CS5204 Intelligent System Applications I

Paul Ahern

May 31, 2007

Abstract

Notes from lectures.

Contents

1 Key Facts to Know

1.1 Exam

- May contain question about sealed-bid auctions

2 Game Models

2.1 Cournot model of duopoly

In the Cournot model, firms simultaneously choose the quantities they will produce (of a homogeneous good), which they then sell at the market-clearing price. Equilibrium is determined by the condition that firms choose the action that is the best response to the anticipated play of their opponents.

Firm 1 and firm 2 simultaneously choose their respective output levels, q_i, from the feasible sets $Q_i = [0, \infty]$, say. They sell their output at the market-clearing price $p(q)$, where $q = q_1 + q_2$. Firm i's cost of production is $c_i(q_i)$,and firm i's total profit is then $u_i(q_1, q_2) = q_i p(q) - c_i(q_i)$.

The feasible sets Q_i and the payoff functions u_i determine the strategic form of the game. The "Cournot reaction functions" $r_1 : Q_2 \to Q_1$ and $r_2 : Q_1 \to Q_2$ specify each firm's optimal output for each fixed output level of its opponent.

For instance, for linear demand ($p(q) = \max(0, 1-q)$) and symmetric linear cost ($c_i(q_i) = cq_i$ where $0 \le c \le 1$), firm 2's reaction function is $r_2(q_1) = \frac{1-q_1-c}{2}$. Firm 1's reaction function is $r_1(q_2) = \frac{1-q_2-c}{2}$.

The Nash Equilibrium satisfies $q_2^* = r_2(q_1^*)$ and $q_1^* = r_1(q_2^*)$ or $q_1^* = q_2^* = \frac{1-c}{3}$.

3 Lecture 29 September 2006

Course will include:

- Game Theory
- Intelligent Systems - helping Humans make difficult decisions
- TAC SCM - team project

Main Points:

1. John Nash - Equilibrium Strategy, maximize payoff no matter what other agents do.

2. Expected Utility Theory - Poorer people tend to be more risk averse (Loss of stake has a bigger impact than a win).
3. Nash - Multiple Equilibria.
4. Extensiform Games - Games over time (Decisions made at different times).
5. Bertrand Competition - Predict competitors decisions based on your own.
6. Auction Games - Many different types of auction (e.g. Combinatorial Auctions).

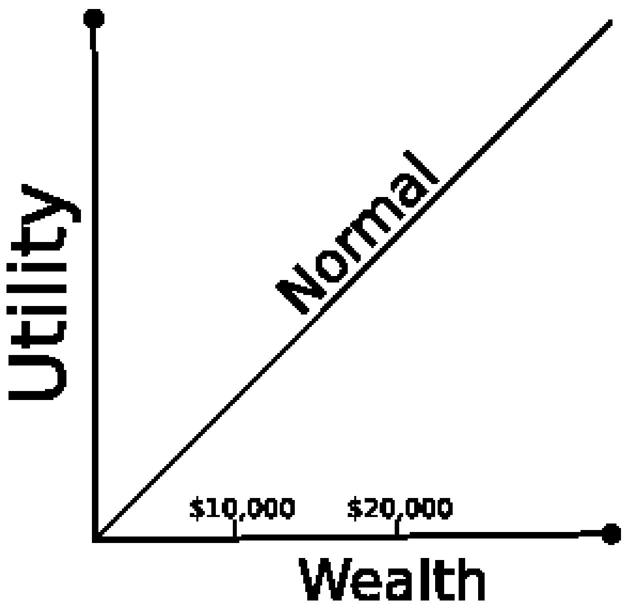

Figure 1: Risk Attitudes

Expected Utility Theory "The determination of the value of an item must not be based on the price, but rather on the utility it yields... There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount." - Nicolaus Bernoulli
A related concept is the *certainty equivalent* of a gamble. The more risk-averse a person is, the more he will be prepared to pay to eliminate risk, for example accepting $1 instead of a 50% chance of $3, even though the expected value of the latter is more. People may be risk-averse or risk-loving depending on the amounts involved and on whether the gamble relates to becoming better off or worse off; this is a possible explanation for why people may buy an insurance policy and a lottery ticket on the same day. However, expected utility as a descriptive model of decisions under risk has in recent years been replaced by more sophisticated variants that take irrational deviations from the expected utility model into account.

4 Lecture 4 October 2006

4.1 Game Theory Experiments

4.1.1 Pick a number between 0..100

Each person in the class picks a number between zero and one hundred.

Objective: Pick a number which is two thirds of the average responses.

If everyone is naive then the average picked is 50 and $\frac{2}{3}$ of that is 33.

However, if everyone has worked this much out and all guess 33 then the average is 33 and $\frac{2}{3}$ of that is 22!

This reasoning can be repeated and the average guess reduced to zero.

Ultimately in such a game played by practitioners of Game Theory everyone chooses zero.

4.1.2 Strategic Behavior

- Being aware of your fellow players existence and trying to anticipate their moves is called strategic behavior.
- Game theory is mainly concerned with models of strategic behavior. In the previous game, the winner has to correctly guess how often his fellow players iterate. Assuming infinite iterations would be consistent but those who bid 0 typically lose badly. Guessing higher numbers can mean two things:
    1. The player does not understand strategic behavior
    2. The player understands strategic behavior but has low confidence in the ability of other players to understand that this is a strategic game.

4.1.3 Auction $20

Each person can submit a sealed bid; Everyone pays their bid; Highest bidder gets the $20. If there are multiple instances of the highest bid then the $20 is assigned randomly amongst those bidders.

There is no equilibrium strategy. No optimum single bid for all players. For all cases where bidder i bids b_i at least one player will regret their choice (regret $\Rightarrow$ loss of money).

Although there is no pure strategy equilibrium, if all players are playing a randomized strategy then an equilibrium can be found (Nash Equilibrium).

If bidders can randomize there can exist an equilibrium so that no single player can profitably deviate from their strategy.

4.2 Game Theoretic Thinking

Game Theory is a formal way to analyze interaction among a group of rational agents.

Key concepts:

Group: More than one decision maker.

Interaction: What one player does directly affects at least one other player.

Strategic: Individual players take account of this interdependence.

Rational: While accounting for this interdependence, each player chooses his/her best action.

Aside: There are ways of modeling bounded rationality - Behavioral Game Theory.

Sample applications:

Trade: Levels of imports, exports, prices depend not only on your own taxes and tariffs but also on those of other countries.

Auctions: Your optimal bidding strategy depends on the actions of your opponents.

Labor: Internal labor market promotions like tournaments: your chances depend not only on effort but also on efforts of others.

Industrial Organization: Price depends not only on your output but also on the output of your competitors.

Elections: Your voting strategy should depend on the electoral rules and your beliefs about other voters preferences.

Example: Assume ten people go to a restaurant. If everyone pays for their own meal then this is just a decision problem.
If the bill is split equally then we have a game.
If the bill is split then the equilibrium bill is higher (good for the restaurant). People tend to order more expensive items as each reasons that they will be paying only 10% of the price difference!

4.3 Normal Form Game

Consists of:

1. A list of players $D = \{1, 2, \ldots i\}$ (we mainly consider $i = 2$).
2. Each player i can choose actions from a strategy set S_i.
3. The outcome of the game is defined by the *strategy profile* consisting of all strategies chosen by the players.
4. Players have preferences over outcomes (not actions).

