

REVIEW ON SUBGAME PERFECT EQUILIBRIA

prepared by Brent Hickman

3/13/08

1. Definitions

This handout is to get you better acquainted with the concept of subgame perfect equilibrium. I'll start with some definitions.

DEFINITION: A game in *normal form* is a simple list of 3 things: (1) the set of all players, (2) the set of all strategies for each player and (3) payoffs for each player resulting from all possible strategy profiles.

It should be noted that for any game there is one and only one normal form representation.

DEFINITION: A *Nash equilibrium* is a profile of strategies, one for each player, such that no one player can achieve a strictly higher payoff by unilaterally deviating to a different strategy. In other words, each player is responding optimally to his opponents' strategies.

DEFINITION: A game in *extensive form*, often represented by a game tree, conveys the same information in a slightly different light. The extensive form representation includes (1) a list of players, (2) a set of decision nodes with a well-defined ordering (a partial ordering, to be exact) (3) a one-to-one and onto correspondence from the set of players to the set of decision nodes and (4) a set of terminal nodes representing the payoffs to each player resulting from play along every possible branch of the game.

It should also be noted that although every game has a unique normal form representation, it may have multiple equivalent extensive form representations.

DEFINITION: *Information sets* are sets of nodes such that (1) a single player is mapped into each member of the set, (2) if the player finds herself at one of the nodes in the set, she cannot distinguish between any of them and (3) the player can distinguish between members of the set and any decision node not in the set.

DEFINITION: An *action* is a choice that a player makes within an information set. For example, if at a given information set, a player has the choice of "fight"

or “cooperate,” then we say she has two available actions there. It is important to remember though that strategies are not the same as actions.

DEFINITION: For a game in extensive form, a *strategy* is a

COMPLETE CONTINGENT PLAN

for actions to be taken *at each information set*, even information sets that will not be reached due to ones own actions.

Thus, for each player the dimensionality of a strategy will be the same as the number of information sets that the player has. As for counting the number of strategies, if a player has I information sets and n_i available actions at the i^{th} information set, then the number of strategies available to that player is given by

$$\prod_{i=1}^I n_i.$$

The main difference between the normal form and the extensive form of a game is in the way in which we think about players’ strategies. In the normal form, strategies are viewed as indivisible objects which are implemented simultaneously with all of ones opponents. When we look at games in the extensive form, we view strategies as being made up of smaller components, which are actions taken at information sets. This subtle difference makes no difference for identifying Nash equilibria of a game; however, it provides us with a useful framework for refining our solution concept in games with a multiplicity of Nash equilibria. The simplest type of refinement of the Nash equilibrium concept is subgame perfect equilibrium.

DEFINITION: A *subgame* is a collection of nodes starting at a single node and containing all of its successor nodes, such that any information set which intersects the collection is contained in the collection.

DEFINITION: A *subgame perfect equilibrium* is a profile of strategies which induces a Nash equilibrium on *every* subgame.

For finite games (*i.e.* games with finitely many players who each have finite strategy sets) the common method of finding subgame perfect equilibria is known as *backward induction*. Starting with the smallest subgame of the game tree, we determine what the Nash equilibria are for those subgame, keeping track of the actions

taken at each information set of the subgame. We then re-write the original game, replacing the smallest subgame with a terminal node having the payoffs involved in its Nash equilibrium. We then repeat the process, tracing it all the way back to the beginning node of the largest subgame of the original game and keeping track of the equilibrium actions taken at each of the information sets of each subgame.

For infinite games, we may not be able to use backward induction, as there may not exist a smallest subgame—that is, there may not exist a subgame which does not contain any other subgame. An example of this would be a game in which actions are chosen sequentially over infinitely many periods. However, we can still think about subgame perfect equilibria, by keeping in mind that any strategy profile must induce a Nash equilibrium on every subgame. In this case a strategy is a **COMPLETE CONTINGENT PLAN** which is unveiled before the game starts, which stipulates each player’s actions at every period in time. Thus, a subgame perfect equilibrium is a strategy profile for which if the game were started at any period $t \in \{1, 2, 3, \dots\}$, no player would have incentive to choose actions other than those prescribed by his equilibrium strategy.

