Lecture I-II: Motivation and Decision Theory

Markus M. Möbius

February 7, 2004

1 Two Motivating Experiments

Experiment 1 *Each of you (the students in this course) have to declare an integer between 0 and 100 to guess "2/3 of the average of all the responses". More precisely, each student who guesses the highest integer which is not higher than 2/3 of the average of all responses, will receive a prize of 10 Dollars.*

How should you play this game? A naive guess would be that other players choose randomly a strategy. In that case the mean in the game would be around 50 and you should choose 33. But you realize that other players make the same calculation - so nobody should say a number higher than 33. That means that you should not name a number greater than 22 - and so on. The winning number was 13. That means that people did this iteration about 3 times. But in fact, the stated numbers were all over the place - ranging from 0 to 40. That means that different students had different estimates of what their fellow students would do.

- Being aware of your fellow players' existence and trying to anticipate their moves is called strategic behavior. Game theory is mainly about designing models of strategic behavior.

- In this 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: (a) the player does not understand strategic behavior or (b) the player understands strategic behavior but has low confidence in the ability of other players to understand that this is a strategic game. Interestingly, most people knew at least one other person in the class (hence there was at least some degree of what a game theorist would call common knowledge of rationality).[1]

Experiment 2 *I am going to auction a textbook (Osborne's book). It costs about 60 Dollars on Amazon. Each of your can bid secretly on the book and the highest bidder wins the auction. However, all of you have to pay your bid regardless of whether you win or lose.*

In this game there is no optimal single bid for all players. You can check that for all cases where each player i bids some fixed bid b_i at least one of the players will regret her decision and try to reverse it - we say that there is no pure strategy Nash equilibrium in this game. Consider for example the case where all player bid 55 Dollars. Then some player should bid 55 and 5 cents. No equilibrium!

There is an equilibrium if we allow players to randomize. You can check that with two players who pick random numbers between 0 and 60 with equal probability no player would want to change her pick - all picks will give her zero profit in expectation.

The more players there are, the more the bid distribution is skewed towards 0 (check)! We will formally discuss mixed strategy Nash equilibria in a few lectures time.

2 What is game theory?

Definition 1 *Game theory is a formal way to analyze interaction among a group of rational agents who behave strategically.*

This definition contains a number of important concepts which are discussed in order:

Group: In any game there is more than one decision maker who is referred to as player. If there is a single player the game becomes a decision problem.

[1] There is common knowledge of rationality between players A and B if A knows that B is rational (and vice versa), if A knows that B knows that A is rational (and vice versa) etc.

Interaction: What one individual player does directly affects at least one other player in the group. Otherwise the game is simple a series of independent decision problems.

Strategic: Individual players account for this interdependence.

Rational: While accounting for this interdependence each player chooses her best action. This condition can be weakened and we can assume that agents are boundedly rational. Behavioral economics analyzes decision problems in which agents behave boundedly rational. Evolutionary game theory is game theory with boundedly rational agents.

Example 1 *Assume that 10 people go into a restaurant. Every person pays for her own meal. This is a decision problem. Now assume that everyone agrees before the meal to split the bill evenly amongst all 10 participants. Now we have a game.*

Game theory has found numerous applications in all fields of economics:

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

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

3. *IO:* Price depends not only on your output but also on the output of your competitor (market structure ...).

4. *PF:* My benefits from contributing to a public good depend on what everyone else contributes.

5. *Political Economy:* Who/what I vote for depends on what everyone else is voting for.

3 Decision Theory under Certainty

It makes sense to start by discussing trivial games - those we play against ourselves, e.g. decision problems. Agents face situations in which they have to make a choice. The actions of other agents do not influence my preference ordering over those choices - therefore there is no strategic interaction going on. Proper games will be discussed in the next lectures.

A decision problem $(A, \preceq)$ consists of a finite set of outcomes $A = \{a_1, a_2, .., a_n\}$ and a preference relation $\preceq$. The expression $a \preceq b$ should be interpreted as "b is at least as good as a". We expect the preference relation to fulfill two simple axioms:

Axiom 1 *Completeness. Any two outcomes can be ranked, e.g.* $a \preceq b$ *or* $b \preceq a$.

Axiom 2 *Transitivity implies that if* $a \geq b$ *and* $b \geq c$ *then* $a \geq c$.

Both axioms ensure that all choices can be ordered in a single chain without gaps (axiom 1) and without cycles (axiom 2).

Although the preference relation is the basic primitive of any decision problem (and generally observable) it is much easier to work with a consistent utility function $u : A \to \Re$ because we only have to remember n real numbers $\{u_1, u_2, .., u_n\}$.

