Chapter 5: extensive form games

December 30, 2009

1 Definition

The general model of n-person games in extensive form is a straightforward extension of the model in sectiion 1.3 for two-person zero sumgames.

Definition

An n-person game in extensive form Γ^e is given by:

1) a set of players $N = \{1, ..., n\}$;

2) a tree (a connected graph without cycles), with a particular node taken as the root;

3) for each non-terminal node, a specification of who has the move (one of real players or "chance");

4) a partition of all nodes, corresponding to each particular player, into information states, which specify what players know about their location on the tree;

5) for each terminal node, a payoff attached to it.

Formally, a rooted tree is a pair (M, σ) where M is the finite set of nodes, and $\sigma : M \to M \cup \emptyset$ associates to each node its predecessor. A (unique) node m_0 with no predecessors (i.e., $\sigma(m_0) = \emptyset$) is the root of the tree. Terminal nodes are those which are not predecessors of any node. Denote by $T(M)$ the set of terminal nodes. For any non-terminal node r, the set $\{m \in M : \sigma(m) = r\}$ is the set of successors of r. We call the edges, which connect m with its successors, "alternatives" at m. The maximal possible number of edges in a path from the root to some terminal node is called the length of the tree.

Given a rooted tree (M, σ), the game in extensive form is specified once we label all the nodes and edges according to the following rules.

(a) Each non-terminal node (including the root) is labeled by number from $\{0, 1, ..., n\}$, where $i \in \{1, ..., n\} = N$ represents a real player in the game, and 0 represents a "nature" or "chance" player. We denote by M_i the set of nodes labeled by the player i. The interpretation is that when the game is played, we start at the root and then for each node $m \in M_i$ the player i is choosing which edge to follow from this node.

(b) The alternatives at a node labeled by the chance player 0 are labeled by numbers from $[0, 1]$, so that those numbers over all the alternatives sum to 1. They represent probabilities that chance would choose those alternatives.

(c) The alternatives at a node $m \in M_i,\ i \in \{1, ..., n\}$ are labeled by "move labels". Different alternatives at the same node are labeled with different labels.

(d) Each M_i, $i \neq 0$, is partitioned into information sets $P^i_1, ..., P^i_{k_i}$, $M_i = \bigcup\limits_j P^i_j$, $P^i_{j_1} \cap P^i_{j_2} = 0$, with the following condition: any two nodes x, y from the same information set must have the same number of successors, and the set of move labels on the alternatives at x should coincide with the set of move labels on the alternatives at y. The interpretation is that when a player i has to choose an alternative at the node $m \in M_i$, he knows in what information set he is, but he does not know at what exact node from this information set he is making his choice.

(e) Each terminal node m is labeled by a vector $u(m) = (u_1, ..., u_n)$ which specifies the payoffs for players $1, ..., n$, if the game ends at this node. This defines the payoff function $u : T(M) \to \mathbb{R}$.

The game starts at the root m_0 of the tree. For each non-terminal node m, $m \in M_i$ means that player i has the move at this node. A move for the chance player consists in choosing the successor of m randomly according to the probability distribution on the alternatives at m. A move for a real player $i \in N$ consists in picking a move label for the successor of this node. Note that when making the move, a player does not know where exactly he stands. He only knows the information set he is at, and hence the set of the move labels. Once a move label is picked, the game moves to the successor of the node m which is connected to m by the alternative with the chosen move label. The game continues until some terminal node m_t is reached. Then a payoff $u(m_t)$, attached to this node, is realized.

An important special case: When each information set of each player consists of a single node, we say that this game has "**perfect information**". This term refers to the fact that, when a player has to move, he possesses perfect information about where exactly in the tree he is.

Normal form games as extensive form games: any normal form game can be represented in extensive form, by ordering the players arbitrarily say $1, 2, \cdots, n$, have player 1 move first, after which the information set of other players "hides" 1's move, then player 2 moves, after which the information set of the remaining players hides the first two moves, etc.. In this fashion we can also represent multi stage games where at some nodes, several players move simultaneously.

