Microeconomics 5, Part II: Solutions
Thursday 11th November 2004.

Question 1

(a) Loosely speaking, we say that an event E is **common knowledge** if for every $n = 1, 2, \ldots$ it is the case that [(everybody knows that)$^n E$ is the case.] Next, we say that there is **complete information** in a game G if it is common knowledge that the game to be played is actually G. For otherwise, we say that the game features **incomplete information**

(b) A dynamic (or sequential) game has **perfect information** if players move once at a time and each player -when it is his move to play- is perfectly informed of all the previous moves (including the realization of previous random events). If some moves are simultaneous but each player can observe all past moves, we say that the game has **almost perfect information** (you might also find the terms "observable actions" or "perfect monitoring").

(c) See Osborne and Rubinstein definition 14.1, p14.

(d) See Osborne and Rubinstein, definition 45.1, p 45.

(e) See Osborne and Rubinstein, definition 97.1, p 45.

Question 2

(a) False. Consider the game G in Figure 1:

	L	R
T	$-1, 1$	$1, -1$
B	$1, -1$	$-1, 1$

G

Figure 1: Matching pennies

The four action profiles (T, L), (T, R), (B, L) and (B, R) are rationalizable but are not Nash equilibrium.

(b) True. The mixed extension of a finite strategic-form game is a game with non-empty, compact, convex action spaces and quasi-concave and continuous payoff function. Hence, by the maximum principle of Berge, best-reply maps are upper hemi-continuous and non-empty valued. By quasi-concavity, best-reply maps are convex valued. Finally, it suffices to apply Kakutani fixed-point theorem to prove the existence of a Nash equilibria. Note, however, that the existence of a Nash equilibrium in pure actions is not guaranteed. For an example of a finite game with no pure Nash equilibrium, see Figure 1.

(c) False. The best-reply BR_i of player i has to map $[0,1]$ to $[0,1]$. However, for $x_{-i} > 1/2$, we have that $BR_i(x_{-i}) = \frac{1-2x_{-i}}{4} < 0$, hence $BR_i(x_{-i}) \notin [0,1]$. Therefore, the true best-reply map of player i in G is

$$BR_i(x_{-i}) = \begin{cases} \frac{1-2x_{-i}}{4} & \text{for } x_{-i} \leq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases}$$

(d) False. In any Nash equilibrium of a finitely repeated prisoner dilemma, the terminal history is $h^* = (\emptyset, (D,D), (D,D), \ldots, (D,D))$. Moreover, it is clear that the terminal history reached at the equilibrium cannot be $(\emptyset, (C,C), (C,C), \ldots, (C,C))$ as a player has a profitable deviation, say playing C at each stage but the last one, and playing D at the last stage.

(e) False. A finite extensive-form game has a subgame-perfect equilibrium. However, uniqueness is not guaranteed. For a trivial counter-example, consider an extensive-form game with a unique player, two actions L and R and the same payoff attached to the two terminal histories of this game $(\emptyset, L)$ and $(\emptyset, R)$.

Question 3

(1). We will prove that $(a) \Rightarrow (b) \Rightarrow (c) \Rightarrow (a)$.

$(a) \Rightarrow (b)$. Since a_i^* is strictly dominant, we have that

$$u_i(a_i^*, a_{-i}) > u_i(a_i, a_{-i}),$$

for all $a_i \in A_i \backslash \{a_i^*\}$, for all $a_{-i} \in A_{-i}$. Consider any mixed action $\alpha_i \in \Delta(A_i)$ with $\alpha_i \neq \alpha_i^*$. We clearly have

$$u_i(\alpha_i, a_{-i}) = \sum_{a_i \in A_i} \alpha(a_i) u_i(a_i.a_{-i}) < u_i(a_i^*, a_{-i}),$$

for all $a_{-i} \in A_{-i}$, hence the desired result.

$(b) \Rightarrow (c)$. Since a_i^* strictly dominates every mixed action $\alpha_i \neq \alpha_i^*$, we have

$$u_i(a_i^*, a_{-i}) > u_i(\alpha_i, a_{-i}),$$

for all $\alpha_i \in \Delta(A_i), \alpha_i \neq \alpha_i^*$, for all $a_{-i} \in A_{-i}$. It follows that

$$\sum_{a_{-i} \in A_{-i}} \mu^i(a_{-i}) u_i(a_i^*, a_{-i}) > \sum_{a_{-i} \in A_{-i}} \mu^i(a_{-i}) u_i(\alpha_i, a_{-i}),$$

for all $\mu^i \in \Delta(A_{-i})$, for all $\alpha_i \in \Delta(A_i), \alpha_i \neq \alpha_i^*$, hence the desired result.

