Game Theory Fall 2006

Problem Set 3

[1] (a) Omitted. The question I wrote here earlier may work but it is harder than I thought . . .

(b) In a repeated game with discounting and with one-period payoff functions defined continuously over the product of finite action spaces, prove that the induced payoff function on action *paths* is continuous in the product topology.

(c) Consider a growth model in which there is an initial stock every period, y, to be divided between consumption c for that period and a capital investment k. Thus $y_t = c_t + k_t$ for each t. There are two agents, A and B. A moves in every even period and chooses (c, k) for that period; B does the same in every odd period. Outputs across time are linked by the production function $y_{t+1} = f(k_t)$, where f is some increasing, smooth, concave function with $f'(\infty) < 1$. A has a discount factor $\alpha \in (0, 1)$ and continuous utility function a, so gets infinite-horizon payoffs $\sum_t \alpha^t a(c_t)$, while the corresponding objects for B are β and $b(.)$.

Verify that the one-shot deviation principle is satisfied for this game.

(d) Consider a finite game tree with 2 players. The only twist is that player 1 receives two-dimensional vector payoffs. Strategy profiles are defined exactly as in class. Say that a strategy (σ_1, σ_2) is a subgame perfect equilibrium if at the initial node and at every subgame, player 1 cannot *unambiguously* gain; that is, there is no alternative strategy which will improve his payoffs along *both* dimensions. Show by example that such games do not satisfy the one-shot deviation principle.

[2] Consider an infinitely repeated game with a finite number of actions for each player and a common discount factor δ. Prove that if δ is close enough to *zero* then every subgame perfect equilibrium must involve the play of a static Nash equilibrium after every t-history. Show that this conclusion may be false if there are infinitely many actions available to each player.

[3] (a) Prove that if the worst punishment to player i in a repeated game is not sustained by the play of a static best response on her part, then player i's continuation payoff after the first period of the punishment must *strictly* exceed the worst punishment payoff.

(b) Provide a formal proof that the lifetime normalized payoff of a player (in any equilibrium of a repeated game) cannot drop below her security level.

[4] Consider n firms competing Bertrand with constant unit cost. There is a demand curve $D(p)$. Each firm i chooses a price p_i. The firm with the lowest price supplies the whole market. If there are two or more firms with the lowest price, then they split the market.

(a) Prove that if $n \geq 2$ there is a unique Nash equilibrium payoff outcome in which each firm makes zero profits. [Note that the equilibrium is not unique in *strategies* if $n \geq 3$; why?]

(b) Repeat this game with discount factor δ. Provide necessary and sufficient conditions on δ to sustain a collusive price p (with market-sharing). Do this for *every* p greater than unit cost.

[5] Formally establish the following properties of the "support mapping" ϕ.

[a] ϕ is isotone in the sense that if $E \subseteq E'$, then $\phi(E) \subseteq \phi(E')$.

[b] If all A_i are compact and f_i continuous, ϕ maps compact sets to compact sets: that is, if E is a compact subset of F^*, then $\phi(E)$ is compact as well.

[6] The use of continuation values (from the self-generating set of perfect equilibrium payoffs) to analyze equilibria of dynamic games has become quite popular in economics, especially among macroeconomists (see, for instance, Ljungqvist and Sargent's recent text, *Recursive Macroeconomic Theory*). A good example is the mutual insurance model studied by several authors, among them Kocherlakota (*Review of Economic Studies* 1996) which is described in Ljungqvist-Sargent Chapter 15. This problem gives you an introduction to that model.

There are two infinitely-lived agents. The incomes of agents 1 and 2 are given by $y^1(s)$ and $y^2(s)$, where s is an exogenously determined state, iid over time. There is a finite set of states S. The probability of $s \in S$ is given by $\pi(s) > 0$. Assume y^1 and y^2 are not perfectly correlated. Let $Y(s)$ denote aggregate income in state s; i.e., $Y(s) = y_1(s) + y_2(s)$. Also assume that both agents are perfectly symmetric in the sense that if some vector of incomes (a, b) has some probability, the permuted vector (b, a) has exactly the same probability.

Each agent has the same strictly concave smooth utility function $u(c)$, where c is consumption in that period (assume $u'(0) = \infty$), and the same discount factor $\delta \in (0, 1)$.

Income is completely perishable and must be consumed at that date or never.

(a) First, forget about any game theory and let us try to understand the set of first-best income sharing schemes. Begin with just one period. Imagine that you are maximizing the expected sum of utility of the two players. Show that you would divide Y equally in each state. More generally, suppose that $\lambda \in (0, 1)$ is the weight on player 1's expected utility and $1 - \lambda$ is the weight on player 2's utility. Now show that the optimal scheme $\{c^1(s), c^2(s)\}$ has the property that

$$u'(c^1(s))/u'(c^2(s))$$

is a constant over all states s. Indeed, under our assumptions, this constancy is the defining feature of all first-best (static) schemes (assuming $c^1(s) + c^2(s) = Y(s)$ for all s; i.e., no output is wasted).

(b) Now suppose that you want to do the same exercise dynamically; i.e., you want to maximize $\lambda \mathbb{E} \sum_t \delta^t u(c_t^1) + (1 - \lambda) \mathbb{E} \sum_t \delta^t u(c_t^2)$, subject to the constraint that $c_t^1 + c_t^2 \leq Y_t$ for all t. Show that the result of part (a) now extends to the description:

$$u'(c_t^1(s))/u'(c_t^2(s))$$

is a constant over all states s *and dates* t. This constancy is the defining feature of all first-best (dynamic) insurance schemes.