Conversely there is a canonical normal form representation Γ of any extetensive form game Γ^e. A strategy for a player i is a complete specification of what move to choose at each and every information set from $P = \{P^i_1, ..., P^i_{k_i}\}$. The set of all such possible specifications is the strategy set C_i for player i in Γ. The payoff $u_i(c_1, ..., c_n)$ is the payoff to player i at the terminal node which is reached after all players have chosen all their moves according to the strategies $c_1, ..., c_n$. It is important to note that, since there are chance players in the extensive form game who make their choice at random, the game could have an uncertain outcome even when all real players use pure strategies. In this case the game could end in different terminal nodes, but we can calculate the probability of our game to end in each terminal node (given choice of strategies $c_1, ..., c_n$). Then, the payoff payoff $u_i(c_1, ..., c_n)$ will be the expected payoff according to those probabilities.

As usual, we assume that players evaluate different outcome on the basis of a VNM (expected) utility function.

As for normal form games we define the mixed strategy $s_i \in \Delta(C_i)$ for player i as a probability distribution on his set of pure strategies C_i. The best response correspondence is defined by $br_i(s_{-i})$ to be the set of strategies for player i that give him the best (expected) payoff against the vector s_{-i} of strategies of other players. A Nash equilibrium of the extensive form game Γ^e is the vector $s = (s_1, ..., s_n)$ of strategies, where each one is a best response to the others.

2 Subgame perfection

In a game in extensive form, the set of the Nash equilibria is often very big and some of those equilibria make little sense.

Consider for instance the extensive form variant of the Nash demand game (example 6 in Chapter 2) with perfect information. Demands are in cents (they divide \$1), player 1 chooses his demand x, which is revealed to player 2, who can only accept or reject it. For any integer $x, 0 \leq x \leq 100$, the pair of strategies where player 1 demands x, and player 2 rejects if $s_1 > x$, and accepts if $s_1 \leq x$, is a Nash equilibrium. But for $x \leq 50$, this equilibrium involves the unrealistic refusal of a fair share of the pie.

The key concept of subgame perfection is an important *refinement* that will eliminate many such unrealistic outcomes. We define it first, before illustrating its predictive power and its limits in a handful of examples.

We assume that our game has *perfect recall*. Thus, in the course of the game each player remembers his past moves. In particular, it implies some restrictions on the information sets. Two nodes x, y cannot belong to the same information set of the player i, if the choices in the game he made before reaching x or y allow him to distinguish between the two. For example, no game path (a path from the root to a terminal node) could contain several nodes from the same informational set.

A proper subgame of an extensive form game Γ^e is a subtree starting from some non-terminal node, with all the labels, such that any information set which intersects with the set of nodes in this subtree, is fully contained in that set of nodes. Thus, the fact that a player knows that a subgame is being played does not give him any additional information to refine his information structure.

Definition 1 *A subgame perfect equilibrium for the extensive form game Γ^e is a Nash equilibrium whose restriction to any subgame is also a Nash equilibrium of this subgame.*

in the variant of the Nas demand game just discussed, there are exactly two subgame perfect equilibria: player 1 demands 100, and player 2 accepts any demand; player 1 demands 99 and player 2 accepts any demand of 99 or less, but rejects the demand 100. Note that an extensive experimental testing of this game reveals that such a strategy typically fails, because the utility of player 2 depends on more than the amount of money he takes home.

Example 1 Consider the following extensive form game with perfect information. Player 1 decides whether to go left or right. Knowing his choice, player 2decides whether to go up or down. The payoffs are $u(left, up) = (3, 1)$, $u(left, down) = (0, 0)$, $u(right, up) = (0, 0)$, $u(right, down) = (1, 3)$.

In this game player 1 has two strategies (left and right), while player 2 has four strategies, since one has to specify for her what to do if player 1 chooses left as well as what to do if he chooses right. Thus, her strategy set is $\{(up_l, up_r), (up_l, down_r), (down_l, up_r), (down_l, down_r)\}$, where subindex l is for her choice after player 1 goes left, and subindex r is for her choice after player 1 goes right. Note that if player 2 would not know the choice of player 1 at a time she makes her own choice, then it would be the Battle of Sexes game, in which each player has just two strategies.

