Econ 200C, Spring 2012
Problem Set #1
Grading Notes and Suggested Solutions

Comments

1. I intended to permit groups of no more than 5 students to submit answers. I apologize for not doing this.

2. Kristy graded Nageeb's probem 9 and my problem 3.

3. Below are suggested answers to Nageeb's 4, 8, and 9 and my 2-4.

4. I intend these answers to do several things: to give a sample answer to the question, to provide a bit more background to what the question was trying to demonstrate, and to be a model for your future answers.

5. On Nageeb's Question 4, note that a fairly simple computation answers the question. (Some supplied an extended computation. The longer computations may not be incorrect, but suggests lack of understanding.)

6. The point of Nageeb's Question 9 was to get a better understanding of the one-shot deviation principle. My answer below attempts to describe just what kinds of deviations you must check to verify Nash (rather that subgame perfect) equilibrium.

7. My Question 2 asked you to think about some simple inequalities that are related to repeated games. I hope that my arguments below are clear. I skimmed some answers: several people wrote down true expressions without supplying details. This would be a poor strategy on exams.

8. On my Question 3, some people gave correct answers that were not in the spirit of the question. The simplest answer to the first part was to invoke the folk theorem and then show that the set of strictly individually rational and feasible payoffs may exist if you add a strictly dominated strategy (an example would be useful). You correctly answer the "letter" of the question by (a) explicitly exhibiting a SGPNE strategy profile for the new game that is not an equilibrium profile in the original game; and (b) showing that any SGPNE strategy profile for the original game induces a SGPNE for the new game (to do this you must extend the original strategy profile so that it responds to histories that involve use of the new stage-game strategy). Next time, I will ask about the relationship between the set of subgame perfect equilibrium **payoffs** of the repeated game with stage game G' and the repeated game with stage game G.

9. Kristy graded my Question 3 and Nageeb's Q uestion 9. She awarded a maximum of 10 points for each (20 points total). For Question 3 she sometimes abbreviated comments with a code:

- A: Missing argument that the same SPE payoffs from repeated G exist for repeated G'.
- Spe: Made no argument about the size of the set of SPE payoffs in repeated G' relative to repeated G
- F-spe Argument about the size of the sets of SPE in repeated G' and repeated G not formalized
- M-ds No argument about dominance in repeated game.
- F-ds Argument about dominance in repeated game not formalized
- NG No game (G') provided
- W Argument about existence/non-existence of repeated-game dominant strategy has wrong conclusion – something missing based on math or not considering the right kind of strategies or best responses
- E Argued that the strategy of always playing the new dominated strategy can't be dominated because there is a subgame perfect equilibrium in which one player's strategy specifies the use of the nominated action in some (but not necessarily all) subgames.

Answers start here.

1. Nageeb's 4:

   Pure strategy minmax $(1, 1)$ Mixed strategy minmax $(3/4, 3/4)$. (Method: to find player one's security level, treat the game as zero sum using the given payoffs for player 1 ie:

	LZ	R
LZ	$0, 0$	$3, -3$
R	$1, -1$	$0, 0$

   and find the equilibrium. Reflect on why this works.)

   Ann gets A, where $A = (1-\delta)3 + (1-\delta)\delta + (1-\delta)\delta^2 A$, so $A = (3 + \delta)/(1-\delta^2)$. (Bob gets $B = (1 + 3\delta)/(1-\delta^2)$.)

   To see what is happening when discount factors are different, imagine $\delta_B = 0$ and δ_A is close to one. Now specify profiles in which Bob gets 3 in the first period and Ann gets three from then on. Ann's average payoff is close to three. Bob's is exactly three. But $(3, 3)$ is not in the feasible set. If players have different discount factors you can guarantee that average payoffs are in the cube bounded by highest and lowest payoffs for each player (in each dimension), but generally you can't say more.

   Nageeb's 8:

   Strategy profile: Play A in the first period, continue with C in second period if opponent plays A in first period. Play C in third period if history is (A, A), (C, C) (that is, both played A in the first period and C in the second). Otherwise, play B. Payoff for following the strategy is $(3 + 2x)/3$ (assuming $\delta = 1$). Payoff for deviation in first period is at most $4/3$. If $2x \geq 1$ deviation is not attractive. There are no profitable deviations in subgames because the equilibrium specifies that agents play NE in the second and third period following any history.

Nageeb's 9:

The feasible and set: convex hull of $\{(3, 3), (-1, 4), (4, -1), (1, 1)\}$ intersected with all points greater than equal to $(1, 1)$ (the min max is $(1, 1)$)

"best punishment": static Nash equilibrium. δ should be large enough to guarantee that

$$3 \geq (1 - \delta)4 + \delta$$

or $\delta \geq 1/3$.