Relation to Dynamic Programming: One special type of subgame perfect equilibrium in games played over infinitely many periods is a *Markov perfect equilibrium*. This is a subgame perfect Nash equilibrium in which players condition their own strategies only on the relevant states in each period. Thus, a Markov-perfect equilibrium is derived as the solution to a set of dynamic programs that each individual player solves. Equilibrium strategies are generated by each players’ policy function. A famous paper that applies this concept to industry dynamics in industrial organization is Ericson and Pakes (*Review of Economic Studies*, 1995). The concept of Markov-perfect equilibrium was first introduced by Maskin and Tirole (*Econometrica* 1988 (two papers) and *European Economic Review*, 1987).

2. Practice Problems

1. **The Centipede Game with Geometric Increase:** There are two players, Lloyd and Harry, deciding whether to share a sum of money. There are 5 periods of play. In the first period, Lloyd chooses between keeping \$1 all for himself or passing it along to Harry. If he passes it along, then the sum increases by a factor of $r = 1.5$. In the second round Harry faces the same decision with the new increased amount, either

keeping it all to himself or passing it along and having it increase again by a factor of $r = 1.5$. The game alternates likewise until the fifth and final period, when Lloyd decides whether to take all of the current sum for himself or to share it evenly. If he shares it, then it will increase once again by the same factor before it is split.

- a. Write down this game in extensive form (HINT: this game is commonly called the centipede game because its game tree resembles a centipede). How many information sets does each player have? How many strategies does each player have?
 - b. Find all pure strategy Nash equilibria of this game.
 - c. Find all subgame perfect equilibria of this game.
 - d. Are there any values of r for which there is a subgame perfect equilibrium where the game reaches the final period?
2. **Prisoner's Dilemma with a Twist:** Bugsy and Jimmy "the Toucan" are suspected of running a bootlegging operation out of Canada and have been picked up by the feds. They are being held in different rooms and the G-men tell each of them individually that if they rat on their partner they will get off easy. The deal is that if one of them keeps quiet while his partner rats him out, the one who squealed gets off with no prison time, while the other does 20 years. Of course, if they both rat each other out, then the feds send them both to The Rock for 15 years. If both keep quiet, then the IRS still has enough evidence to lock them both up for 5 years on tax evasion. Thus, the payoff matrix is:

	R	Q	
R	$-15, -15$	$0, -20$.
Q	$-20, 0,$	$-5, -5$	

- a. Write down two distinct extensive form representations of this game.

Now consider an alternative game in which the feds decide to have a little more fun with Bugsy and Jimmy. They split the interrogation up into two rounds in which the above game is repeated twice and players get to observe the outcome of the first round before choosing their actions in the second round. At the end of the game, the feds add the prison sentences from both rounds, just for fun.

- b. Write down this game in extensive form. How many information sets does each player have? How many strategies does each player have?

- c. Write down a payoff matrix representing the normal form of the two-stage game (HINT: the dimensions of the matrix will be determined by how many strategies each player has).
- d. Find all pure strategy Nash equilibria of this game.
- e. Find all subgame perfect equilibria of this game. Is there a subgame perfect equilibrium in which both players keep quiet in the first round?

Just as the G-men are about to play their nasty little prank, Elliot Ness walks in and recognizes Jimmy “the Toucan” from the time he bumped off his partner. He still holds a grudge with Jimmy, so he decides to make the game a little more interesting. At the end of the first round they decide to tell Bugsy whether or not the Toucan sang (that’s gangster talk for “ratted him out”), but they don’t tell Jimmy anything.

- f. Write down the extensive form form of this game. How many information sets and strategies does each player have now?
 - g. Write down a matrix representing the normal form of this game.
 - h. Find all pure strategy Nash equilibria and determine which of them are subgame perfect. Does Bugsy have an advantage now? Is Jimmy “the Toucan” any better or worse off?
3. **Repeated Negotiation with an Ice Cream Cake:** There are two players negotiating over how to divide an ice cream cake. Player 1 makes the first offer in the form of a fraction of the cake $\pi^2 \in [0, 1]$ which player 2 may eat. Player 2 can either accept the offer or reject it and propose a counter offer. However, if 2 decides to extend negotiations into the second round by making a counter offer for his own portion, then δ of the cake melts away. The game continues in like fashion for at most 3 rounds of negotiation. If no accord has been reached by the end of the 3rd round, then part of the cake melts away and player 2 gets to simply divide the cake any way he likes without question.
- a. How many strategies does each player have in this game?
 - b. Find the subgame perfect equilibrium of this game by backward induction.
 - c. How does the subgame perfect equilibrium change if 1 gets the final say?
 - d. How does the subgame perfect equilibrium change if play can extend to an arbitrary T rounds of negotiation? Answer this question for the cases where 1 and 2 have the final say.