This game has two proper subgames, in each only player 2 is to make a move. The whole game has three Nash equilibria in pure strategies. They are $(left, (up, down))$, $(left, (up, up))$, $(right, (down, down))$. However, only first of them is subgame perfect. Player 2 would prefer the last one, where she gets 3, by threatening player 1 to choose terminal node with zero payoffs if he goes left. But it is not sustainable under the subgame perfection assumption, since if player 1 actually moves left player 2 will have strong incentive to choose the node with payoffs $(3, 1)$ and she has no way to pre-commit herself not to do it.

Theorem 2 *Any finite (i.e., based on a finite tree) game Γ^e in extensive form has at least one subgame perfect equilibrium.*

The proof is by induction in the number of proper subgames the game Γ^e has. If it has no proper subgames, then any Nash equilibrium of the corresponding normal form game will be a subgame perfect equilibrium of Γ^e. Now, consider a subgame $\Gamma^{e\prime}$ of Γ^e which has no its own proper subgames. It has (at least one) Nash equilibrium; pick up one of those. Substitute this whole subgame $\Gamma^{e\prime}$ by a new terminal node for Γ^e, located at the root of this subgame $\Gamma^{e\prime}$. Label this new terminal node with the payoffs from the chosen Nash equilibrium of Γ^e. We thus constructed a new game Γ^e_1 which has less proper subgames, and hence has a subgame perfect equilibrium vector of strategies by induction hypothesis. Now, we add to the strategies in this equilibrium for Γ^e_1 vector the specification for each player of what to do in $\Gamma^{e\prime}$, namely the prescription to play according to the Nash equilibrium we have picked for $\Gamma^{e\prime}$. It is easy to check that the resulting vector of strategies will be the subgame perfect equilibrium of Γ^e.

Theorem 3 *Any finite game Γ^e in extensive form with perfect information has at least one subgame perfect equilibrium in pure strategies. If for any player all payoffs at all terminal nodes are distinct, then this equilibrium is unique.*

It is easy to see that such subgame perfect equilibrium in pure strategies can be always found by backward induction, starting from the end (by seeing for every node, whose all successors are terminal nodes, what should be the choice there, and then proceeding by induction).

Leader-follower equilibrium

Given a two person game in normal form (S_1, S_2, u_1, u_2), the extensive form game where player i chooses his strategy s_i first, this choice is revealed to player j who then chooses s_j, is called the $i-$Leader,$j-$Follower game. When we speak below of the $i-$Leader,$j-$Follower equilibrium, we always mean its subgame perfect equilibrium, or equilibria.

Comparing the $i-$L,$j-$F equilibrium with the Nash equilibrium (or equilibria) of the initial normal form game, gives useful prediction about commitment tactics in that game. Clearly player i always prefers (sometimes weakly) the $i-$L,$j-$F equilibrium to any of the Nash equilibria. But there are no other restrictions on the comparison of i's $j-$L,$i-$F equilibrium payoff with the two above.

In two-person zero sum games with a value, or in a game with (strictly) dominant strategy, the L-F equilibrium and Nash equilibrium coincide: it does not matter if we choose strategies simultaneously and independently, or sequentially with the first choice being revealed.

In the Battle of the Sexes, in the war of attrition (example 7 chapter 2 and example 6 chapter 3), as well as in the simple Cournot duopoly of example 10 chapter 2, both players prefer to be the leader. In the former two, the leader-follower equilibria coincide with the pure strategy Nash equilibria; in the latter case i's payoff in the $i-$L,$j-$F equilibrium is larger than at the unique Nash equilibrium, whereas j's payoff is lower.

In two-person zero sum games without a value, both players obviously prefer to be follower. The same is true in the following game of timing.

Example 2 *grab the dollar*

This is a symmetrical game of timing with two players. Both functions a and b increase with $a(t) > b(t)$ for all t, and $b(1) > a(0)$. Recall that $a(t)$ is the payoff to the player who cries stop at t. If both stop at $t = 0$, thay both get $a(0)$; if they both stop at $t = 1$, they both get $b(1)$. The normal form game has a unique Nash equilibrium; the Leader-Follower equilibrium favors the Follower, but they both prefer it to the Nash equilibrium of the normal form.

A common difficulty with the interpretation of subgame perfect equilibrium selection is that it involves *imprudent* strategies.

