14.126
Fall 2003
Problem Set 1
Suggested Solutions[1]

Problem 1.

Let us deal first with game A. To determine the set of subgame-perfect equilibria we use backward induction. Conditional on player 1 moving U, player 2 has a unique best-response L. Conditional on player 1 moving D, player 2 has the strict best response R. Thus player 1 essentially chooses between the outcomes (U, L) and (D, R); her payoff is better in the first case. Hence the unique subgame-perfect equilibrium is for player 1 to play U, and for player 2 to play LR, that is L if U and R if D. The equilibrium path is then L, U.

In order to determine the set of all Nash-equilibria, we need to rewrite the game in strategic form. Observe that player 1 has two strategies, U and D, and player 2 has four, LL, LR, RL and RR, where each strategy specifies what she does depending on whether player 1 played U or D. The payoff matrix is easily calculated to be

	LL	LR	RL	RR
U	$2,1$	$2,1$	$0,0$	$0,0$
D	$4,-1$	$1,1$	$4,-1$	$1,1$

.

Assume first that player 1 is playing a pure strategy. If player 1 is playing U, then player 2 has two best responses, LL and LR. If she plays LR, then it is optimal for player 1 to play U, if she plays LL then not. Thus (U, LR) is a Nash equilibrium. However, it is still optimal for player 1 to play U as long as player 2 is mixing between LR and LL, with probability at most $1/3$ on LL. Thus this is a one-parameter continuum component of the set of Nash-equilibria.

If player 1 is playing D, then player 2 again has two best-responses, LR and RR. Similarly to the previous case, it turns out that if player 2 is mixing between LR and RR with probability at most $1/2$ on LR, then it is (at least weakly) optimal for player 1 to play D; thus this is again a one-parameter continuum set of Nash-equilibria.

It is easy to see that if player 2 is playing a pure strategy, then player 1 has a unique best-response which is again a pure strategy. Thus the only Nash-equilibria where player 2 is playing a pure strategy are such that player 1 is also playing a pure strategy, and hence they have been found above.

It remains to find those Nash-equilibria where both players are strictly mixing. But if player 1 is strictly mixing then LR is a unique (pure) best response of player 2, so there are no completely mixed equilibria.

Summing up, the set of Nash-equilibria of game A constitutes of two one-parameter continuum components.

[1] Thanks to previous TAs for the solutions to problems 1-3.

Let us now turn to game B. Player 1 has again strategies U and D; player 2 has four strategies LL, LR, RL and RR, each strategy specifying what she does depending on whether the signal she sees is s' or s''. The payoff matrix is as follows

	LL	LR	RL	RR
U	$2, 1$	$2-2\varepsilon, 1-\varepsilon$	$2\varepsilon, \varepsilon$	$0, 0$
D	$4, -1$	$1+3\varepsilon, 1-2\varepsilon$	$4-3\varepsilon, 2\varepsilon-1$	$1, 1$

.

It is easy to check that as $\varepsilon \to 0$ this payoff matrix converges to that of the previous game. Assume that ε is small but positive. To find the set of equilibria, suppose first that player 1 is playing a pure strategy. If it is U, player 2 has a unique best-response LL, but this is not an equilibrium. If it is D, player 2 has a unique best-response RR, and it is an equilibrium.

If player 2 is playing a pure strategy (and $\varepsilon < 1/5$), then for each strategy of player 2 player 1 has a unique best-response which is a pure strategy; thus any such equilibrium has already been found in the previous paragraph.

Let us see what happens if both players are strictly mixing. Since LR dominates RL, there can be no best-response mix involving RL. As to the other three possible pairs, we determine the mixing probability p player 1 needs to put on U in order to make player 2 indifferent.

Case (LL, LR).

The payoff of player 2 from LL is $p-(1-p)$, the payoff from LR is $p(1-\varepsilon)+(1-p)(1-2\varepsilon)$. These two are equal if $p = (2-2\varepsilon)/(2-\varepsilon)$.

Case (LL, RR). $p = 2/3$.

Case (LR, RR). $p = 2\varepsilon/(1+\varepsilon)$.

Now player 2 can be mixing over more than two strategies only if the corresponding values of p are equal for any pair of the strategies involved. Since for ε small the numbers $2/3$, $2\varepsilon/(1+\varepsilon)$ and $(2-2\varepsilon)/(2-\varepsilon)$ are all different, it follows that in equilibrium player 2 is only mixing over two strategies.

In an (LL, RR) mix player 1 would prefer to play D with certainty, so this cannot be an equilibrium.

For (LL, LR), the probability q player 2 needs to put on LL in order to make player 1 indifferent is $q = (1-5\varepsilon)/(3-5\varepsilon)$. For (LR, RR) the probability q player 2 needs to put on LR to make her opponent indifferent is $q = 1/(2-5\varepsilon)$.

To show that in the latter two cases we do have an equilibrium, it is also necessary that check that player 2 would not deviate to a strategy not in the support of her mix. This is easy to see.

