Game Theory Fall 2003

Problem Set 5

[1] 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.

[2] (a) Prove that if the worst punishment to player i in a repeated game is not sustained by a static Nash equilibrium profile, then player i's continutaion payoff after the first period of the punishment must *strictly* exceed the worst punishment payoff.

(b) Prove that the lifetime normalized payoff of a player (in any equilibrium) cannot drop below her security level.

[3] (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.

[4] 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.

[5] Explain why it is easy to describe the set of all perfect equilibrium payoffs in a Bertrand game with homogeneous goods and constant marginal cost of production, but why this is harder for a Cournot game. [The next problem takes up Cournot games in some detail.]

[6] In general, calculating maximal punishments isn't an easy task. But one can make some headway if one restricts attention to symmetric strategies (in symmetric games). The purpose of this long problem is to develop some of these ideas in a symmetric game (what follows is based on Abreu [JET 1986]).

Consider a game G in which everyone has an action set A_i which is given by an interval of real numbers, unbounded above. Further, assume symmetry (the action set is the same for everyone). Assume, moreover, that payoff functions are continuous and bounded above, and symmetric in the sense that for every permutation p of the set of players $\{1, \ldots, n\}$,

$$f_i(a) = f_{p(i)}(a_p)$$

for all i and action vectors a, where a_p denotes the action vector obtained by permuting the indices of a according to the permutation p.

More conditions on the payoff function follow:

Condition 1. The payoff to symmetric action vectors (captured by the scalar a), denoted $f(a)$, is quasiconcave, with $f(a) \to -\infty$ as $a \to \infty$.

Condition 2. The best payoff to any player when all other players take the symmetric action a, denoted by $d(a)$, is nonincreasing in a, but bounded below.

A strategy profile σ is *strongly symmetric* if for every t-history $h(t)$, $\sigma_i(t)[h(t)] = \sigma_j(t)[h(t)]$ for all i and j.[1]

(a) As you can tell, Conditions 1 and 2 are set up to handle something like the case of Cournot oligopoly. Even though the action sets do not satisfy the compactness assumption, the equilibrium payoff set is nevertheless compact. How do we prove this?

[a1] First note that a one-shot equilibrium exists. To prove this, use Condition 1 and Kakutani's fixed point theorem.

[a2] This means that the set of strongly symmetric perfect equilibrium payoffs V is nonempty (why?). Now, look at the infimum perfect equilibrium payoff. Show that it is bounded below, using Condition 2. Using Condition 1, show that the supremum perfect equilibrium payoff is bounded above.

[3] Now show that the paths supporting infimum punishments indeed are well-defined, and together they form a simple strategy profile which is a SGPE.

[4] Finally, prove the compactness of V by using part [3].

(b) With the above exercise worked out, we can claim that there exists best and worst symmetric payoffs v^* and v_* respectively, in the class of all strongly symmetric SGPE. The following theorem then applies to these payoffs.

Theorem 1. *Consider a symmetric game satisfying Conditions 1 and 2. Let v^* and v_* denote the highest and lowest payoff respectively in the class of all strongly symmetric SGPE. Then*

[a] The payoff v_* can be supported as a SGPE in the following way: Begin in phase I, where all players take an action a_* such that

$$(1-\beta)f(a_*) + \beta v^* = v_*.$$

If there are any defections, start up phase I again. Otherwise, switch to a perfect equilibrium with payoffs v^*.

[b] The payoff v^* can be supported as a SGPE using strategies that play a constant action a^* as long as there are no deviations, and by switching to phase 1 (with attendant payoffs v_*) if there are any deviations.

[1] Note that the symmetry is "strong" in the sense that players take the same actions after all histories, including "asymmetric" ones.

Prove this theorem.

(c) Notice that by virtue of Theorem 1, the problem of finding the best strongly symmetric equilibrium therefore reduces, in this case, to that of finding two numbers, representing the actions to be taken in two phases.

Something more can be said about the punishment phase, under the assumptions made here.

THEOREM **2.** *Consider a symmetric game satisfying Conditions 1 and 2, and let (a_*, a^*) be the actions constructed to support v_* and v^* (see statement of Theorem 1). Then*

$$d(a_*) = v_*$$

Prove this theorem.

[7] 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 schmes, 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.

[8] Finally, to keep your interest going in coalitional games, here are a couple of examples that extend the characteristic function form described in Chatterjee *et al*. Both are taken from Ray and Vohra (*Games and Economic Behavior* 1999). The idea is that characteristic functions are inadequate in many situations because they do not capture interactions across players. Often, a better and more general tool is the (transferable utility) *partition function*, which describes the aggregate payoff to very coalition, but that could vary depending on which coalition structure that coalition lives in.

Formally, let N be a set of players. A *partition* of N is a division of N into disjoint coalitions $\pi = \{S_1, \ldots, S_m\}$ which exhaust the entire space. A partition function assigns to *every* partition π and every coalition $S \in \pi$ a number $v(S, \pi)$.

(a) Suppose that three individuals face a linear demand curve for their homogeneous product, $A - bx$, where x is the aggregate quantity sold and A, b are positive. Each has a common marginal cost of production, c. Construct the partition function for all coalition structures.

(b) Now consider a very special bargaining model in which each player can make an offer to a coalition, just as in class, but must restrict herself (for simplicity) to *equal division* of that coalition's payoffs. What is the equilibrium coalition structure?

(c) Can you re-do this example for 4 or 5 players?

(d) As another example of a partition function game, consider the provision of a public good by three symmetric agents. Suppose that the three players together can get a total payoff of 3. If one player stands alone and the other two are together, the standalone player gets 2 and the remainder get a total of 0.5. Finally, assume that if all are separate, each get 0. Describe the partition function and the equilibrium of the bargaining game.

(e) What happens in the above examples (as $\delta \to 1$) if players are not constrained to equal division but can make arbitrary offers, as in the characteristic function bargaining model studied in class?

(f) For more on this sort of model, see the Ray-Vohra paper. For the public goods model in particular, see Ray and Vohra (*Journal of Political Economy*, 2001).