Consider *Kalai's hat game*: a hat passes around the n players; each can put a dollar or nothing in the hat; if all do, they get back \$2 each; if one or more put nothing in the hat, all the money in the hat is lost. There are two Nash equilibria: all put \$1 or nobody does; the former is the s.p. equilibrium, but, unlike the latter, its strategies are imprudent.

The next example is a celebrated paradoxical game.

Example 3 *Selten's chain store paradox*

There are $20 + 1$ players. The incumbent meets successively the 20 *small* potential entrants. At every meeting, the following game takes place: first stage: the small firm chooses to enters or stay out; in the latter case payoffs are $(0, 100)$ to small firm and incumbent; if small firm enters, the incumbent chooses to collude or fight, with corresponding payoffs $(40, 50)$ and $(-10, 0)$ respectively. The only s.p. equilibrium is that all small firms enter, and collusion occurs every time.

Now suppose you are small firm #17 and the incumbent has been challenged 5 times and has fought every time, what do you do? It is certainly imprudent to enter! The other Nash equilibrium where the incumbent is committed to fight every period seems more plausible.

On the other hand, the s.p. equilibrium may display excessive prudence as in the following game.

Example 4 *Rosenthal's centipede game*

This is a multi-stage version of grab the dollar (example 2 above), where the pot starts empty, and grows by 1 cent every period. In odd periods, player 1 can grab half of the pot plus one cent, and leave the rest of the pot to player 2, or do nothing and let the pot grow till next period; in even periods player 2 can grab half of the pot plus one cent, and leave the rest to player 1, or do nothing and let the pot grow till next period. The game lasts for 100 periods. In the last period player 2 gets $51c$ and player 1 gets $49c$.

In the subgame perfect equilibrium, player 1 grabs $1c$ in period 1 and player 2 gets nothing. This is actually the only Nash equilibrium of the game!

3 Subgame perfect equilibrium in infinite games

When the number of stages in the game is infinite, the computation of s.p. equilibria becomes more tricky, and can lead to much indeterminacy or to a deterministic prediction. A famous example follows.

Example 5 *Rubinstein's alternating offers bargaining*

The two players divide a dollar by taking turns (starting with player 1) making offers. The first accepted offer is final. No money is handed out until an offer is accepted. Player i's discount rate is $\delta_i, 0 \leq \delta_i \leq 1$: receiving $\$x$ in period k is worth $\$x(\delta_i)^{k-1}$ in period 1. (Alternative interpretation: after each rejected offer, there is a chance $(1 - \delta)$ that the game ends with no one getting any money).

Case 1: no impatience, $\delta_1 = \delta_2 = 1$ (or no risk of the game terminating). If the number of periods is finite, whoever makes the last offer acts as the Leader in a Nash demand game, therefore keeps essentially the whole dollar. If the game never stops, infinite number of periods (and disagreement for ever yields zero profit to both players), any division $(x, 1 - x)$ of the dollar is a subgame perfect equilibrium outcome. It is achieved by the *inflexible* strategies where player 1 (resp. 2) refuses any offer below x (resp. below $1 - x$) and accepts any offer weakly above x (resp. weakly above $1 - x$), and the first offer is $(x, 1 - x)$.

Case 2: impatient players, $\delta_1 < \delta_2 < 1$

Check first that the inflexible strategies around $(x, 1-x)$ described above, form a Nash equilibrium, but not *not* a subgame perfect equilibrium. Say in his first move player 1 offers $(y, 1-y)$ where $1-x > 1-y > \delta_2(1-x)$. Player 2's inflexible strategy is to say No, however in the subgame starting in period 2 where inflexible strategies are used, player 2 cannot hope any more than $\delta_2(1-x)$, therefore No to $(y, 1-y)$ is not part of any equilibrium strategy in this subgame. The inflexible strategies are not subgame perfect because they contradict the equilibrium rationale in some out of equilibrium subgame.

We show now the equilibrium is unique, and compute the corresponding shares.

Observe first that in any s.p.eq. outcome, agreement takes place immediately. Indeed suppose for instance agreement takes place in period 2 at $(z, \delta - z)$, then player 1 can offer $(z + \frac{1-\delta}{2}, \frac{1+\delta}{2} - z)$ to player 2, a better result for both players, which player 2 should accept under subgame perfection.