Mathematically, preferences over outcomes are written as: $U_i : S \rightarrow R$ (where U_i is the utility or payoff for player i; R is a set of rational numbers).

E.g. In the New York game $U_i = 1$ if they choose the same location.

4.3.1 New York Game

Two people are to meet in New York either at Central Park (C) or at the Empire State building (E). The players want to meet but do not care where. They have to choose their individual destinations simultaneously (i.e. before knowing where the other has chosen to go).

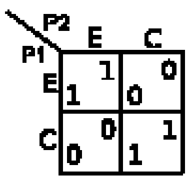

Figure 2: New York game

$S_1 = \{E, C\}$; $S_2 = \{E, C\}$

Strategy Profiles: $\{(E, E), (E, C), (C, E), (C, C)\}$ (here S has order 4)[1]

[1] If player 1 had 5 possible actions and player 2 had 10 then the set of profiles has order 50. $S = S_1 \times S_2$.

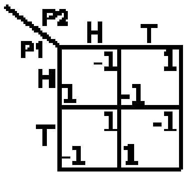

Figure 3: Matching Pennies

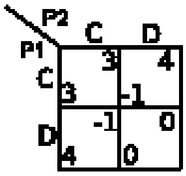

Figure 4: Prisoners' Dilemma

5 Lecture 6 October 2006

5.1 Normal Form (Strategic) Games

We say $S = S_1 \times S_2 \times \ldots \times S_I$.

Let $s \in S$, so s is one Strategy Profile. Set of chosen profiles within an instance of the game.

We denote the strategy choices of all other agents (except player i) as S_{-i} for $(S_1, S_2, \ldots S_{i-1}, S_{i+1}, \ldots S_I)$.

5.2 Matching Pennies

Zero Sum Games - these are games of conflict.

Matching Pennies game is a classic example. a two-player game where Player 1 gets one penny from Player 2 if both choose the same action. Player 2 gets one penny from Player 1 if they choose differently.

5.3 Prisoners' Dilemma

A cooperation game.

- Two prisoners interrogated separately.
- If both cooperate (with each other) - each will spend a year in prison.
- If both defect (give each other away) - each will spend three years in prison.
- If one cooperates and the other defects - Cooperator gets ten years; defector gets released.

The outcome (D,D) is Pareto dominated by (C,C). So (D,D) is not a good outcome for both (it is, however, a Pure Strategy Nash Equilibrium).

Some examples of Prisoners' Dilemma are:

Arms Race If two countries engage in an expensive arms race, this corresponds to (D,D). Both would prefer to spend their money elsewhere. (Likewise buying larger, safer, vehicles - neighbor has an incentive to get a larger car too).

Price War Two companies competing in a market can harm each other by cutting prices but fare better if they cooperate.

5.4 Cournot Competition

Two firms choose output levels (for a commodity product) q_i and have cost functions $c_i(q_i)$. Market price is $p(q_1 + q_2)$.

$D = \{1, 2\}$ (Demand is 1 or 2)

Utility to player equals quantity produced times the price minus the cost of production.

$u_1(q_1, q_2) = q_1 p(q_1 + q_2) - c_1(q_1)$

$u_2(q_1, q_2) = q_2 p(q_1 + q_2) - c_2(q_2)$

Game Theory can predict how much each company will produce.

6 Lecture 11 October 2006

6.1 Bertrand Competition

- In some respects the opposite of Cournot competition.
- Firms compete in a homogeneous product market but set prices.
- Consumers buy from the lowest cost firm.

NB: Under perfect competition firms are price-takers. In this case there is no strategic interaction. Each firm solves a decision problem to maximize profit.

Alternatively, imagine if both firms set equal prices above marginal cost, firms would get half the market at a higher than MC price. However, by lowering prices just slightly, a firm could gain the whole market, so both firms are tempted to lower prices as much as they can. It would be irrational to price below marginal cost, because the firm would make a loss. Therefore, both firms will lower prices until they reach the MC limit.

6.1.1 Normal Form

Two possible interpretations:

1. A game played once in time between a set of players.
2. One instance of a repeated game played between a large population of player 1's and player 2's who are randomly matched together to play this stage game. Random matching is important - if the normal form game is repeated with the same two players $\Rightarrow$ Repeated Extensive Form Game.
   In a repeated game you can arrive at a different equilibrium than in a one-off game.

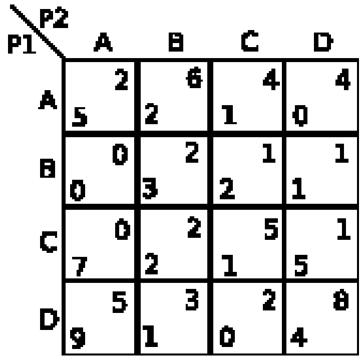

Figure 5: IDDS Normal Form Game

6.2 Iterated Deletion of Dominated Strategies

1. For P1, C weakly dominates ($\succeq$) A (each possible outcome is better or equal), $[7,2,1,5] \succeq [5,2,1,0] \Rightarrow$ Remove P1-A.
2. For P2, D strictly dominates ($\succ$) A (following earlier removal), $[1,1,8] \succ [0,0,5] \Rightarrow$ Remove P2-A.
3. For P1, C $\succ$D, $[2,1,5] \succ [1,0,4] \Rightarrow$ Remove P1-D.
4. For P2, B $\succ$D, $[2,2] \succ [1,1] \Rightarrow$ Remove P2-D.
5. For P1, B $\succ$C, $[3,2] \succ [2,1] \Rightarrow$ Remove P1-C.
6. For P2, B $\succ$C, $[2] \succ [1] \Rightarrow$ Remove P2-C.

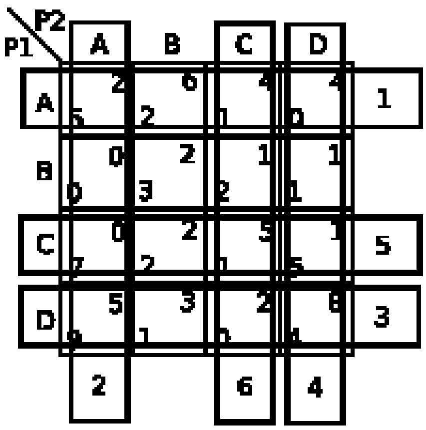

Figure 6: IDDS Solved Game

The game has now been solved. Nash Equilibrium outcome is {B,B}.

Warning Not all games can be solved using iterated deletion of dominated strategies:
Games can have multiple Nash Equilibria.
Not all equilibria are Pure Strategy Nash Equilibria (Mixed Strategy Nash Equilibria require randomization - e.g. rock-paper-scissors game)

7 Lecture 13 October 2006

7.1 Rationality

Player i is rational with beliefs μ_i if

$$s_i \in argmax_{s_i'} E_{\mu_i(s_{-i})} u_i(s_i, s_{-i})$$

where each agent has a belief (a probability distribution over the strategy set S_{-i}) μ_i about the play of her opponents. A player faces a simple decision problem as soon as this belief has been formed.

It is rational to choose a strategy that maximizes expected utility given μ_i. The choice of your strategy depends on your belief of what your opponent will do (i.e. what strategy he will follow).

7.1.1 New York Game

Assume in the New York game that I believe my friend will choose the Empire State with a 60% probability and Central Park with 40%.

If I go to Central Park I induce the following lottery:

- With 60% probability I'll see (C,E), and
- With 40% probability I'll see (C,C).