Definition 2 *A utility function* $u : A \to \Re$ *is consist with the preference relationship of a decision problem* $(A, \preceq)$ *if for all* $a, b \in A$*:*

$$a \preceq b \quad \textit{if and only if} \quad u(a) \leq u(b)$$

Theorem 1 *Assume the set of outcomes is finite. Then there exists a utility function u which is consistent.*

Proof: The proof is very simple. Simple collect all equivalent outcomes in equivalence classes. There are finitely many of those equivalence classes since there are only finitely many outcomes. Then we can order these equivalence classes in a strictly increasing chain due to completeness and transitivity.

Note that the utility function is not unique. In fact, any monotonic transformation of a consistent utility function gives another utility function which is also consistent.

We can now define what a rational decision maker is.

Definition 3 *A rational decision maker who faces a decision problem* $(A, \preceq)$ *chooses an outcome* $a^* \in A$ *which maximizes his utility (or, equivalently, for each* $a \in A$ *we have* $a \preceq a^*$*).*

Remark 1 *When there are infinitely many choices we want to make sure that there is a continuous utility function. This requires one more axiom which makes sure that preferences are continuous. For that purpose, one has to define topology on the set of outcomes. We won't deal with that since we won't gain much insight from it.*

4 Decision Theory under Uncertainty

Lotteries are defined over the of outcomes A (which is again assumed to be finite to keep things simple).

Definition 4 *A simple lottery is defined as the set* $\{(a_1, p_1), (a_2, p_2), .. (a_n, p_n)\}$ *such that* $\sum_{i=1}^{n} p_i = 1$ *and* $0 \leq p_i \leq 1$. *In a simple lottery the outcome* a_i *occurs with probability* p_i.

When there are up to three outcomes we can conveniently describe the set of lotteries in a graphical way (see triangle).

Under certainty the preference relationship can still be written down explicitly for finite A (simply write down all of the $\frac{n(n+1)}{2}$ rankings). Under uncertainty there are suddenly infinitely many lotteries. This poses two problems. First of all, it's impractical to write a large number of lottery comparisons down. A second (and deeper) point is the observation that the preference relationship is in principle unobservable because of the infinite number of necessary comparisons.

John von Neumann and Oscar Morgenstern showed that under some additional restrictions on preferences over lotteries there exists a utility function over outcomes such that the expected utility of a lottery provides a consistent ranking of all lotteries.

Definition 5 *Assume a utility function* u *over the outcomes* A. *The expected utility of the lottery* $L = \{(a_1, p_1), (a_2, p_2), .. (a_n, p_n)\}$ *is defined as*

$$u(L) = \sum_{i=1}^{n} u(a_i) p_i$$

Before we introduce the additional axioms we discuss the notion of compound (two stage) lotteries.

Definition 6 *The compound lottery* $\tilde{L}$ *is expressed as* $\tilde{L} = \{(L_1, q_1), (L_2, 1 - q_1)\}$*. With probability* q_1 *the simple lottery* L_1 *is chosen and with probability* $1 - q_1$ *the simple lottery* L_2 *is chosen.*

Note, that we implicitly distinguish between simple and compound lotteries. Therefore, we allow that a simple lottery L might have the same outcome distribution as the compound lottery $\tilde{L}$ but $L \prec \tilde{L}$.

The first axiom assumes that only outcomes matter - the process which generates those outcomes is irrelevant.

Axiom 3 *Each compound lottery is equivalent to a simple lottery with the same distribution over final outcomes.*

In some books the equivalence of simple and compound lotteries is assumed in the definition of a lottery. However, it is useful to keep those types of lotteries separate because we know that the framing of a decision problem influences how people make choices (i.e. both the process and the final outcome distribution matter).

The next axiom is fairly uncontroversial.

Axiom 4 *Monotonicity. Assume that the lottery* L_1 *is preferred to lottery* L_2*. Then the compound lottery* $\{(L_1, \alpha), (L_2, 1 - \alpha)\}$ *is preferred to* $\{(L_1, \beta), (L_2, 1 - \beta)\}$ *if* $\alpha > \beta$*.*

Axiom 5 *Archimedian. For any outcomes* $a < b < c$ *there is some lottery* $L = \{(a, \alpha), (c, 1 - \alpha)\}$ *such that the agent is indifferent between* L *and* b*.*

The substitution axiom (also know as independence of irrelevant alternatives) is the most critical axiom.

Axiom 6 *Substitution. If lottery* L_1 *is preferred to lottery* L_2 *then any mixture of these lotteries with any other lottery* L_3 *preserves this ordering:*

$$\{(L_1, \alpha), (L_3, 1 - \alpha)\} \geq \{(L_2, \alpha), (L_3, 1 - \alpha)\}$$

This axiom is also known as independence of irrelevant alternatives.