4. **A Twice-Repeated Game with Multiple Static Equilibria:** Consider a dynamic game in which the following static game is repeated twice (final payoffs are the sum of payoffs from the two rounds):

	H	M	L
H	100, 100	1, 110	-10, -10
M	110, 1	20, 20	-100, -100
L	-10, -10	-100, -100	0, 0

- a. How many information sets and strategies does each player have in this game?
 - b. What is the dimensionality of a strategy for each player?
 - c. Find all subgame perfect equilibria of this game by backward induction. HINT: rather than writing out strategies as n -tuples, express them in the following form: In round 1 do “X;” in round 2 do “Y” if “W” was observed in round 1 and do “Z” otherwise. Can you explain why we can do this WLOG?
 - d. Find at least two Nash equilibria that are not subgame perfect and explain why they are not.
5. **Infinitely-Repeated Bertrand Game with Grim Trigger Strategies and Finite Punishment Paths:** There are 3 firms in an oligopolistic market in which production and allocation occurs in each period $t \in 0, 1, 2, \dots$. Market demand is given by

$$D(p) = 1000 - 2p,$$

where p denotes the prevailing market price. Per-unit production cost is \$1, firms have no capacity constraints (*i.e.* each firm can produce enough output to serve the entire market) and price competition follows a Bertrand scheme. The three firms simultaneously announce a price, the lowest of the three prices becomes the prevailing market price and the firm that announced that price captures all market demand. If there is a tie for the lowest price, then all firms involved in the tie divide market demand among themselves evenly.

- a. What are the Nash equilibrium strategies and profits of the static Bertrand game? is this equilibrium unique?
- b. Now suppose, hypothetically, that the three firms colluded and behaved as a single monopolist, dividing demand and profits evenly. What would be the prevailing market price, quantity and individual firm profits then?

- c. Suppose that this static Bertrand game were repeated for T periods. Firms maximize their discounted sum of profits, with discount factor $\beta \in (0, 1)$. Is there any Nash equilibrium in which the 3 firms collude at any point in time? HINT: try reasoning by backward induction.

Now suppose that the game were repeated for infinitely many periods. We wish to discover whether there are any collusive agreements which would induce all three competitors to collude in monopolistic pricing in every period. We say that such an agreement is self-sustaining. One possibility is called a *grim trigger strategy* (GTS). With GTS, firms agree to collude in monopolistic pricing until one firm is observed cheating, after which they revert back to Bertrand competition forever after. Another option is called a *finite punishment path strategy* (FPPS). With FPPS, the firms collude in monopolistic pricing until one of them is observed cheating. Starting with the next period, the firms revert back to Bertrand competition for K periods, after which they return to monopolistic collusion.

- d. What is the minimal discount factor needed for there to exist a subgame perfect Nash equilibrium with GTS? In other words, what is the minimal discount factor needed for collusion to be self-sustaining under GTS? HINT: For simplicity, you may assume that in-period price deviations can be arbitrarily small, or in other words, you may assume that a deviating firm is able to charge the monopoly price and still capture all market demand, leaving his confederates with nothing in that period.
- e. How does the answer to part d change if there is an arbitrary number of firms, N ? Another way to think about this question is, for any value of β , does there exist some large N such that collusion is not sustainable?
- f. Returning to the 3-firm case, what is the minimal discount factor needed for there to exist a subgame perfect Nash equilibrium with FPPS? That is, what is the minimal value of β for which collusion is self-sustaining under FPPS? HINT: for this problem, you not only need to make sure that firms have no incentive to deviate from collusion, but you must also verify that firms also have no incentive to deviate from the punishment path either. As before, you may assume that in-period price deviations are arbitrarily small.
- g. Is this discount factor different from the one in part d? If there is a difference, can you explain why it is so?

- h. How does the answer to part f change if there is an arbitrary number of firms, N ?