(c) Now for some game theory. Imagine that we are trying to "support" one of these schemes as an equilibrium. The resulting description is a repeated game. The actions are as follows. At each date, after incomes are realized and commonly observed by both agents, each agent simultaneously and unilaterally transfers some nonnegative amount to the other player (of course, one or both transfers may be zero). Formally define strategies for this game.

(d) Prove that the strategy profile in which *no* transfers are ever made in any history is a subgame-perfect equilibrium of this game, and indeed is the worst subgame perfect equilibrium of the game for either player. Let the expected lifetime utility for each player under this equilibrium be written as A (for "autarky").

(e) Of course, "better" equilibria may be supportable, with the equilibrium in (d) as a (perfect) threat. To do this, think of *consumption allocation schemes* that depend on each t-history (note that a description of a t-history should also include the *current* realization of incomes at date t). For instance, the first-best schemes studied in part (b) can be written as allocation schemes of this type (formally do so).

(f) Prove that a (possibly history-dependent) scheme is supportable as a subgame perfect equilibrium of the repeated game if and only if for every date t and every state s,

$$(1-\delta)u(c_t^j(s)) + \delta \mathbb{E} \sum_{\tau=1}^{\infty} u(c_{t+\tau}^j) \geq (1-\delta)u(y^j(s)) + \delta A,$$

for $j = 1, 2$, where $\{c_{t+\tau}^j\}$ denotes the continuation of the allocation scheme for all future dates.

(g) Confirm that every first-best scheme which yields a player strictly more than his autarkic payoff A can be supported as a subgame perfect equilibrium if δ is sufficient close to to 1. Of all such first-best schemes, which one do you think is supportable for the *least* restrictions on the discount factor?

(h) If no first-best scheme is supportable, then continuation values do well in describing what second-best schemes look like. For this, study Kocherlakota's paper.

[7] (a) For any $p \in F^*$, the convex hull of the set of one-shot payoffs, and any $\epsilon > 0$, prove that there is $\delta^* \in (0,1)$ such that for every $\delta \in (\delta^*, 1)$, there is p' in the ϵ-neighborhood of p and a *periodic action path* (one that involves only a finite number of distinct action profiles that periodically recur) that generates a normalized lifetime payoff of p'.

(b) Use this observation to formally add details to the folk theorem that ensure its validity even when there is no mixed action profile that supports the desired payoff vector.

[8] A borrower takes loans L as working capital; these are converted to output by means of a production function $F(L)$ satisfying standard assumptions. A lender advances L and specifies a repayment R. Assume that the borrower cannot be asked to repay more than the total output produced from the loan; thus a *contract* is a pair (L, R) with $R \leq F(L)$.

The borrower's payoff under a contract is $F(L) - R$, while the lender's payoff is $R - (1+r)L$, where r represents the opportunity interest rate per unit of funds advanced. Both borrower and lender have a common discount factor $\delta \in (0,1)$.

At any date the borrower can default by not repaying the loan. Let $v \geq 0$ be the per-period payoff to the borrower if the lender lends nothing to him at all.

(a) Provide two formalizations of this model, one interpretable as a repeated game and the other not. In the repeated game interpretation do assumptions [G.1] and [G.2] hold? Explain why both formalizations will lead to the same results regarding the set of supportable payoffs or paths.

(b) Assume that the same contract is offered period after period. It is *incentive compatible* if the lender gets nonnegative return and the borrower always repays. Prove that an incentive-compatible contract exists *if and only if* the following maximization problem

$$\max_{x \geq 0} [F(x) - \frac{1+r}{\delta} x]$$

has a value that's at least as large as v.

(c) We can easily extend the definition of incentive-compatibility for a *sequence* of contracts which might vary over time. Do so. Show by means of an example that it is possible to Pareto-dominate every stationary incentive-compatible contract by means of a (nonstationary) sequence of incentive-compatible contracts.

(d) [Optional] Consider a sequence of contracts that is efficient in the class of all incentive-compatible sequences. That is, there is no other incentive-compatible sequence which makes both lender and borrower better off at the initial date. Prove that allsuch sequences must converge to the same contract over time. What is this contract? (See Ray, *Econometrica* (2002).)

[9] Now for a different repeated relationship. A laborer faces two seasons of equal length in every year: a slack season with wage w_* and a peak season with wage w^*. Assume $w^* > w_*$. *No savings are possible.* The laborer has strictly concave utility function u defined on seasonal consumption and an inter-season discount factor of δ.

(a) Assuming that wages can neither be saved nor borowed upon, write down the laborer's lifetime utility.

Now suppose that an employer with a linear payoff function offers the laborer a contract (x_*, x^*), which is a vector of slack and peak payments. The contract can be committed but the laborer may default on the contract. There is no stigma associated to the default: the employer can continue to offer him contracts; if not, he gets the spot wages as before.

(b) Interpret this relationship as a repeated game. Be careful to write the payoff functions at each date; are they continuous?

(c) Even assuming that the offer is made in the slack season, there are still two constraints that must be respected for the laborer to accept and (later) honor this contract. What are they?

(d) Use the constraints in part (ii) to prove that a mutually profitable incentive-compatible contract exists if and only if

$$\delta^2 u'(w_*) > u'(w^*). \tag{1}$$

Notice that "fluctuation-aversion" — which is just $u'(w_*) > u'(w^*)$ would be enough to guarantee a mutually profitable contract had there not been enforcement constraints (in fact, it would have been enough to guarantee full smoothing of consumption). If (1) fails, why isn't *any* smoothing — however small — profitable?