If I go to Empire State:

- With 60% probability I'll see (E,E), and
- With 40% probability I'll see (E,C).

The rational choice is to go to Empire State. If I am rational then I shall choose Empire State (with probability 1 - i.e. all the time) .

7.1.2 Prisoners Dilemma

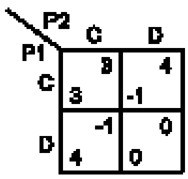

Figure 7: Prisoners Dilemma

Assume that I believe in the Prisoners Dilemma game that my companion will cooperate with 80% probability and defect with 20% probability.

If I cooperate I induce the following lottery L^C over outcomes of the game: with 80% probability I will see the outcome (C, D) and with 20% (C, C): $L^C = 0.8(C,C) + 0.2(C,D) = 2.4 - 0.2 = 2.2$.

Similarly, if I were to choose defect: $L^D = 0.8(D,C) + 0.2(D,D) = 3.2$.

So, given this belief, I would defect.

Definition Strategy s_i is strictly dominated for player i if there is some $s_i' \in S_i$ such that $u_i(s_i', s_{-i}) > u_i(s_i, s_{-i}) \forall s_{-i} \in S_{-i}$.[2] Rational players don't play strictly dominated strategies.

The hardest task is to determine players' beliefs. A lot of games can be simplified by rationality and the knowledge that my opponent is rational. To see that look at the Prisoners Dilemma. Cooperating is a dominated strategy so a rational player would therefore never cooperate. This solves the game since every player will defect. This is the worst outcome in terms of joint surplus and it would be Pareto improving if both players would cooperate.

7.1.3 Removal of Strictly Dominated Strategies

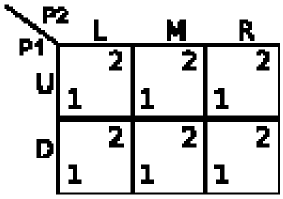

Figure 8: LMR Game

1. Eliminate M (L≻M).
2. Eliminate D (U≻D).
3. Eliminate R (L≻R).

Pure Strategy Nash Equilibrium (PSNE) is (L,U).

Many (most) games are not dominance solvable (e.g. Coordination games, zero-sum games etc.) But it can make a game smaller and more manageable.

Note that we have not specified the order in which strategies are eliminated. For strictly dominated strategies this does not matter - the order of elimination of weakly dominated strategies can make a difference.

8 Lecture 17 October 2006

8.1 Removal of Weakly Dominated Strategies

Order of removal of weakly dominated strategies can lead to different outcomes.

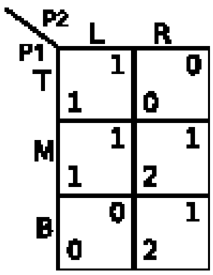

Figure 9: Weakly Dominated Strategies

[2]Weak dominance is $u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i}) \forall s_{-i} \in S_{-i}$.

- We can eliminate T and then L in which case we know that (2, 1) is played for sure.
- However, if we eliminate B first and then R we know that (1, 1) is being played for sure.
- Weak elimination does not deliver consistent results.

8.2 Cournot Competition

Cournot competition with two firms can be solved by iterated deletion in some cases.

We examine a linear demand function $p = \alpha - \beta(q_i + q_j)$ and constant marginal cost c such that the total cost of producing q_i units is cq_i.

We need to determine a "best-response" function $BR(q_i)$ of each firm i to the quantity choice q_j of the other firm.

By taking the first-order condition of the profit function you can derive the best-response function for both firms (they are symmetric) is:

$$BR_i(q_j) = \begin{cases} \frac{\alpha-c}{2\beta} - \frac{q_j}{2} & , if\ q_j \leq \frac{\alpha-c}{\beta} \\ 0 & otherwise \end{cases}$$

If the opponent (player j) increases the quantity he produces (q_j) above $\frac{\alpha-c}{\beta}$ the price drops to $c - \beta q_i$. The cost of producing an item is c. If player chooses a $q_i > 0$ the price drops below c. You only want to produce goods when the price they command is greater than the cost of production.

The best-response function is decreasing in my belief of the other firm's action.

Note, that for $q_j \leq \frac{\alpha-c}{\beta}$ firm i makes a loss even if it chooses the profit maximizing output. It therefore is better to stay out of the market and choose $q_i = 0$. Initially, firms can set any quantity, i.e. $S_1^0 = S_2^0 = \mathcal{R}$.

However, the best-responses of each firm to any belief has to lie in the interval $[\underline{q}, \bar{q}]$ with $\underline{q} = 0$ and $\bar{q} = \frac{\alpha-c}{\beta}$. All other strategies make negative profit (and are thus eliminated).

In the second stage only the strategies $S_1^2 = S_2^2 = [BR_1(\bar{q}), BR_1(\underline{q})]$ survive.

9 Lecture 19 October 2006

Nash Equilibria: Expected outcomes in games (even ones without dominated strategies).

Iterated Dominance is an attractive concept. It only assumes that players are rational and that this is public knowledge. It is useful in reducing (simplifying) many games.

However, rarely can you solve a game via iterated dominance.

9.1 Pure Strategy Nash Equilibrium

Nash Equilibrium (NE) is the most important concept in game theory.

Definition: A strategy profile S^* is a pure strategy NE of game G iff $u_i(S_i^*, S_{-i}^*) \geq u_i(S_i, S_{-i}^*)$ for all players $i \forall s_i \in S_i$.

If there is a set of strategies with the property that no player can benefit by changing her strategy while the other players keep their strategies unchanged, then that set of strategies and the corresponding payoffs constitute the Nash Equilibrium.

E.g. (Defect,Defect) in the Prisoners Dilemma.

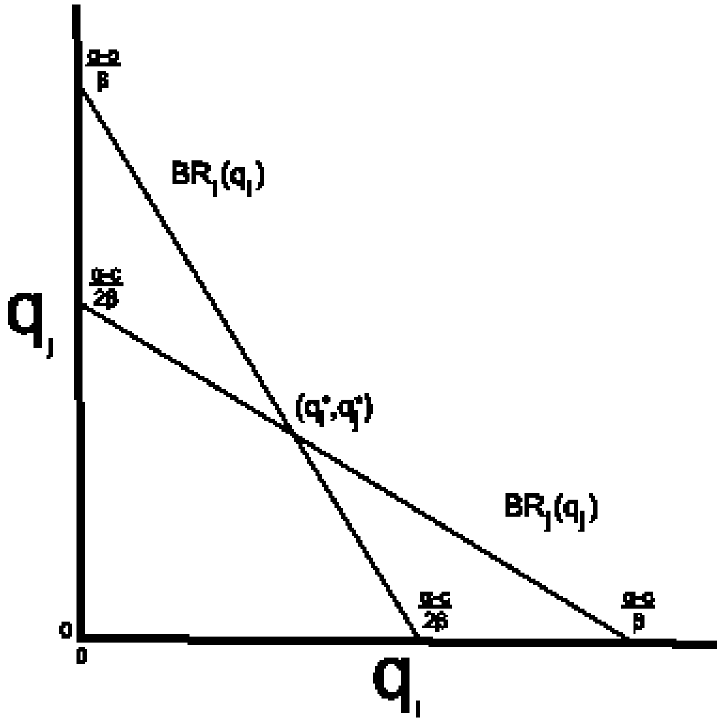

Figure 10: Cournot Competition

P1 \ P2	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Figure 11: Rock-Paper-Scissors

9.2 Mixed Strategy Nash Equilibrium

Determine probabilities of choosing strategies in Mixed Strategy.

