_1^i) < \rho(E/P_2^i).$$

Taking unions on both sides, we must conclude that

$$\rho(E/F) < \rho(E/F),$$

which is obviously absurd.

(76.2) Follow the same sort of reasoning.