$(c) \Rightarrow (a)$ Since a_i^* is the unique mazimizer of $u_i(a_i, \mu^i)$ for every $\mu^i \in \Delta(A_{-i})$, it follows that

$$u_i(a_i^*, \mu^i) > u_i(a^i, \mu^i),$$

for all $a^i \in A_i \setminus \{a_i^*\}$, for all $\mu^i \in \Delta(A_{-i})$. In particular, it is true for every $\mu^i \Delta(A_{-i})$ such that $\mu^i(a_{-i}) = 1$ for all $a_{-i} \in A_{-i}$, hence the desired result.

(2). This result generalizes the one-shot deviation principle to δ-discounted infinitely repeated game. The proof is taken from Osborne and Rubinstein. Let $s := (s_i)_{i \in N}$ be a profile of strategies in the repeated game and let $(v^t)_{t=1}^{\infty}$ be the infinite sequence of payoff profile of G that s induces, let $U_i(s) = (1-\delta) \lim_{T \to \infty} \sum_{t=1}^{T} \delta^{t-1} v_i^t$, player i's payoff in the repeated game when the players adopt the strategy profile s. For any history $h = (\emptyset, a^1, \ldots, a^t)$, let

$$W_i(s, h) := (1-\delta) \sum_{k=1}^{\infty} \delta^{t-1} u_i(a^{t+k}),$$

where $(a^{t+k})_{k=1}\infty$ is the sequence of action profiles that s generates after the history h. That is, $W_i(s, h)$ is player i's payoff, discounted to period $t+1$, in the subgame following the history h when the players use the strategy s.

If a player can gain by a one-shot deviation then the strategy profile is obviously not a subgame-perfect equilibrium.

Now assume that no player has a one-shot profitable deviation from s after any non-terminal history h, but there is a history h^* after which player i can profitably deviate to s_i^*. Assume that $h^* = \emptyset$, the initial history, hence $U_i(s_i^*, s_{-i}) > U_i(s)$.

Since we consider δ-discounted payoffs, for any $\epsilon > 0$, there exists a $T < \infty$ such that any change in player i's payoffs in any period following T does not change player i's payoff in the repeated game by more than ϵ. Hence, we can assume, that there exists some T^* such that s_i^* differs from s_i only in the first T^* periods.

For any positive integer t, let $h^t = (a^1, \ldots, a^t)$ be the sequence of outcomes of G induced by the profile of strategies (s_i^*, s_{-i}) in the first t periods of the repeated game. Then since s_i and s_i^* differ only for the first t^* periods, we have

$$U_i(s_i^*, s_{-i}) = (1-\delta)\sum_{k=1}^{T^*} \delta^{k-1} u_i(a^k) + \delta^{T^*} W_i(s, h^{T^*}).$$

Now, since we have assumed that **no player has a profitable one-shot deviation**, player i cannot profitably deviate from s_i in the first period of the subgame following the history h^{T^*-1}. Thus $u_i(a^{T^*}) + \delta W_i(s, h^{T^*}) \leq W_i(s, h^{T^*-1})$, hence

$$U_i(s_i^*, s_{-i}) \leq (1-\delta)\sum_{k=1}^{T^*-1} \delta^{k-1} u_i(a^k) + \delta^{T^*-1} W_i(s, h^{T^*-1}).$$

By backward induction, we reach the conclusion that

$$U_i(s_i^*, s_{-i}) \leq W_i(s, \emptyset) = U_i(s),$$

a contradiction with our assumption that s_i^* is a profitable deviation. Therefore, in δ-discounted infinitely repeated game, the one-shot deviation principle does hold. The very loose intuition for this result is that in δ-discounted infinitely repeated game, the very "far' future does not matter too much, hence the game essentially behaves as a finitely repeated game.

(3). Let us first prove that the best-reply BR_i in any simple game with strategic complements is non-decreasing, that is, for any $x_{-i} \geq x'_{-i}$, $BR_i(x_{-i}) \geq BR_i(x'_{-i})$. ($\geq$ is the component-wise order.) By contradiction. Suppose that $x_{-i} \geq x'_{-i}$, but $BR_i(x_{-i}) < BR_i(x'_{-i})$. Since (SC) holds, we have

$$
\begin{aligned}
0 > u_i(BR_i(x'_{-i}), x_{-i}) - u_i(BR_i(x_{-i}), x_{-i}) \geq \\
u_i(BR_i(x'_{-i}), x'_{-i}) - u_i(BR_i(x_{-i}), x'_{-i}) > 0,
\end{aligned}
$$

a contradiction. The first inequality comes from the fact that $BR_i(x_{-i})$ is the unique best-reply to x_{-i} and the last inequality comes from the fact that $BR_i(x'_{-i})$ is the unique best-reply to x'_{-i}. Thus, BR_i is non-decreasing.