Definition: Let G be a game with strategy spaces $S_1, S_2, ..., S_I$. A mixed strategy σ_i for player i is a probability distribution on S_i i.e. for S_i a mixed strategy is a function: $\sigma_i : S_i \to \mathbb{R}^{+3}$ such that $\Sigma_{s_i \in S_i} \sigma_i(s_i) = 1$.

For example in the Matching Pennies Game: $\delta_1(H) = \frac{1}{2}$, $\delta_1(T) = \frac{1}{2}$, $\delta_2(H) = \frac{1}{2}$, $\delta_2(T) = \frac{1}{2}$.

Mixed Strategy Nash Equilibrium (MSNE): We write Σ_i for the set of probability distributions on S_i. We let $\Sigma = \Sigma_1 \times \Sigma_2 \times ... \times \Sigma_I$.

A mixed strategy profile $\sigma \in \Sigma$ is an is an i-tuple $(\sigma_1, \sigma_2, ..., \sigma_I)\sigma_i \in \Sigma_i$.

We write $u_i(\sigma_i, \sigma_{-i})$for player *i's* expected payoff when he uses mixed strategy σ_i and all other players σ_{-i}.

[3] Set of Positive Rational numbers.

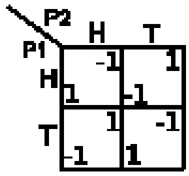

Figure 12: Matching Pennies

10 Lecture 1 November 2006

10.1 Mixed Strategy Nash Equilibria

Example Coordination Game.

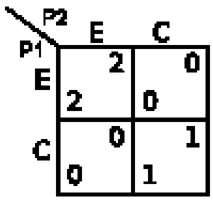

Figure 13: New York Game

In the New York Game: Let α represent the probability that Player 1 plays E(mpire State). Therefore $1-\alpha$ is the probability that Player 1 plays C(entral Park). Likewise β and $1-\beta$ are the probabilities that Player 2 plays E or C.

Expected payoffs for player 1:

$$Exp_1(E) = 2\beta + 0(1-\beta) = 2\beta$$

$$Exp_1(C) = 0\beta + 1(1-\beta) = 1-\beta$$

Expected payoffs for player 2:

$$Exp_2(E) = 2\alpha + 0(1-\alpha) = 2\beta$$

$$Exp_2(C) = 0\alpha + 1(1-\alpha) = 1-\alpha$$

For a player to be indifferent between strategies their utilities must be equal: $u_1(E, \sigma_2^*) = u_1(C, \sigma_2^*)$.

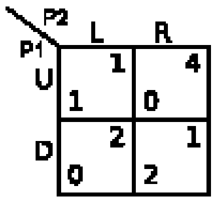

Figure 14: Asymmetric Game

In this Asymmetric Game: P1 is indifferent between U and D. P2 is indifferent between L and R. Assume P1 plays U with probability α. Assume P2 plays L with probability β.

Because $u_1(U, \sigma_2^*) = u_1(D, \sigma_2^*) : \beta = 2(1-\beta) \Rightarrow \beta = \frac{2}{3}$,

$u_2(L, \sigma_1^*) = u_2(R, \sigma_1^*) : \alpha + 2(1-\alpha) = 4\alpha + (1-\alpha)) \Rightarrow \alpha = \frac{1}{4}$.

Given your opponent is playing an equilibrium strategy.

Interpretations of MSNE:

1. Sometimes players flip coins (e.g. poker, soccer, tennis).

2. Large populations of players, each playing a fixed strategy and randomly matched.

11 Lecture 3 November 2006

11.1 Nash Equilibrium

There is a profile of strategies for which play is expected to be stable. In a NE no single player can increase their payoffs by deviating from their strategy.

Theorem Every finite strategic-form game has a mixed-strategy equilibrium. (Nash 1950)

Much research followed into refining the notion of NE.

There can be many NE in a game.

Behavioural Game Theory tried to weaken the joint assumptions of rationality and common knowledge.

Nash existence can be proven using:

1. Elementary geometric techniques in a 2x2 game;

2. A fixed point approach (Kakutani theorem).

11.1.1 Elementary Geometric Proof

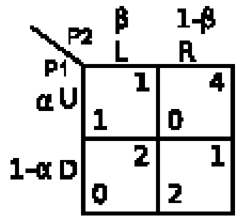

Figure 15: Sample Game (with probabilities)

Draw Best Response Curves: $\sigma_1 = (\alpha)U + (1-\alpha)D$ and $\sigma_2 = (\beta)L + (1-\beta)R$.

The Best Response of P2 to P1 playing α:

$$U_2(L, \alpha U + (1-\alpha)D) = 2 - \alpha$$

$$U_2(R, \alpha U + (1-\alpha)D) = 3\alpha + 1$$

$\therefore$ P2 strictly prefers to play L over R when $2 - \alpha > 3\alpha + 1$ or $\frac{1}{4} > \alpha$.

The Best Response for P2:

$$BR_2(\alpha) = \begin{cases} 1 & if\, \alpha < \frac{1}{4} \\ [0,1] & if\, \alpha = \frac{1}{4} \\ 0 & if\, \alpha > \frac{1}{4} \end{cases}$$

The Best Response for P1:

$$BR_1(\beta) = \begin{cases} 1 & if\, \beta < \frac{2}{3} \\ [0,1] & if\, \beta = \frac{2}{3} \\ 0 & if\, \beta > \frac{2}{3} \end{cases}$$

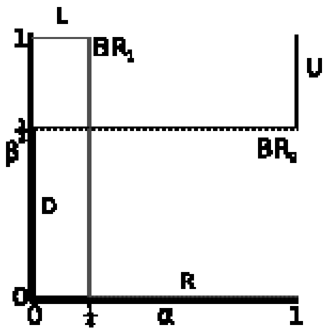

Figure 16: Best Response Graphs

Both correspondences intersect in the single point $(\alpha = \frac{1}{4}, \beta = \frac{2}{3})$. Therefore this is the unique MSNE.

This approach is useful because it generalises to a proof that any 2x2 game has at least one NE.

11.2 Extensive Form Games

Models presented thus-far assume simultaneous moves. This misses some common features of some games. For example: Auctions (sealed bid/oral); Patent Race; Price Competition; Central Bank (Monetary Policy); Deciding on Manufacturing Output; Poker; . . .

The Extensive Form Game is a complete description of:

1. Set of Players.
2. Who moves when and what their choices are.
3. The players' payoffs as a function of the choices that are made.
4. What the players know when they move.

11.2.1 Model of Entry

Firm 1 is an incumbent monopolist. A second Firm, 2, has the opportunity to enter that market. After Firm 2 decides to enter, Firm 1 can choose to fight back with aggressive pricing, or accommodate with higher prices.

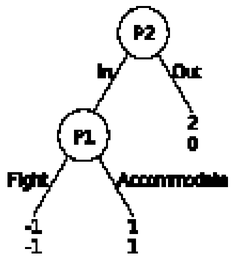

Figure 17: Market Entry Game

12 Lecture 8 November 2006

12.1 Stackelberg Competition

Firm 1 develops a new technology before firm 2. Firm 1 has the opportunity to build a factory and commit to an output level q_1 before firm 2 starts.

Firm 2 observes firm 1's decision, before choosing output level q_2.