If Ann plays always shirk against tit for tat she earns 4 in the first period and 1 thereafter, so average payoff is $4(1 - \delta) + \delta$. If Ann shirks in the first period and then plays tit for tat thereafter, then her payoff alternates between 4 and -1, so average payoff is $(4 - \delta)/(1 + \delta)$ for Ann. Bob's payoff alternates between -1 and 4 so is equal to $(-1 + 4\delta)/(1 + \delta)$.

When is Tit for Tat Nash: What is the optimal deviation? In checking for Nash equilibrium, there is no guarantee that the one-shot deviation principle applies. For this problem, you need to think about what deviations are relevant. Look at the problem from Player 1's point of view. Let V_i be Player 1's optimal value function assuming that Player 2 is playing tit for tat and, in the first period, Player 2 plays i ($i = E$ or S). V_i is a decision theoretic quantity (with player 2's strategy fixed, it depends only on player 1's choices). Furthermore, V_E and V_S are the only possible "continuation" values for Player 1: Either Player 1 played E in the previous period (or it is the initial period of the game), in which case Player 2 will begin by playing E or Player 1 played S in the previous period, in which case Player 2 will begin (the rest of the game) by playing S.

We have formulas for V_E and V_S:

$$V_E = \max\{(1 - \delta)3 + \delta V_E, (1 - \delta)4 + \delta V_S\}$$

and

$$V_S = \max\{(1 - \delta)(-1) + \delta V_E, (1 - \delta)1 + \delta V_S\}.$$

These are standard dynamic programming expressions. Take the first equation: Assuming that Player 2 begins by playing E, Player 1 has two choices: he can begin by playing E, if so he receives 3 in the first period and in the continuation receives the most he can get against an opponent who plays E first or Player 1 can shirk immediately, receive 4, and then play against an opponent who begins by shirking, so Player 1's continuation payoff is V_S. (Since Player 1 assumes that Player 2 is playing tit-for-tat, Player 1 knows the implications of his first action choice.) If playing tit-for-tat is Nash, then

$$V_E = 3 \geq (1-\delta)4 + \delta V_S \tag{1}$$

(this comes from the first maximization, using the assumption that E is at least as good as S). The second equation leaves two possibilities for V_S. It is either equal to $(1-\delta)(-1) + \delta V_E$ or $(1-\delta)1 + \delta V_S$. In the first case, $V_S = (1-\delta)(-1) + \delta 3$. In the second case, $V_S = 1$. To check for NE you must verify the inequality in (1) for both possible choices of V_S. That is, δ must be such that:

$$3 \geq (1-\delta)4 + \delta\left((1-\delta)(-1) + \delta 3\right)$$

and

$$3 \geq (1-\delta)4 + \delta.$$

It is straightforward to check that the first inequality requires $\delta \geq 1/4$, with the second inequality requires $\delta \geq 1/3$. Hence you need $\delta \geq 1/3$ for tit-for-tat to be NE.

Notice that the one-shot deviation is not the best deviation: If you attempted to use the one-shot deviation principle you would conclude that the Tit for Tat profile is an equilibrium for $\delta \geq 1/4$. The reason for the difference is that the proof of the one-shot deviation principle uses the property (true for SGPE but not necessarily for NE) that continuation strategies induce equilibrium behavior. In the case of subgame perfect equilibria, after you make a single deviation there is no way to profit in the continuation game by playing a non-equilibrium strategy (assuming that the original strategies induce an equilibrium on the subgame). In the case of NE (and, in particular, this example), after you deviate their may be gains from further deviations because your original strategy does not necessarily respond optimally to your opponent's strategy.

Tit for Tat not subgame perfect: Consider a subgame after Ann defected (and Bob did not). Tit for tat specifies that Bob begins by defecting and then plays tit for tat and that Ann begins by cooperating and then plays tit for tat. Hence in this subgame Ann's average payoff is $(-1+4\delta)\delta/(1+\delta)$ and Bob's is $(4-\delta)/(1+\delta)$. If Bob instead cooperated in the first period, he would receive 3 instead. So Bob is best responding only if $\delta \leq 1/4$. However, it must also be the case that neither player wishes to deviate to always defect in the first period. That is it must be $3 > 4(1-\delta) + \delta$, but this is incompatible with $\delta \leq 1/4$.

2. Variations on min max. These problems require patience and an understanding of the definitions.

Consider the following quantities ($\Delta(X)$ is the set of all probability distributions on X):

$$\max_{\sigma_i \in \Delta S_i} \min_{\sigma_{-i} \in \Delta S_{-i}} u_i(s_i, \sigma_{-i}) \tag{2}$$

and

$$\min_{\sigma_{-i} \in \Delta S_{-i}} \max_{\sigma_i \in \Delta S_i} u_i(s_i, \sigma_{-i}) \tag{3}$$

(a) What is the relationship between (2) and (3)?