Summing up, game B has three Nash-equilibria: a pure one (D, RR) and two strict mixes. Because there are no proper subgames, these are also the subgame perfect equilibria of the game. The first mix converges to U for sure against $1/3$ on LL, $2/3$ on LR; the second mix converges to D for sure against

1/2 on LR, 1/2 on RR as $\varepsilon \to 0$. The limits of the equilibria are indeed equilibria of the limiting game (i.e., game A); however as we have seen game A has more equilibria than these three. Thus the Nash-equilibrium correspondence is not lower hemi-continuous in this example. In particular, the subgame perfect equilibrium of game A does not have any nearby equilibrium in the nearby games. The subgame perfect equilibrium correspondence is neither lower, nor upper hemi-continuous.

The set of Nash equilibrium payoffs of game B, however, does converge to the set of Nash equilibrium payoffs of game A, that is, the Nash payoff correspondence in this example is continuous. The SPE payoff correspondence is lower but not upper hemi-continuous: the payoff vector $(1, 1)$ will not be an SPE payoff of game A.

It is known (Theorem 12.1 in FT) that in almost all games, any Nash equilibrium has the property that nearby games have nearby Nash equilibria. This condition fails here, so game A is in a sense non-generic. This may seem surprising given that the extensive form payoffs look generic (this is very loose of course); but then the normal form payoffs do not look generic at all. One morale is that generic extensive form games may have non-generic normal form representations.

Problem 2.

Claim: $K_i(E) = K_i(K_i(E))$

Pf: First we note that $K_i(E) = \{\omega | h_i(\omega) \subseteq E\} \subseteq \{\omega | \omega \subseteq E\} = E$ implies $K_i(E) \supseteq K_i(K_i(E))$. Next, suppose $\omega \in K_i(E)$. By definition of K_i, we know $h_i(\omega) \subseteq E$. Moreover, $\forall \omega' \in h_i(\omega), h_i(\omega') = h_i(\omega) \subseteq E$. Hence applying the definition of K_i, $\forall \omega' \in h_i(\omega), \omega' \in K_i(E)$. Therefore, $h_i(\omega) \subseteq K_i(E)$ and we can conclude $\omega \in K_i(K_i(E))$.■

Claim: $\neg K_i(\neg K_i(E)) \subseteq K_i(E)$.

Pf: Given $\omega \in \neg K_i(\neg K_i(E))$, we know $h_i(\omega) \nsubseteq \neg K_i(E)$, which implies $\exists \omega' \in h_i(\omega)$ s.t. $\omega' \notin \neg K_i(E)$, i.e. $\omega' \in K_i(E)$. Hence $h_i(\omega) = h_i(\omega') \subseteq E$, and thus $\omega \in K_i(E)$.■

Problem 3. (FT 14.1)

(a) We prove the first claim by contradiction. Suppose there was a Nash equilibrium where at state ω_1 each player played D every period. In that case, at information set (ω_1, ω_2), player 1 learns the true state after the first period, because in equilibrium player 2 goes D in the first state and goes C in the second. Given player 2's strategy, what will player 1 do, if he sees that his opponent played C in the first period? He then knows that the opponent is playing tit for tat. A quick examination of the payoffs proves that (conditional on the opponent indeed playing D next period) he is going to play C and D in the second and the third periods. More precisely, equilibrium requires that if player 1 is at information set (ω_1, ω_2) and sees the opponent playing C in the first period, he is going to play C next period, and if the opponent plays D next period, player 1 will play D in the final period.

However, if such is the equilibrium strategy of player 1, then in state ω_1 player 2 can do better then playing D each period (along the equilibrium path). Indeed not deviating gives a payoff of only zero; however, if he plays C the first period, then player 1 is going to play C next period (by the previous paragraph). So if player 2 goes D in the second and the third periods, his total payoff will be $-1+3+0=2$, thus it is worth for him to deviate. Therefore the strategy profile cannot be an equilibrium. It follows that there is no Nash equilibrium with both playing D in every period at ω_1.

Note: many of you argued that given the strategy of (D,D,D) to be played by player 1 in (ω_1,ω_2) he can do better by playing C in the first period, given player 2 plays (D,D,D) in ω_1. The problem with this argument is that the question asked you to rule out the (DD,DD,DD) *outcome*. The question does not ask you to show that the strategies "always play D" do not constitute a Nash equilibrium.

(b) As to the second part of the problem, the following strategy profile is a Nash equilibrium. Player 1 at (ω_1,ω_2) plays tit for tat in the first two periods and defects in the last period; at (ω_3) he plays C,C,D. Player 2 plays C,D,D at (ω_1). More formally, player 1 plays (C,CD,DD) (at each information set) and player 2 plays (C,DD,DD) at (ω_1).

Let us check that these strategies are mutual best responses. At (ω_3) player 1 is clearly best-responding to the tit for tat strategy of his opponent. It is also clear that player 1 is going to defect in the last period no matter what the state is, since D is a dominant strategy. Similarly, player 2 will defect in the third period at (ω_1).

Is player 2 best-responding at (ω_1)? He knows that player 1 is going to defect last period, so essentially he faces a tit-for tat for two periods. The unique best response to this is to cooperate in the first period and defect in the second. Thus his strategy is indeed a best response.