Under these axioms we obtain the celebrated result due to John von Neumann and Oskar Morgenstern.

Theorem 2 *Under the above axioms an expected utility function exists.*

Proof: First of all we find the best and worst outcome b and w (possible because there are only finitely many outcomes). Because of the Archimedian axiom we can find a number α_a for each outcome a such that $L = \{(b, \alpha), (w, 1-\alpha)\}$. We can define a utility function over each outcome a such that $u(a) = \alpha$. Using the monotonicity axiom it can be shown that this number is unique. For each lottery we can now calculate its expected utility. It remains to be shown that this expected utility function is consistent with the original preferences. So take two lotteries L_1 and L_2 such that $L_1 \preceq L_2$. We can write $L_1 = \{(a_1, p_1), (a_2, p_2), ..(a_n, p_n)\}$ and $L_2 = \{(a_1, q_1), (a_2, q_2), ..(a_n, q_n)\}$. Now replace each outcome a_i by the above lotteries. The compound lottery can be rewritten as $L_1 = \{(b, \sum_{i=1}^{n} p_i u(a_i)), (w, 1 - \sum_{i=1}^{n} p_i u(a_i))\}$. Similarly, we get $L_2 = \{(b, \sum_{i=1}^{n} q_i u(a_i)), (w, 1 - \sum_{i=1}^{n} q_i u(a_i))\}$. By the monotonicity axiom we can deduce that $\sum_{i=1}^{n} p_i u(a_i) \leq \sum_{i=1}^{n} q_i u(a_i)$, e.g. $EU(L_1) \leq EU(L_2)$. QED

From now on all payoffs in our course will be assumed to represent vNM utility values. The expected payoff will be the expected utility.

4.1 Puzzles

EUT forms the basis of modern micro economics. Despite its success there are important behavioral inconsistencies related to it. Some of those we are going to discuss briefly before we turn our attention to proper games.

4.2 Allais Paradox

Consider the following choice situation (A) among two lotteries:

- Lottery A1 promises a sure win of 3000,
- Lottery A2 is a 80 percent chance to win 4000 (and zero in 20 percent of the cases).

Typically, A1 is strictly preferred to A2. Now, consider two further choice pairs (B) and (C):

- Lottery B1 promises a 90 percent chance of winning 3000,
- Lottery B2 is a 72 percent chance to win 4000.

This choice is included to see if there is a certainty effect.

- Lottery C1 promises a 25 percent chance of winning 3000,
- Lottery C2 is a 20 percent chance to win 4000.

Most people in our class now preferred C2 over C1.

It can be checked that the lotteries B_i and C_i are derived from A_i just by mixing the original lotteries with an irrelevant alternative - in the case of (B) there is a 10 percent chance of getting nothing and a 90 percent chance of getting (A), and in case of (C), there is a 75 percent chance of getting nothing.

The Allais paradox is the most prominent example for behavioral inconsistencies related to the von Neumann Morgenstern axiomatic model of choice under uncertainty. The Allais paradox shows that the significant majority of real decision makers orders uncertain prospects in a way that is inconsistent with the postulate that choices are independent of irrelevant alternatives.

There is an alternative explanation for the failure of EUT in this case. Assume, that agents face the compound lottery instead of the simple lotteries (B) and (C). Now the relationship to (A) is much more transparent - in fact, one could tell a story such as: "with 75 percent probability you are not invited to choose between these two outcomes, and with 25 percent probability you can choose either A1 or A2". It's likely that choices would be much more consistent now.

The standard explanation for the failure of EUT is peoples' inability to keep small probability differences apart. 80 percent and 100 percent 'looks' quite different and people focus on the probabilities. 20 percent and 25 percent 'looks' the same - so people focus on the values instead. Prospect theory (Kahnemann and Tversky) can deal with the Allais paradox by weighting probabilities accordingly.

4.3 Framing effects

Framing effects are preference reversals induced by changes in reference points.

Consider the following choice situation (A):

Pair 1: 600 people are struck with a disease that could kill. Vaccine 1 will save 400 lives for sure while the second one will either save no one (1/3) or will save everyone (with probability 2/3).

Table 1: McNeil, Pauker and Tversky (1988)

	Survival		Mortality		Both	
	Radiation	Surgery	Radiation	Surgery	Radiation	Surgery
immediate	100	90	0	10		
1 year	77	68	23	32		
5 year	22	34	78	66		
US	16	84	50	50	44	56
Israeli	20	80	45	56	34	66

Pair 2: 600 people are struck with a disease that could kill. Vaccine 1 will kill 200 people for sure while the second one implies a 2/3 chance that no one will die and a 1/3 chance that everyone will die.

Note that both situations are identical because save is equal to not kill. However, people tend to be risk averse in saving lives and risk loving if it is phrased in terms of losses (kills).