Suppose $q_1 \in \{0, 1, 2\}$ and market demand is $p(Q) = 3 - Q$ ($Q = q_1 + q_2$). The marginal cost of production is zero.

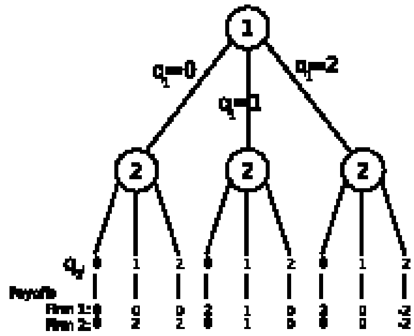

Figure 18: Stackelberg Game

12.2 Sequential Matching Pennies

So far we have assumed that players observe previous moves.

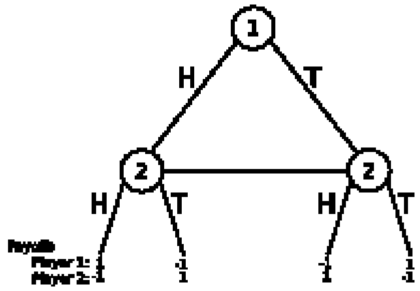

Figure 19: Sequential Matching Pennies

Player 2 doesn't know which node he is at in the information set (sometimes called information partition).

12.3 Normal Form Analysis

In an Extensive Form Game, write H_i for the information sets at which player i moves. $H_i = \{SCT \mid S = h(t) \text{ for some } t \in T \text{ with } i(t) = i\}$ where an information partition for each node t, $h(t)$ is the set of nodes which are possible given what player $i(x)$ knows.

The partition must satisfy the following

$$t' \in h(x) \Rightarrow i(t') = i(t), A(t') = A(t), h(t') = h(t)$$

$A(i)$ is the set of actions available.

In Sequential Matching Pennies:

1. There are two players.
2. $S_1 = S_2 = \{H, T\}$.
3. The tree defines the terminal nodes and show that t_2 and t_3 are successors to t_1.

4. $h(t_1) = \{t_1\}$ and $h(t_2) = h(t_3) = \{t_2, t_3\}$.

Definition: A pure strategy for player 1 in an extensive form game is a function: $S_i \to A_i$ such that $S_i(h) \in A(h) \forall h \in H$ (H is the set of information sets at which A moves).
A strategy is a *complete contingent plan* explaining what a player will do in all possible situations.

Might earlier actions make it impossible to reach certain sections of a tree? Why do we specify how players would play at nodes that cannot be reached?

The reason is that play off the equilibrium path is essential to determine if a set of strategies is a NE. Off equilibrium threats are crucial.

Definition: A mixed behaviour strategy for player i is a function $\sigma_i : H \to \Delta(A_i)$ such that *supp*$(\sigma_i(h)) \subset A(h) \, \forall \, h \in A$ (*supp* is the support set).
Note: an independent randomisation at each information set.

13 Lecture 10 November 2006

Extensive form games can be represented in normal form.

Entry Game: We can find the pure strategy sets: S_1 ={Fight, Accommodate}; S_2 ={Out, In}.

P1 \ P2	O	I
F	2, 0	-1, -1
A	2, 0	1, 1

Figure 20: Normal Form of Entry Game

Stackelberg Competition: Firm 1 chooses q_1 and firm 2 q_2.

With three possible output levels, firm 1 has three strategies. Firm two has $3 \times 3 = 9$ strategies. So the normal form of the game has $3^3 = 27$ strategies.

P1 \ P2	$q_1=0$, $q_2=0$	$q_1=1$, $q_2=0$	$q_1=2$, $q_2=0$	$q_1=0$, $q_2=1$	$q_1=1$, $q_2=1$	$q_1=2$, $q_2=1$	$q_1=0$, $q_2=2$	$q_1=1$, $q_2=2$	$q_1=2$, $q_2=2$
$q_1=0$	0, 0	0, 0	0, 0	0, 2	0, 2	0, 2	0, 2	0, 2	0, 2
$q_1=1$	2, 0	2, 0	2, 0	1, 1	1, 1	1, 1	0, 0	0, 0	0, 0
$q_1=2$	2, 0	2, 0	2, 0	0, 0	0, 0	0, 0	-2, -2	-2, -2	-2, -2

Figure 21: Normal Form of Stackelberg

Sequential Matching Pennies: $S_1 = \{H, T\}$.

Player 2 has four strategies as they choose two actions at each information set. Strategy HH implies that Player 2 chooses H at both nodes while HT implies that Player 2 chooses H in the left node (after observing H) and T in the right node (after observing T).

P1 \ P2	HH	HT	TH	TT
H	1, -1	1, -1	-1, 1	-1, 1
T	-1, 1	1, -1	-1, 1	1, -1

Figure 22: Normal Form of Matching Pennies

Recall that Player 2 wins if different sides announced.

These represent complete contingent plans.

We can apply NE in extensive form games by looking at the normal form representation.

However, this is not an appealing *solution concept* because it allows too many profiles to be Equilibria.

14 Lecture 22 November 2006

Presentations.

14.1 Deep Maize

Controlling a Supply Chain Management agent using value-based decomposition.

Predict demand and supply curves and calculate marginal value (cost of components and price of completed) of PCs.

Instead of optimizing the overall profit margin, the sales decision optimizes the margin between expected revenue and the value of the PCs sold. Similarly the purchasing decision optimizes the margin between the value of he components purchased and the total cost.

Overly elaborate strategy requiring far too much computing resource given the marginal benefits received.

14.2 Tac-Tex-05

Utilises a modular approach, takes advantage of the artificial nature of having a fixed length game with a definite end, and can adapt its strategies both in an individual game and throughout a series of games when it encounters the same agents in multiple games.

The Supply Manager handles all planning related to component inventories and purchases, and requires no information about computer production except for a projection of component use over a future period, which is provided by the Demand Manager.

The Demand Manager, in turn, handles all planning related to computer sales and production. The only information about components required by the Demand

Manager is a projection of the current inventory and future component deliveries, along with an estimated replacement cost for each component used. This information is provided by the Supply Manager.

14.3 SouthamptonSCM

The design and evaluation of SouthamptonSCM, the runner-up in the 2005 International Trading Agent Supply Chain Management Competition (TAC SCM). In particular, the way in which the agent purchases components using a mixed procurement strategy (combining long and short term planning) and how it sets its prices according to the prevailing market situation and its own inventory level (because this adaptivity and flexibility are key to its success).

SouthamptonSCM is composed of three sub-agents. The customer agent receives RFQs from the customers and decides what offers to respond with. It also communicates with the factory agent to obtain the updated inventory levels and to send the relevant customer PC orders. The component agent decides which RFQs and which orders to send to which suppliers. The factory agent receives the supplies delivered from the suppliers, decides based on the available resources (computer components and factory cycles) in what order the customer orders should be produced, and determines the schedules for delivering the finished PCs to the customers.

14.4 Simple search methods for finding Nash Equilibrium

Presented two simple search methods for computing a sample Nash equilibrium in a normal-form game: one for 2-player games and one for n-player games. These algorithms were tested on many classes of games, showing that they perform well against the state of the art: the Lemke-Howson algorithm for 2-player games, and Simplicial Subdivision and Govindan-Wilson for n-player games.

The basic idea behind the new search algorithms is simple. Recall that, while the general problem of computing a NE is a complementarity problem, computing whether there exists a NE with a particular support for each player is a relatively easy feasibility program. Our algorithms explore the space of support profiles using a backtracking procedure to instantiate the support for each player separately. After each instantiation, they prune the search space by checking for actions in a support that are strictly dominated, given that the other agents will only play actions in their own supports.