Next the set of s.p.eq. outcomes can be shown to be closed, hence compact, so we can talk of the best or worst s.p. share for either agent.

In a s.p.eq. where 1 speaks first, if his offer is rejected we go to a s.p.eq. where 1 speaks second. Hence the best s.p.eq. for 1 when he speaks first is the one followed by the worst s.p.eq. of 2 in the game where 2 speaks first. Let x be 1's share in his best s.p.eq. when he speaks first, and y be player 1's share in his best s.p.eq. when he speaks second. Because the offer $1-x$ is accepted by 2 in that s.p.eq., we have

$$1 - x = \delta_2(1 - y)$$

Next consider player 2: the worst s.p.eq. for 2 when he speaks first is the one followed by the best s.p.eq. of 1 in the game where 1 speaks first. Because the offer y is accepted by 1 in that s.p.eq., we have

$$y = \delta_1 x$$

We can symmetrically look at the worst s.p.eq. share x' for 1 in the game where he speaks first, and worst s.p.eq. share y' in the game where he speaks second. Check that x', y' satisfies the same system of equations as x, y, implying $x = x'$ and $y = y'$, i.e., the s.p. equilibrium outcome is unique. When player 1 speaks first it is

$$(x, 1-x) = (\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2})$$

It remains to show that in the game where player 1 speaks first, the following strategies form a s.p.eq.:

player 1 always offers $(x, 1-x)$, rejects any offer below y, accepts any offer y or more;

player 2 always offers $(y, 1-y)$, rejects any offer below $1-x$, accepts any offer $1-x$ or more;

Our last example involves only two stages but many players. It illustrates the techniques to compute s.p.equilibria in this context.

Example 6 *durable goods monopoly*

A monopolist produces at zero cost a durable good. There are 1000 consumers, with reservation prices for the good uniformly distributed in the interval $[0, 100]$. The common discount rate of the monopolist and consumers is δ. If the monopolist can commit himself to a fixed pricing policy at the beginning of the game, his best choice is a constant price of 50. Consumers are impatient, so the upper half will buy immediately, for a monopolist profit of 25,000 and consumer surplus 12,500. However it is more realistic to assume the monopolist cannot commit ex ante for both periods; in period 2, he wants to cut his price to extract a little more surplus from the consumers who did not buy in the 1st period. But $p_1 = 50, p_2 = 25$ is not an equilibrium, because a consumer who values the object at \$51 prefers to wait for the "sale" rather than buying immediately.

Say p_1 is the price charged in period 1, and all consumers with valuation in $[\varphi(p_1), 100]$ buy in period 1; then in period 2, a regular monopoly situation, the price will be $\frac{\varphi(p_1)}{2}$ and all agents in $[\frac{\varphi(p_1)}{2}, \varphi(p_1)]$ will buy. Equilibrium conditions in period 1:

for consumers

$$\varphi(p_1) - p_1 = \delta(\varphi(p_1) - \frac{\varphi(p_1)}{2}) \Longleftrightarrow \varphi(p_1) = \frac{p_1}{1 - \frac{\delta}{2}}$$

for the monopolist

$$p_1 \text{ maximizes } (100 - \varphi(p_1))p_1 + \delta(\frac{\varphi(p_1)}{2})^2$$

hence

$$p_1 = \frac{(1 - \frac{\delta}{2})^2}{1 - \frac{3\delta}{4}} 50 < 50; \varphi(p_1) = \frac{1 - \frac{\delta}{2}}{1 - \frac{3\delta}{4}} 50 > 50$$

finally both monopoly profit and consumer surplus go up, relative to the non strategic $p_1 = 50, p_2 = 25$.

4 Other refinements of Nash equilibrium

When we represent an extensive form game in the normal form, the normal form could have multiple equilibria which are "behaviorally" the same. For example, assume that player 1 makes move two times. The first time he chooses a or b, and the second time he chooses c or d. This results in four strategies (a, c), (a, d), (b, c), and (b, d). The unique (behaviorally) mixed strategy "play a or b, with probability 1/2 each, at the first move, and play c or d, with probability 1/2 each, at the second move", can be represented in a continuum ways as a mixed strategy in normal form representation, as $p(a, c) + (1/2 - p)(a, d) + (1/2 - p)(b, c) + p(b, d)$ for any $p \in [0, 1/2]$.