Finally, since each BR_i is non-decreasing, $BR : X \to X$, with $X := \prod_{i \in N} X_i$ and $BR := BR_1 \times \cdots \times BR_n$ is non-decreasing as well. A fixed-point of BR is a Nash equilibria in pure actions of G. Existence of a fixed-point is a consequence of Tarski fixed-point theorem. Note that the key difference with the usual existence theorem of Nash equilibria is that we didn't assume convexity of X_i. This result can in fact be generalized to the class of supermodular games. See Milgrom and Roberts, Econometrica, 1990.

(4). We first show that for $N = 3$, the set of Nash equilibria of a simple game with strategic complements is completely ordered. From the above question, we know that the set of Nash equilibria in pure actions of G is non-empty. If G has a unique Nash equilibrium, our claim is obviously true. Now, consider two Nash equilibria of G, $x^* = (x_1^*, x_2^*, x_3^*)$ and $x^{**} = (x_1^{**}, x_2^{**}, x_3^{**})$. Suppose that x^* and x^{**} are not comparable and that $x_1^* \geq x_1^{**}$, $x_2^{**} < x_2^{**}$ and $x_3^* \geq x_3^{**}$. It follows that $(x_1^*, x_3^*) \geq (x_1^{**}, x_3^{**})$, hence $BR_2(x_1^*, x_3^*) \geq BR_2(x_1^{**}, x_3^{**})$ from the property of best-reply of a game with strategic complements, a contradiction. The other cases are similar.

Next, consider two Nash equilibria of G, $x^* = (x_1^*, x_2^*, x_3^*)$ and $x^{**} = (x_1^{**}, x_2^{**}, x_3^{**})$ with $x^* \geq x^{**}$. We want to show that $u_i(x^*) \geq u_i(x^{**})$ for all $i \in N$. Since the two equilibria are ordered, we have $x_{-i}^* \geq x_{-i}^{**}$, hence

$$
u_i(x_i, x_{-i}^*) \geq u_i(x_i, x_{-i}^{**}),
$$

for all $x_i \in X_i$ from the property (ND). In particular, it is true for $x_i = x_i^{**}$, hence we have

$$u_i(x_i^{**}, x_{-i}^{**}) \leq u_i(x_i^{**}, x_{-i}^{*}) \leq u_i(x_i^{*}, x_{-i}^{*}),$$

where the least inequality follows from the fact that x_i^* is the best-reply to x_{-i}^*. The last inequality is strict if $x_i^* \neq x_i^{**}$.

Question 4

This problem is adapted from Hamilton and Slustky, Games and Economic Behavior, 1990. The sketch of analysis below follows Bade, Haeringer and Renou, 2004.

(1). We first carefully describe Γ. The set of players of Γ is $N = \{1, 2\}$. The set of non-terminal histories $H \setminus Z$ of Γ is composed of the initial history $\emptyset$ and the histories $(\emptyset, (X_1, X_2))$ for all $(X_1, X_2) \in \mathcal{P}([0,1]) \times \mathcal{P}([0,1])$ with $\mathcal{P}([0,1])$ the collection of singletons of $[0,1]$ together with $[0,1]$. The set of terminal histories is composed of all histories of the form $((\emptyset, (X_1, X_2)), (x_1, x_2))$ with the constraint that $x_i \in X_i$ for all $i \in N$. For all non-terminal histories, both players are active, hence $P(\emptyset) = P((\emptyset, (X_1, X_2))) = N$. Given a terminal history $((\emptyset, (X_1, X_2)), (x_1, x_2))$, the payoff to player i is $u_i(x_i, x_{-i})$.

Observe that the (proper) subgame of Γ that follows any history of the form $(\emptyset, (X_1, X_2))$ is a two-player strategic-form game $G(X)$, with for each player $i \in N$, the set of actions is X_i and the payoff function is u_i. For any such a game, denote $N(G(X))$ its set of Nash equilibria. Given our assumptions, we obviously have that $N(G(X))$ is non-empty.