15 Lecture 29 November 2006

15.1 NE in Extensive Form games

Can be applied by looking at the normal form representation. There can be an infinite number of NE. Some NE rely on empty threats.

Stackelberg example: For any $q_i' \in [0,1]$ the game has a NE where Player 1 produces q_i (relies on Player 2 flooding the market - an empty threat).

15.2 Sub-game Perfect Equilibria

These are a subset of NE. Many NE are unreasonable - because they are based on empty threats (e.g. flooding the market in the entry game). In equilibrium this threat is *not* carried out.

Sub-game perfection is a refinement of NE and rules out non-credible threats.

Definition: A sub-game G' of an extensive form game G consists of the following:

1. A subset of all the nodes of G: T' consisting of a single node plus all of its successors which has the property $t \in T'; t' \in h(t)$ then $t' \in T'$. ($h(t)$ is the set of nodes that are possible given what player $i(x)$ knows)
2. Information Sets, feasible moves and payoffs at terminal nodes are the same as in G.

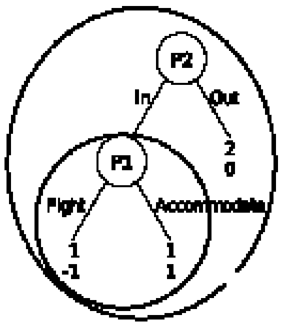

Figure 23: Example 1 Entry Game

Whole game is always the first sub-game.
Start a new sub-game when player knowledge is updated.

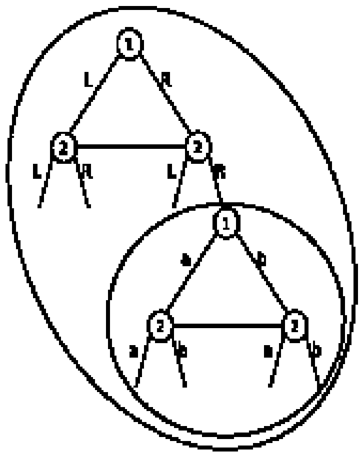

Figure 24: Example 2

This sub-game has no further sub-games, otherwise the information set of Player 2 would be separated.

Definition: A strategy profile S^* is a Sub-game Perfect Equilibrium (SPE) of G if it is a NE of every sub-game.

Note that a SPE is a NE.

Example Stackelberg.

We claim that the *unique* SPE is $q_2^* = \frac{1}{2}$ and $q_1^*(q_2) = \frac{1-q_2}{2} = \frac{1}{4}$.

Proof: A SPE must be a NE in the sub-game after firm 1 has chosen q_1. This is a 1-player game so NE is equivalent to firm 1 maximizing its payoff (i.e. $q_1^*(q_1) \in argmax_{q_1}[1-(q_1+q_2)]$).

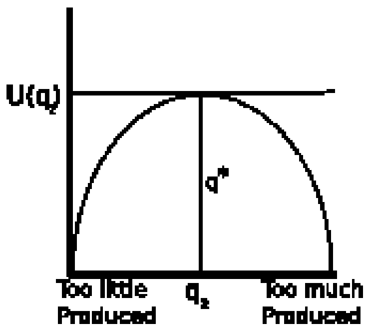

Figure 25: Payoff for Player 2

P1 \ P2	L	C	R
T	3, 1	0, 0	5, 0
M	2, 1	1, 2	3, 1
B	1, 2	1, 1	4, 4

Figure 26: Example Normal Form Game

This implies that $q_1^*(q_2) = \frac{1-q_2}{2}$. Equivalently, Player 2 plays on its best response curve.

A SPE must also be a NE in the whole game, so q_2 is a Best Response to q_1^*.

$U_2(q_1, q_2^*) = q_2(1 - (q_2 + q_1^*(q_2))) = q_1\frac{(1-q_1)}{2}$

16 Lecture 1 December 2006

16.1 Repeated Games

Dynamic games can result in large numbers of Nash Equilibria (NE).

Sub-game Perfect Equilibrium (SPE) dealt with this by ruling out non-credible threats.

However, dynamic games can also have large numbers of SPEs.

Many finite-horizon games contain Credible Threats which can cause multiple SPEs.

Example: Normal Form Game

The game has three NE: (T,L), (M,C) & ($\frac{1}{2}$T+$\frac{1}{2}$L,$\frac{1}{2}$M+$\frac{1}{2}$C).

Suppose the players were to play the game twice and observe the first period actions before the second period actions.

One way to get a SPE is to play any of the three profiles above followed by any other[4].

We can also, however, use credible threats to get other actions played in the first period. Such as:

```
Play (B,R) in Period 1
```

[4] May be the same one.

```
IF Player1 plays B in Period 1
   Play (T,L) in Period 2
Otherwise
   Play (M,C) in Period 2
```

No single-period deviation helps here.

16.2 Repeated Prisoner's Dilemma

The Prisoner's Dilemma game has a single NE. So there are no easy threats to use. Therefore the unique SPE is (D,D).

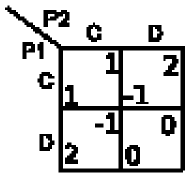

Figure 27: Prisoners' Dilemma

In the infinite-horizon game we get many SPEs because other types of threats are credible. Examples:

- Tit-for-tat strategy ("You hurt me and I'll hurt you").
- Both Players Cooperate in every period.
- Defect in first period and Cooperate in all future periods.
- Defect in every even period and Cooperate in every odd period.
- (C,D) in even periods and (D,C) in odd periods.

16.3 Recipe for checking for SPE

- First try to classify all histories (information sets) on and off the equilibrium path.
- Apply the Single Period Deviation Principle (relates to gains in a single period).

For Example, assume that you want to check that Cooperate in even periods and Defect in odd periods and *Grim Trigger Punishment* in case of deviation.

There are three types of histories... (Outside the scope of this Module).

16.4 Discounting Factor

Previous examples suggest that the Prisoner's Dilemma can have a large number of Equilibria, but this is only true if a discounting factor δ is employed (and is large). This is typically true of repeated games.

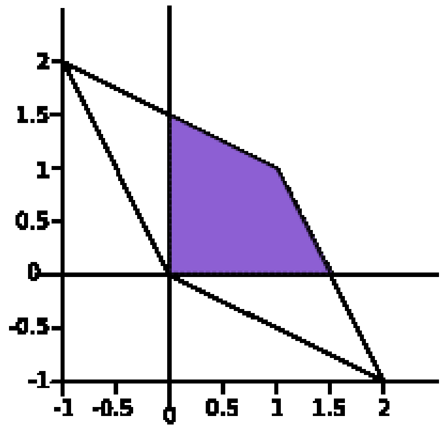

Figure 28: Repeated Prisoners' Dilemma Payoffs

17 Lecture 6 December 2006

17.1 Infinite Games with Discounting

δ discounting factor.

$U_i(S_i, S_{-i}) = U_{i0}(S_i, S_{-i}) + \delta U_{i1}(S_i, S_{-i}) + \delta^2 U_{i2}(S_i, S_{-i}) + \dots$

Where $U_{it}(S_i, S_{-i})$ is a payoff received at time t when the strategies are followed.