Another way to view Nash equilibrium of an extensive form game is looking at its multiagent representation. Namely, assume that each player i is represented by several agents, one for each of his information sets. All those agents have the same payoffs (same as player i). Each agent acts at most once in the game — if and when the game path goes through the corresponding information set — and at the moment this agent acts he has no additional information compared with what he knew before the game started. Hence, we can regard our game as a game where all players (i.e., all agents) simultaneously and independently choose each a strategy from his strategy set (which is the set of move labels for the information set for which an agent is responsible). This game hence can be viewed as a normal form game.

The Nash equilibria of the original extensive form game can be defined as the Nash equilibria of thus constructed normal form game which is called its multiagent representation. The problem with this definition is that agents are precluded from cooperation. Thus, we get unrealistic equilibria.

Example 7 Consider an extensive form game where agent 1 first chooses a or b. Without knowing his choice, agent 2 then chooses x or w. If agent 1 has chosen b initially, then the game ends there. If agent 1 has chosen a initially, he has now to choose between y and z, without knowing the choice of agent 2. The payoffs are $(3,2)$ for (b,x), $(2,3)$ for (b,w), $(4,1)$ for (a,x,z), $(2,3)$ for (a,x,y), $(0,5)$ for (a,w,z), and $(3,2)$ for (a,w,y). It is easy to check that (b,w,z) is a Nash equilibrium of the multiagent representation of this game, but not an equilibrium of its normal form. The last follows from the fact that player 1's best response to w is not (b,z), but (a,y).

The way to deal with this is to consider as Nash equilibria of extensive form game Γ^e only those equilibria of its multiagent representation which survive as Nash equilibria of the normal representation of initial game Γ^e. These equilibria are called Nash equilibria in behavioral strategies. They always exist for finite extensive form games.

4.1 Sequential rationality

Sequential rationality is a generalization of the subgame-perfect equilibrium): the idea that the choice in each information set should be rational (i.e. a best response), given what the player believes about what are the chances for him to be at each particular node from this information set. These beliefs are assumed to be formed by Bayesian update. This idea results in the notion of sequential equilibrium (they always exist for finite extensive form games).

A **sequential equilibrium** is (s,π), a vector of behavioral strategies plus a vector of Bayesian consistent beliefs for all nodes (conditional probabilities that we are at each particular node, given that we are in the information set including this node), such that given those beliefs it is sequentially rational for the players to follow the prescribed strategies. The proper definition includes the way to define the consistency of π for the nodes that have a zero probability to be on the game path under s. It is done by assuming that there exists a sequence of "tremblings" of s, which assign a positive probability to each pure strategy and converge to s, such that the belief about the node with zero probability is the limit of the Bayesian updated beliefs for those tremblings (see below the definition of trembling hand equilibrium).

4.2 Trembling hand perfect equilibrium

Trembling hand perfection is the refinement of Nash equilibrium which applies to the normal form games. Consider $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$ with all C_i finite, $S_i = \Delta(C_i)$. A vector of mixed strategies $s \in S$ is a (trembling hand) perfect equilibrium of this game if there exists a sequence of $s^k \in S$, $k = 1, 2, ...$, such that

(1) Any s_i^k is completely mixed strategy, i.e. all pure strategies from C_i belong to its support (are used with positive probability)

(2) $\lim\limits_{k\to\infty} s_i^k(c_i) = s_i(c_i)$ for all $i \in N$, all $c_i \in C_i$ (i.e., s^k converges to s)

(3) $s_i \in \arg\max\limits_{t_i\in S_i} u_i(t_i, s_{-i}^k)$ for all $i \in N$, and all k

I.e., s is a (trembling hand) perfect equilibrium if there exists a sequence of "tremblings" (completely mixed strategies, ones which could end up in using any pure strategy with positive probability, even the most unreasonable one), such that this sequence converges to s, and that each strategy s_i in s is a best response to any of those tremblings made by all players other then i.