A strategy s_i for player i in Γ is a map from the set of non-terminal histories to the set of actions. We have $s_i(\emptyset) \in \mathcal{P}([0,1])$ and $s_i(h) \in X_i$ for an history $h = (\emptyset, (X_i, X_{-i}))$.

Observe that since we use subgame perfection as our solution concept, a necessary condition for a profile of strategies (s_i, s_{-i}) to qualify as an equilibrium is that $s_i(h) \in N_i(G(X))$ following any history $h = (\emptyset, (X_i, X_{-i}))$. In other words, following any history $(\emptyset, (X_i, X_{-i}))$, $(s_i(h), s_{-i}(h))$ has to be a Nash equilibrium of the proper subgame $G(X)$. Note that uniqueness is not guaranteed at all.

(2). To prove the existence of a subgame perfect equilibrium, we construct a profile of strategies that is an equilibrium. Consider the game $G([0,1] \times [0,1])$ and pick a Nash equilibrium y^* of this game. Define the profile of strategies s as follows: for all $i \in N$, $s_i(\emptyset) = \{y_i^*\}$ and $(s_i(h), s_{-i}(h))$ for any non-terminal histories that is different from the initial history, $\emptyset$. Following any history h of the form $h = (\emptyset, (X_1, X_2))$, $(s_1(h), s_2(h))$ is obviously a Nash equilibrium of the subgame following h, as required by subgame perfection. Finally, since y_i^* is the (unique) best-reply of player i to y_{-i}^*, player i has no profitable deviation (given the strategy of player $-i$). It follows from this argument that all Nash equilibria of the strategic-from game $G([0,1] \times [0,1])$ are equilibrium outcomes of Γ.

(3). Consider the following profile of strategies (s_i^*, s_{-i}^*) with $s_i^*(\emptyset) = X_i^*$ and $(s_i^*(h), s_{-i}^*(h)) \in N(G(X))$ for any history $h = (\emptyset, (X_i, X_{-i}))$. We want to show that if $s_i^*(h^*) \neq BR_i(s_{-i}(h^*))$, $h^* = (\emptyset, (X_i^*, X_{-i}^*))$ for all $i \in N$, then s^* is not an equilibrium of Γ.

The proof proceeds by contradiction. Suppose that the profile of strategies s^* as defined above is an equilibrium of Γ. First, observe since for all players, $s_i^*(h^*) \neq BR_i(s_{-i}(h^*))$ following the history $h^* = (\emptyset, (X_i^*, X_{-i}^*))$, we have that $s_i^*(\emptyset) = X_i^* \neq [0,1]$ for all player $i \in N$. For otherwise, the profile of strategies $(s_i^*(h^*), s_{-i}^*(h^*))$ is not a Nash equilibrium of the subgame of Γ following the history h^*. (This subgame being the strategic-form game G.) .

Thus, X_i^* has to be a singleton for all players $i \in N$. Without loss of generality, assume that $s_i^*(\emptyset) = \{x_i^*\}$ for all player $i \in N$. Second, consider the profile of strategies (s_1, s_2^*) with $s_1(\emptyset) = BR_1(x_2^*)$ and $s_1(h) = s_1^*(h)$ for all non-terminal histories that is not the initial history.

Note that $(s_1(h), s_2^*(h))$ is clearly the unique Nash equilibrium of the subgame following the history $h := (\emptyset, (\{BR_1(x_2^*)\}, \{x_2^*\})$. As $s_1(h)$ is the best-reply in G to x_2^*, it follows that

$$u_1(s_1, s_2^*) > u_i(s_1^*, s_2^*),$$

a contradiction.

(4). We already know that all Nash equilibria of G are equilibrium outcomes of Γ, that is, choose a Nash equilibria y^* of G and define the profile of strategies s with for all player $i \in N$, $s_i(\emptyset) = \{y_i^*\}$ and $(s_i(h), s_{-i}(h))$ a Nash equilibrium of the subgame following the history h. Let x_i^* be a solution of

$$\max_{x_i \in [0,1]} u_i(x_i, BR_i(x_i)).$$

Then, the profile of strategies $s_i(\emptyset) = \{x_i^*\}$, $s_{-i}(\emptyset) = [0, 1]$ and $(s_i(h), s_{-i}(h))$ a Nash equilibrium of the subgame following the history h for all histories $h \neq \emptyset$ is obviously a subgame perfect equilibrium of Γ. Finally, it is easy to show that no other profiles of strategies are equilibria.