Interpretations of δ:

Interest Rate $\delta = \frac{1}{1+r}$ (e.g. ECB rate 4% $\rightarrow \delta = \frac{1}{1+0.04} = 0.96$). By discounting future payoffs by δ we correct for the fact that future payoffs are worth less.

Probabilistic End-Of-Game The game is really finite, but the end of the game is not deterministic. Instead, given that stage t is reached, there is a probability δ that the game continues (or a probability of $1 - \delta$ that the game ends).

17.2 Folk Theorem

Suppose that the set of feasible payoffs of G is I-dimensional (I is the number of players). Then for any feasible and strictly individually rational payoff vector v, there exists $\underline{\delta} < 1$ such that $\forall \delta < \forall \underline{\delta} \exists$ a SPE x^* of G^∞ such that the average payoff to s^* is v_i.

$u_i(s^*) = \frac{v_i}{1-\delta}$

The average payoff $p = (1 - \delta) u_i(s^*)$.

Example: Repeated Prisoners' Dilemma

Possible payoffs: (0, 0), (2, -1), (-1, 2), (1, 1).

The Average Payoff must be in the trapezoid.

The minimax payoff for each player is 0 (The NE is (D, D) with payoffs (0, 0)). The other player can at most punish his rival by defecting so that player gets 0.

17.3 Applications of Repeated Games

A market with several competitors we call an oligopoly. Classical economics assumes perfect competition.

17.3.1 The Bertrand Paradox

In a Bertrand duopoly - firms set prices. The price arrived at is closer to the higher monopoly price than that suggested by classical economics. Many economists believe that this is counter-intuitive (recouping fixed costs not included).

Repeated games provide an explanation:

- Firms can "cooperate".
- Set prices above marginal cost in each period.
- If a firm defects then they both revert to static Nash pricing at the marginal cost.
- This is an example of tacit collusion.

17.4 Stiglitz's Efficient Wages Model

Why do firms pay workers more than they have to pay in order to prevent the workers from leaving?

17.4.1 One Stage shirking Model

Firm and worker play a 2-period game.

1. Firm sets wage $= w$.
2. Worker observes wage and decides to accept or reject the job.
    - In case of rejection, gets w_0 (social welfare?)
    - In case of acceptance, decides whether to make an effort or shirk.

18 Lecture 8 December 2006

If the worker makes no effort then she will produce output $y > 0$ with probability p and 0 otherwise.

If the worker makes an effort then she will produce y for sure. (So $output = 0 \Rightarrow$ *no effort*)

Exerting effort has a cost e to the worker.

Note that *the firm cannot enforce effort* - there will be shirking. The firm only has to pay w_0 in order to employ the worker - paying higher wages makes no sense because the worker shirks anyway.

18.1 Repeated Interactions in Stiglitz's Model

How can the firm prevent the worker from shirking? The following grim-trigger type strategy dissuades shirking.

1. The firm pays the worker a higher wage $w_* > w_0$ (otherwise the worker has nothing to lose).
2. The firm has to fire the worker if it detects shirking. Since $w_* > w_0$, the worker has an incentive not to shirk.

A worker who exerts effort over their career gets a surplus: $V_e = \frac{w_* - e}{1-\delta}$ (formula derived from $w_* + \delta w_* + \delta^2 w_* + \ldots$; δ is the discounting factor).

If the worker were to shirk during one period, she gets the following payoff: $V_s = w_* + \delta[pV_e + (1-p)\frac{w_0}{1-\delta}]$ (where p is the probability that the shirking not being detected).

So V_e is the expected payoff in your lifetime if you never shirk; While V_s is the expected payoff in your lifetime if you shirk in period 1 and work every subsequent period.

In equilibrium, each worker should have a greater expected payoff from working than shirking $\therefore V_e \geq V_s$.

So $\frac{w_* - e}{1-\delta} \geq w_* + \delta[pV_e + (1-p)\frac{w_0}{1-\delta}]$.

And $w_* \geq w_0 + \frac{(1-\delta p)e}{\delta(1-p)}$.

Firm's best strategy is to set wage $= w_0 + \frac{(1-\delta p)e}{\delta(1-p)}$ (i.e. to the minimum necessary).

19 Lecture 13 December 2006

19.1 Games of Incomplete Information

In the real world, players usually don't have complete information on their opponents.

Sometimes the game being played itself is not common knowledge.

Unknowns:

1. Payoffs: In a price or quantity competition model, you may know that your rival attempts to maximize profit, but what are his costs? (Unknown Costs $\Rightarrow$ Profits Unknown).

2. Identity of other players (e.g. Buying a house; R&D race).

3. What moves are possible (e.g. Levels of quality that are possible - a rival might offer a lower quality good or service at a lower price).
   Lack of information about possible strategies can inhibit your payoffs.

4. How does the outcome depend on action (e.g. Stiglitz Model - the workers may not know the probability of being caught; TV license enforcement adverts on TV designed to increase people's perception of the probability of being caught if without a license).

19.2 Examples

19.2.1 Crazy Incumbent

Entry Game: Incumbent may choose to fight (and get higher utility out of this choice) because of their personality.

Sub-game Perfect Equilibrium (SPE) differs if game includes crazy incumbent.

19.2.2 Auction

Two bidders are trying to purchase the same item.

- Sealed bid auction.

- They simultaneously choose bids b_1 and b_2.

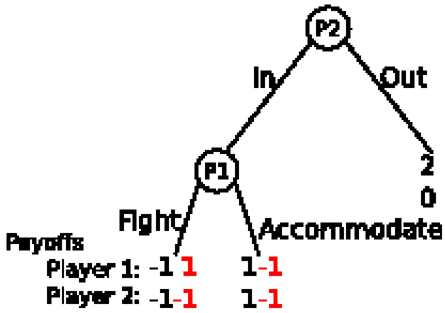

Figure 29: Entry Game with Crazy Incumbent (Red)

- Good is awarded to the highest bidder at that bid price (assume coin-flip if $b_1 = b_2$).
- Suppose that the player utilities are:

$$U_i(b_i, b_{-i}) = \begin{cases} U_i - b_i & if\, b_i > b_{-i} \\ \frac{1}{2}(U_i - b_i) & if\, b_i = b_{-i} \\ 0 & if\, b_i < b_{-i} \end{cases}$$

Where $U_i$is the value of the item to player i.

The missing information is the other players' valuations and their perceptions (beliefs) about other agents' valuations.

19.2.3 Public Good

Two advisers of a graduate student each want the student to get a job at company X. Each can ensure this by calling someone and exaggerating the abilities of the student.

Suppose the payoffs are as shown

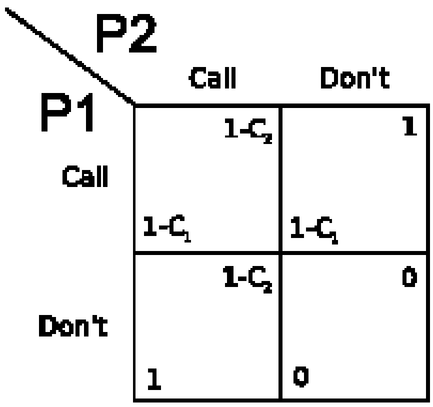

Figure 30: Public Good

- Assume the actions are chosen simultaneously and players know only their own costs.

- They believe that $C_{-i} \in [\underline{C}, \overline{C}]$.
- Alternatively, we could have player 1's cost known to all (say $C_1 = \frac{1}{2}$) but $C_2 \in [\underline{C}, \overline{C}]$.
- Or player 1 is a senior faculty member and knows from experience that the cost of such calls ($C_1 = \frac{1}{2}, C_2 = \frac{2}{3}$). Player 2 is inexperienced and has prior beliefs: $C_1, C_2 \in [0, 2]$ (uniformly distributed).