The following theorems we will not prove.

Theorem 4 *For any* $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$ *with all* C_i *finite there exists at least one (trembling hand) perfect equilibrium.*

Theorem 5 *If* Γ^e *is an extensive form game with perfect recall, and* s *is a trembling hand perfect equilibrium of the multiagent representation of* Γ^e*, then there exists a vector of beliefs* π*, such that* (s, π) *is a sequential equilibrium of* Γ^e*.*

Note that the existence of sequential equilibria follows from these two theorems.

5 Problems for Chapter 5

Problem 1 *leader follower equilibrium*

In each case, compare the two leader follower equilibria with the Nash equilibrium (or equilibria) of the normal form game. If the game defined earlier is among n players, simply consider the two player case.

a) variant of the grab the dollar game (example 2 chapter 4) where a and b increase and $b(t) > a(t)$ for all t.

b) in the coordination game example 8 chapter 2

c) in the public good provision game of example 20 chapter 2

d) in the war of attrition with mixed strategies, example 6 chapter 3

e) in the (mixed strategies) lobbying game of example 7 chapter 3.

Problem2 *King Solomon*

King Solomon hears from two mothers A and B who both claim the baby but only one of them is the true mother. Both mothers know who is who, but Solomon does not. However Solomon knows that the baby s worth v_1 to the true mother and v_2 to the false one, with $v_2 < v_1$. He has them play the following game.

Step 1 Mother A is asked to say "mine" or "hers". If she says "hers" mother B gets the baby and the game stops. If she says "mine" we go to Step 2. Mother B must "agree" or "challenge". If she agrees mother A gets the baby and the game stops; if she challenges, mother B pays v and keeps the baby, whereas mother A pays w. These two numbers are chosen so that $v_2 < v < v_1$ and $w > 0$.

Show that in the subgame perfect equilibrium of the game, the true mother gets the baby. What about the money?

Problem 3 *Bertrand duopoly*

Two firms are located town A and B respectively; in each town there is d units of inelastic demand with reservation price p (the same in each town); transportation cost between a and B is t. Thus we have a symmetrical game with strategy set $[0, p]$ and payoff

$$\begin{aligned} u_1(s_1, s_2) &= ds_1 \text{ if } |s_1 - s_2| \leq t \\ &= 2ds_1 \text{ if } s_1 + t < s_2; = 0 \text{ if } s_2 + t < s_1 \end{aligned}$$

(note that when t is exactly the price difference, customers does not travel; the opposite assumption would do just as well).

a) Show the game has no Nash equilibrium if $2t < p$. Compute the Nash equilibrium (or equilibria) if $p \leq 2t$.

b) Compute the Leader-Follower equilibria and show that a firm always prefers to be Follower.

Problem 4 *leader-follower equilibrium*

In this problem we restrict attention to finite two-person games (S_1, S_2, u_1, u_2) in pure strategies, such that the mappings u_1 and u_2 are one-to-one on $S_1 \times S_2$. Therefore the best reply strategies are unique, and so are the 1−L,2−F and 2−L,1−F equilibria. Denote the corresponding payoffs L_i and F_i.

Suppose $L_i = F_i$ for $i = 1, 2$. Show that the 1−L,2−F and 2−L,1−F equilibria coincide, are a Nash equilibrium, Pareto superior to any other Nash equilibrium.

Problem 5 *three way duel (Dixit and Nalebuff)*

Larry, Mo and Curly play a two rounds game. In the 1st round, each has a shot, first Larry then Mo then Curly. Each player, when given a shot, has 3 options: fire at one of the other players, or fire up in the air. After the 1st round, any survivor is given a second shot, again beginning with Larry then Mo then Curly.

For each duelist, best outcome is to be the sole survivor; next is to be one of two survivors; inthird place is the outcome where no one gets killed; dead last is that you get killed.

Larry is a poor shot, with only 30% chance of hitting a person at whom he aims. Mo has 80% accuracy, and Curly has 100% accuracy.

Compute the subgame perfect equilibrium of this game, and the equilibrium probabilities of survival.