Player 1 at (ω_1,ω_2) has a more complicated problem. Let us suppose first that he plays C in the first period. His opponent played C in that period too, so he still has a uniform posterior on the two states. Suppose he plays C in the second period (he will certainly defect in the last period). Then his expected payoff in the rest of the game is

$$\frac{1}{2}(-1+0)+\frac{1}{2}(2+3)=2.$$

On the other hand if he plays D in the second period, then his expected payoff for the rest of the game is

$$\frac{1}{2}(0+0)+\frac{1}{2}(3+0)=1.5.$$

So it is better to play C in the second period. The argument also shows that along this path player 1's overall payoff in the game is $2+2=4$.

How much can player 1 get if he plays D in the first period? In that case he gets 3 in that period. If he plays C next period, his payoff in the rest of the game is

$$\frac{1}{2}(-1+0)+\frac{1}{2}(-1+3)=0.5.$$

If he plays D next period, then his payoff in the rest of the game is

$$\frac{1}{2}(0+0)+\frac{1}{2}(0+0)=0.$$

Thus the overall payoff he can get if he starts out playing D is $3+0.5=3.5<4$. It follows that player 1 is best-responding to the strategy of player 2. Also note that we need not check the off the equilibrium path actions, because we are not looking for a subgame-perfect equilibrium.

We have shown that the proposed strategy profile indeed constitutes a Nash-equilibrium of the game, and it also has the property that in the first period the outcome is (C, C) in every state.

Problem 4[2]

The first thing to do is to establish the result suggested in the hint: for any set X with $|X|>1$, there is no onto map $d: X \to \mathcal{N}(X)$. This is a simple variant of a famous result due to Cantor (that there is no onto map $d': X \to \mathcal{P}(X)$ where $\mathcal{P}(X)$ is the power set of X). The standard proof of the result works fine: suppose such a map d exists, and consider the 'diagonal' set $D=\{x \in X : x \notin d(x)\}$. First suppose D is nonempty: then $D \in \mathcal{N}(X)$ and so there exists $y \in X$ such that $d(y)=D$. Then if $y \in D$, by definition of D it would follow that $y \notin d(y)=D$, contradiction. And if $y \notin D$, then again it follows that $y \in d(y)=D$, contradiction. Hence no such map d exists. The only remaining possibility is that D is empty. Then $x \in d(x)$ for each $x \in X$. Now d is onto, so consider what the preimage of $\{x\}$ is: clearly it must be x itself. But then the image of d consists only of singletons, so if $|X|>1$ then d is not onto, contradiction. (Note: there are lots of variants on this argument.)

Next, suppose there is a complete possibility structure $(S^a, S^b, T^a, T^b, \nu^a, \nu^b)$. Then

1. By assumption, there is an onto map from T^a to $\mathcal{N}(S^b \times T^b)$.
2. There is an onto map (the projection map) from $\mathcal{N}(S^b \times T^b)$ to $\mathcal{N}(T^b)$.
3. There is an onto map from $\mathcal{N}(T^b$ to T^b (choose a map such that $\{x\} \mapsto x$for each x and then define the map arbitrarily for other elements of T^b).
4. By assumption, there is an onto map from T^b to $\mathcal{N}(S^a \times T^a)$.
5. Similarly to 2. above, there is an onto map from $\mathcal{N}(S^a \times T^a)$to $\mathcal{N}(T^a)$.

[2]The result of this problem comes from Brandenburger (2002): On the Existence of a 'Complete' Possibility Structure.

Composing these five maps gives an onto map from T^a to $\mathcal{N}(T^a)$, so it follows that $|T^a| = 1$. Suppose $|S^b| > 1$. Then clearly $|\mathcal{N}(S^b \times T^b)| \geq |S^b| > 1$, but this contradicts the existence of the onto map in step 1. above. Hence $|S^b| = 1$. Similarly $|S^a| = 1$.

Discussion: This non-existence result is in contrast with the existence result of Brandenburger and Dekel, who provide a construction of the (universal) type space T and show that there is a homeomorphism (in particular, an onto map) between T and $\Delta(S \times T)$. (The idea of this result is that the universal type space is indeed universal: adding beliefs about types is not necessary since beliefs, beliefs about beliefs, ... suffices.) Here S is the underlying space of uncertainty, assumed complete separable metric (hence not necessarily a singleton), but in the construction of T we restrict considerations to beliefs which are (built up of) probability measures on the Borel field of the underlying space (endowed with the weak topology). This shows first that the cardinality of $\Delta(S \times T)$ is smaller than that of $\mathcal{N}(S \times T)$, but more importantly, it suggests something about 'how big' the 'space of all beliefs' that it is useful to consider in this situation is.

(It's instructive to work out where the proof of Brandenburger and Dekel would fail if we tried to apply it unthinkingly to a definition of beliefs using possibility sets rather than probability measures. The obvious point is the absence of an analogue to Kolmogorov's existence theorem for the construction using possibility sets. I leave it up to you to consider this a little more.)