(b) How do (2) and (3) if you replace (a) $\Delta(S_i)$ by S_i; (b) $\Delta(S_{-i})$ by S_{-i}; (c) $\Delta(S_{-i})$ by $\Pi_{j \neq i} \Delta(S_j)$? (Answer each part independently.)

Consider

$$\min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma).$$

This is the "min max".

In this quantity Player i can pick σ_i to BR to σ_{-i} (the inside maximization), but then the other agents can pick their strategies to hurt Player i (the outside minimization). This is the quantity that is relevant for repeated games.

This quantity is greater than or equal to

$$\max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma),$$

which is the security level. This is the "max min" that is relevant for zero-sum games. In the second quantity, the other players can pick a strategy that hurts Player i as much as possible (this is the inside minimization), but Player i can select the strategy that is the best possible under these circumstance. Note that when you solve the second problem and get σ^*, there is no guarantee that σ_i is a best response to σ^*_{-i}. To prove

$$\min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma) \geq \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma)$$

note that for any σ'_{-i} and σ_i,

$$u_i(\sigma_i, \sigma'_{-i}) \geq \min_{\sigma_{-i}} u_i(\sigma)$$

and so

$$\max_{\sigma_i} u_i(\sigma_i, \sigma'_{-i}) \geq \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma)$$

and since the left-hand side is greater than or equal to the right hand side for all σ'_{-i} it is true when you minimize with respect to σ'_{-i}, which gives the desired result.

Part (b) asks what happens if you change the domains over which the optimization takes place. If you can maximize over a larger set, the value of a problem will go up (or stay the same). If you can minimize over a larger set, the value of a problem will go down (or stay the same). If $\Delta(S_i)$ replaces S_i in (1), the value goes down (because you are maximizing over a smaller set). You can give examples in which the value goes down strictly (matching pennies). If you replace $\Delta(S_i)$ by S_i in (2), the value stays the same. Why can't the value go down? Given σ_{-i} the maximization problem in (2) is linear in the mixing probabilities that determine σ_i. Linear optimization problems (subject to linear constraints) always have boundary solutions (a pure strategy best response will do as well as any mixed strategy because the payoff to the mixed strategy is an average of payoffs from pure strategies). The outer maximization problem is not linear (think about why this is true). The answers to (b) and (c) are similar: In (b) and (c) the change shrinks the domain over which you are minimizing. This must weakly increase the value of the problem (and strictly increases it for some games in (2) but not (1)). In (c) you are replacing all correlated strategies of opponents with the smaller set of independent randomizations.

Informal moral: your opponents can create stronger punishments if they are allowed to randomize and even stronger punishments if they are allowed to correlated their randomizations.

You should construct examples the illustrate each of the strict changes that I claimed.

3. The purpose of this exercise is to get you to think about the relationship between a dominated strategy in a stage game and a dominated strategy in a repeated game. Given a game G, form a new game G' by adding a strictly dominated strategy for Player 1 such that the infinitely repeated game derived from G' has a strictly larger set of subgame perfect equilibria than the repeated game derived from G when the discount factor is sufficiently close to one. Does the repeated game derived from G' have any strictly dominated strategies? If so, find one (and prove that it is strictly dominated). If not, explain why not.

   If you add a strictly dominated strategy to G, then you could lower the individually rational payoff vector $\underline{v}$. This increases the set of strictly individually rational, feasible points and hence, by the folk theorem, enlarges the set of subgame perfect equilibrium payoffs. A strategy for the repeated game can specify that a player use a strictly dominated stage game strategy in a SGPNE of the repeated game (for example you can cooperate in the first period in the prisoner's dilemma). These strategies are not dominated. Whether the repeated game has a strictly dominated strategy depends on the example.