Problem 6

Ten pirates (ranked from 10 to 1 from the oldest to the youngest) share 100 gold coins. The oldest first submits an allocation of his choice to a vote. If at least half of the pirates (including the petitioner) approves of this allocation, it is enforced. Otherwise, the oldest pirate walks away with no coin, and the same game is repeated with nine pirates, etc. How would you recommend the players to play? (Find the subgame-perfect Nash equilibrium outcomes)

Problem 7

In an extensive form game, a behavior strategy for player i specifies a probability distribution over alternatives at each information set of player i. Mixed strategy, as always, is a probability distribution over the set of pure strategies. Two strategies of player i are called equivalent if they generate the same payoff for player i for all possible combinations c_{-i} of pure strategies of other players. Prove that in a game of perfect recall, mixed and behavior strategies are equivalent.

More precisely: every mixed strategy is equivalent to the unique behavior strategy it generates, and each behavior strategy is equivalent to every mixed strategy that generates it.

Problem 8 *(difficult!)*

Prove that for a zero-sum game any Nash equilibrium is subgame perfect. More precisely, for any outcome which is the result of a Nash equilibrium strategy profile, there is a subgame perfect equilibrium strategy profile which results in the same outcome (an outcome is a probability distribution over the terminal nodes).

Problem 9 *grab a shrinking dollar*

One dollar is placed in the "pot" in period 1; its value will diminish by a discount of δ at each period (after k periods, it is worth δ^{k-1} to both players). The two players take turns, starting with player 1. When i has the move, she has 2 choices: to stop the game, in which case 40% of the pot goes to i and 60% to player j, or to let player j have the next move. The game goes on until someone stops, or if no one does both players get zero.

a) Show that if δ is small enough, the only Nash equilibrium of the game is that player 1 grabs the dollar immediately. Explain "small enough".

b) Is there any value of δ such that in some Nash equilibrium of the corresponding game, someone grabs the dollar after each player has declined to do so at least once?

c) Show that if δ is large enough, there is a subgame perfect equilibrium where player 1 does not grab the dollar, and player 2 does in the next turn. Explain "large enough".

Problem 10 *bargaining with alternating offers*

In this variant of Rubinstein's model (example 5), the only difference is that after an offer is rejected, the flip of a fair coin decides the player who makes the next offer. Successive draws are independent.

a) Assume first the players have a common discount factor δ. Find the symmetrical subgame perfect equilibrium of the game, and show it is the unique s.p. equilibrium.

b) Now we have 2 different discount factors. Compute similarly the s.p. equilibrium or equilibria.

Note: for both questions you must describe the equilibrium offer and acceptance strategies of both players.

Problem 11 *durable goods monopoly*

In the model of example 6, we now assume the good is infinitely durable and the game lasts for ever. A strategy of the monopolist is a stream of prices $(p_1, p_2, \cdots)$ and his profit is $\sum_{t=1}^{\infty} \delta^t p_t q_t$, where q_t is the quantity sold in period t. A consumer with valuation v gets utility $\delta^t(v - p_t)$ if she buys in period t.

Look for a *linear stationary* s.p. equilibrium: facing price p_t at time t, all consumers with valuation λp_t or above (if any are left) buy, others don't. facing an unserved demand $[0, v]$ at time t, the monopolist charges the price μv. Naturally the two constant λ, μ are such that $\lambda \geq 1, \mu \leq 1$.

write the equilibrium condtions resulting in a system to compute λ, μ. Solve the system numerically for $\delta = 0.9$ and $\delta = 0.5$. Deduce the optimal sequence $(p_1, p_2, \cdots)$ and discuss its rate of convergence. Compute the equilibrium profit and consumer surplus.

Problem 12 *last mover advantage in a first price auction*

In the game of Example 12 chapter 3 with *two* bidders, recall that the unique symmetrical equilibrium has a bid function $x(t) = \frac{t}{2}$, and an expected gain of $\$\frac{100}{6}$ for each player.

Suppose now player 2 has the last mover advantage: he observes player 1's bid before bidding himself. Compute the unique subgame perfect equilibrium of this game, and the corresponding expected gains of the players. Compare to the case of simultaneous bids.

Suppose next that player 2 sees player 1's bid but player 1 is unaware of this (and so he plays as in the case of simultaneous bids). Compute the corresponding expected gains of both players.