20 Lecture 15 December 2006

20.1 Bayesian Games

Definition: a game with incomplete information, $G = (\Phi, S, P, u)$ consists of:

1. A set $\Phi = \Phi_1 \times \Phi_2 \times \ldots \times \Phi_i$ is the set of possible types (values).
2. A set $S = S_1 \times S_2 \times \ldots \times S_i$ giving possible strategies for each player.
3. Joint probability distribution $p(\phi_1, \phi_2, \ldots, \phi_i)$ over the types for finite type space.
   $p(\phi_1) > 0$ probability for each type is greater than zero.
4. A payoff function $u_i : S \times \Phi \to R$[5]

NB: Payoffs not only depend on your type but also on your rivals. Players know their own type but not the other players' types.

20.2 Public Good

Suppose Player 1 (P1) is known to have cost $c_1 < \frac{1}{2}$. Player 2 (P2) has cost $\underline{c}$ (e.g. $\frac{1}{4}$) with probability p and $\overline{c}$ (e.g. $\frac{5}{6}$) with probability $(1-p)$[6].

Assume $0 < \underline{c} < 1 < \overline{c}$ and $p < \frac{1}{2}$.

P2 knows his cost $>$ P1's and vice-versa.

The unique Bayesian Nash Equilibrium (BNE), or best response functions, is:

- f_1^* =Call
- f_2^* =Don't Call

Proof: In a BNE each type of player must play a best response. So for the type $\overline{c}$ of P2 calling is strictly dominated $u_2(S, Call, \overline{c}) < u_2(S, Don't, \overline{c}) \, \forall S_1$.

For P1:

$u_1(Call, f_2^*, c_1) = 1 - c_1$

$u_1(Don't, f_2^*, c_1) = pu_1(Don't, f_2^*(\underline{c}), c_1) + (1-p)u_1(Don't, Don't, c_1)$

$< p + 1 - p(0) = p$

But $1 - c_1 > p \therefore f_1^*(c) = Call$.

For the type $\underline{c}$ of P2: $u_2(f_1^*, Call, \underline{c}) = 1 - \underline{c}$ and $u_2(f_1^*, Don't, \underline{c}) = 1$ because f_1^* =Call, so $f_2^*(\underline{c})$ =Don't.

[5] Remember the Expected Utility Theorem in first Lecture

[6] P2's cost is either the upper or lower bound but not in between.

20.3 Auctions

Check out www.paulklemperer.org

- You have n participants.
- Each participant has valuation v_i and submits bid b_i (his action)

The rules of the auction determine the probability $q_i(b_1, \dots, b_n)$ that agent i wins, and the expected price $p_i(b_1, \dots, b_n)$. His utility is $u_i = q_i v_i - p_i$.

21 Lecture 5 February 2007

Sample questions and answers.

21.1 Find all PSNE

P1 \ P2	X	Y	Z
A	30 / 10	20 / 0	30 / 20
B	35 / 15	40 / 10	40 / 10
C	25 / 25	25 / 5	25 / 5

Figure 31: Question 1 game

Answers: (A, Z), (B, Y) and (C,X); no one player can benefit from diverging from one of these.

21.2 Find Unique MSNE

P1 \ P2	X	Y	Z
A	10 / 20	20 / 10	1 / 1
B	20 / 10	10 / 20	1 / 1
C	1 / 1	1 / 1	0 / 0

Figure 32: Question 2 original game

Strategies X and Y strictly dominate strategy Z and strategies A and B strictly dominate strategy C. So strategies Z and C can be removed, simplifying the game.

Let α be the probability that player 1 chooses strategy A, and $1-\alpha$ is then the probability that player 1 chooses strategy B. Let β be the probability that player 2 chooses strategy X, and $1-\beta$ is then the probability that player 2 chooses strategy Y.

The utility for player 1 choosing A is: $U_1(A) = 20\beta + 10(1-\beta) = 10\beta - 10$. The utility for player 1 choosing B is: $U_1(B) = 10\beta + 20(1-\beta) = -10\beta + 20$. In a

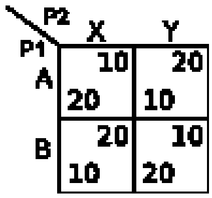

Figure 33: Question 2 - simplified game

mixed strategy Nash equilibrium for player 1 over strategies A and B the expected utilities are the same. So $10\beta - 10 = -10\beta + 20$ or $\beta = \frac{1}{2}$.

By a similar series of calculations $\alpha = \frac{1}{2}$. So the MSNE is $\left\{\left(A, \frac{1}{2}\right)\left(B, \frac{1}{2}\right)\left(X, \frac{1}{2}\right)\left(Y, \frac{1}{2}\right)\right\}$.

21.3 Extensive Form Game I

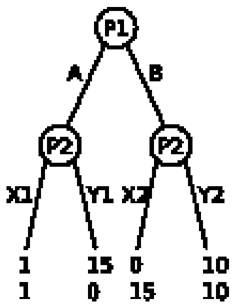

Figure 34: Question 3 game

For Player 2: $X1 \succ Y1$ and $X2 \succ Y2$.

$\therefore$ for Player 1: $A \succ B$.

This is an extensive form of the Prisoner's Dilemma.

21.4 Extensive Form Game II

Each time a player continues the game, the pot increases by 3. If a player stops the game at any time she keeps two thirds of the pot.

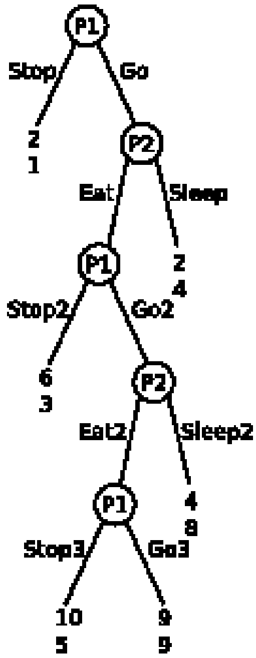

Figure 35: Question 4 game

Analysis: Stop3 $\succ$ Go3; Eat2 $\succ$ Sleep2; Stop2 $\succ$ Go2; Eat $\succ$ Sleep; $2 = 2$ so neither Stop nor Go dominate.

So there are two PSNE: {Stop} and {Go, Eat}.

21.5 Stackelberg Argument

After Fernando Vega-Redondo.

Question: Provide a rigorous verbal argument for the following general assertion: The leading firm in the model of Stackelberg always obtains at least as much profits as it would obtain in a Cournot framework with the same underlying data (demand, costs, etc.).

Answer: In the Stackelberg model if the leader chooses its Cournot output the followers will respond optimally by choosing their Cournot outputs. Thus, the Cournot outcome can be chosen by the Stackelberg leader, but since the leader chooses what is best for him then the Stackelberg outcome cannot yield lower profits for him than the Cournot one does.

21.6 Iterated Removal of Strictly Dominated Strategies

Question: Show that the IRSDS can be completed in a finite number of steps.

Answer: At each step IRSDS removes at least one strategy (or stops). There is a finite number of strategies $\therefore$ a finite number of steps.

Number of Steps: $< \Sigma_{i=1}^{n}(|S_i| - 1)$ where n is the number of players.