	Cooperate	Defect
Cooperate	$5, 5$	$0, 8$
Defect	$8, 0$	$1, 1$
Extra	$a, -40$	$-1, -40$

In the example above, "Extra" is strictly dominated in the stage game provided that $a < 5$. Adding the strategy lowers player 2's minmax, so enlarges the set of SGPE payoffs in the discounted game. If $a < 0$, the strategy of playing "extra" following all histories in the repeated game is strictly dominated (this strategy will have a negative average payoff and will be strictly dominated by playing any other pure strategy in all periods). If $a > 0$, then no strategy is strictly dominated. To see this, suppose f_1 is a strictly dominated strategy in the repeated game and it is dominated by g_1. In order for g_1 to strictly dominate f_1 there must be a history h_t and a strategy f_2 for Player 2 such that (a) The strategy profiles (f_1, f_2) and (g_1, f_2) both generate the history h_t;[1] and (b) $f_1^{t+1}(h_t) \neq g_1^{t+1}(h_t)$. Condition (a) says that f_1 and g_1 generate the same history for a while. (Notice that $t = 0$ is possible. If so, then the strategies specify different actions in the first stage of the repeated game.) Condition (b) says that eventually the strategies specify different actions. If (b) never happened, then g_1 would not dominate f_1. Now modify f_2 so that it defects in all periods following histories generated by g_1 and cooperates otherwise. In this case, a patient player one receives higher average payoff playing f_1 than g_1. So g_1 does not strictly dominate f_1.

4. The first two parts of this question should be straightforward. The final part requires a bit of thought.

   The questions that follow refer to the repeated Prisoner's Dilemma game with stage-game payoff matrix:

	Cooperate	Defect
Cooperate	$5, 5$	$0, 8$
Defect	$8, 0$	$1, 1$

   I will call this game G.

   (a) Find the smallest $\delta_0 \in (0, 1)$ such that for all $\delta > \delta_0$, there exists a subgame perfect equilibrium to the infinitely repeated game with discount factor δ that gives each player average payoff 5.

   Need $5 \geq (1 - \delta)8 + \delta 1$ (the right-hand side is the best you can get from deviating). $\delta_0 = \frac{3}{7}$.

[1] Informally, I mean that if Player 1 plays either f_1 or g_1 against f_2, then the first t periods of play agree with h_t. Formally, the history generated by (f_1, f_2) is defined inductively by $h'_0 = \emptyset, h'_1 = (f_1^1(h_0), f_2^1(h_0)), \ldots h'_k = (f_k^1(h_{k-1}), f_2^k(h_{k-1}))$. I want $h_k = h'_k$ for $k = 0, \ldots, t$. Similarly for the history generated by (g_1, f_2).

(b) How would your answer to part a change if you replaced "subgame perfect" by "Nash"?

For this game, nothing changes. This is because the way to give players maximal punishment is to play a stage-game NE. In general, you can achieve cooperative payoffs with lower discount factors for NE than for SGPE.

(c) Suppose that there is a population of $2N$ players. N players always play Row; N players always play Column. In each period, Row and Column players are paired and each of the N pairs play G. After they play G, they are paired with another player and the process continues. Assume that pairing is deterministic (each Row player cycles through the N Column players in a fixed order, and similarly for the Column players). This is the extended game.

i. Suppose that players can perfectly observe all past plays. Is it possible to construct a subgame perfect equilibrium of the extended game that gives each player average payoff 5? If so, can this be done for any $\delta > \delta_0$?
Consider the strategy of playing Cooperate provided that everyone after all histories cooperated and defect otherwise. In this case, when all past moves are observed the analysis is exactly as in the first part: same algebra, same δ_0.

ii. Repeat the previous part assuming that players can observe only the outcome of stage games in which they played. (That is, they do not observe the outcome of a stage game played by a distinct pair of players.)
Here things are different. One thing you can do is have independent behavior against each of the N opponents. That is, for each opponent you cooperate at first and until there is a defection by either you or that opponent when playing together. This will be an equilibrium if

$$5 \geq (1-\delta^N)8 + \delta^N. \tag{4}$$

This is almost like the original computation, but since you only meet a particular partner every N periods, the discount factor is changed. Hence a candidate for δ_0 is $(3/7)^{1/N}$. There are other possibilities. Suppose, following the hint, everyone plays the strategy of playing cooperate provided that every previous observation was cooperate. If you consider cheating, you will eventually be punished, but it won't be immediate.
Here is a particular matching procedure: Label the players $C1, \ldots, CN$ and $R1, \ldots, RN$. In period 1 (mod N), Ci plays Ri. In period k (mod N), Ci plays $R(i+k-1(\text{mod } N))$. Here $R1$ gets RN's partner the period after RN. If $R1$ cheats in period 1, then $C1$ cheats thereafter. If $C1$ cheats in period 2, then RN cheats

thereafter. In general $C1$ thru Ck will cheat in period $k+1$ so $R1$ will be playing against cheaters after period $N/2$. Hence you can support an equilibrium provided that

$$5 \geq (1-\delta^k)8 + \delta^k N$$

where k is the smallest integer greater than $N/2$. Since $k < N$ (for $N > 2$) this condition is easier to satisfy than (4).
Subgame Perfect Equilibrium might be different because it may not pay for player 1's opponent to defect after the first period.