Preference reversals have real effects and do not just appear in cute examples. McNeil, Pauker and Tversky (1988) asked American doctors and Israeli medical students about how they would choose between two cancer treatments (surgery and radiation) - they presented one group with survival statistics, a second group with mortality statistics and a third group with both. Table 1 sums up their choices.

Lecture III: Normal Form Games, Rationality and Iterated Deletion of Dominated Strategies

Markus M. Möbius

February 19, 2004

Readings:

- Gibbons, sections 1.1.A and 1.1.B
- Osborne, sections 2.1-2.5 and section 2.9

1 Definition of Normal Form Game

Game theory can be regarded as a multi-agent decision problem. It's useful to define first exactly what we mean by a *game*.

Every normal form (strategic form) game has the following ingredients.

1. There is a list of players $D = \{1, 2, .., I\}$. We mostly consider games with just two players. As an example consider two people who want to meet in New York.

2. Each player i can choose actions from a strategy set S_i. To continue our example, each of the players has the option to go the Empire State building or meet at the old oak tree in Central Park (where ever that is ...). So the strategy sets of both players are $S_1 = S_2 = \{E, C\}$.

3. The outcome of the game is defined by the 'strategy profile' which consists of all strategies chosen by the individual players. For example, in our game there are four possible outcomes - both players meet at the Empire state building (E, E), they miscoordinate, (E, C) and (C, E), or

they meet in Central Park (C, C). Mathematically, the set of strategy profiles (or outcomes of the game) is defined as

$$S = S_1 \times S_2$$

In our case, S has order 4. If player 1 can take 5 possible actions, and player 2 can take 10 possible actions, the set of profiles has order 50.

4. Players have preferences over the outcomes of the play. You should realize that players cannot have preferences over the actions. In a game my payoff depends on your action. In our New York game players just want to be able to meet at the same spot. They don't care if they meet at the Empire State building or at Central Park. If they choose E and the other player does so, too, fine! If they choose E but the other player chooses C, then they are unhappy. So what matters to players are outcomes, not actions (of course their actions influence the outcome - but for each action there might be many possible outcomes - in our example there are two possible outcomes per action). Recall, that we can represent preferences over outcomes through a utility function. Mathematically, preferences over outcomes are defined as:

$$u_i : S \to R$$

In our example, $u_i = 1$ if both agents choose the same action, and 0 otherwise.

All this information can be conveniently expressed in a game matrix as shown in figure 1:

A more formal definition of a game is given below:

Definition 1 *A normal (strategic) form game G consists of*

- *A finite set of agents $D = \{1, 2, .., I\}$.*
- *Strategy sets $S_1, S_2, .., S_I$*
- *Payoff functions $u_i : S_1 \times S_2 \times ..S_I \to R$ $(i = 1, 2, .., n)$*

We'll write $S = S_1 \times S_2 \times .. \times S_I$ and we call $s \in S$ a strategy profile $(s = (s_1, s_2, .., s_I))$. We denote the strategy choices of all players except player i with s_{-i} for $(s_1, s_2, .., s_{i-1}, s_{i+1}, ..s_I)$.

Figure 1: General 2 by 2 game

	E	C
E	1,1	0,0
C	0,0	1,1

2 Some Important Games

We already discussed coordination games. These are interesting games, because players have an incentive to work together rather than against each other. The first games analyzed by game theorists were just the opposite - zero sum games, where the sum of agents' utilities in each outcome sums up to zero (or a constant).

2.1 Zero-Sum Games

Zero-sum games are true games of conflict. Any gain on my side comes at the expense of my opponents. Think of dividing up a pie. The size of the pie doesn't change - it's all about redistribution of the pieces between the players (tax policy is a good example).

The simplest zero sum game is matching pennies. This is a two player game where player 1 get a Dollar from player 2 if both choose the same action, and otherwise loses a Dollar:

	H	T
H	1,−1	−1,1
T	−1,1	1,−1

2.2 Battle of the Sexes

This game is interesting because it is a coordination game with some elements of conflict. The idea is that a couple want to spend the evening together. The wife wants to go to the Opera, while the husband wants to go to a football game. Each get at least some utility from going together to at least one of the venues, but each wants to go their favorite one (the husband is player 1 - the column player).

	F	O
F	2,1	0,0
O	0,0	1,2

2.3 Chicken or Hawk versus Dove

This game is an anti-coordination game. The story is that two teenagers drive home on a narrow road with their bikes, and in opposite directions. None of them wants to go out of the way - whoever 'chickens' out loses his pride, while the tough guy wins. But if both stay tough, then they break their bones. If both go out of the way, none of the them is too happy or unhappy.

	t	c
t	−1,−1	10,0
c	0,10	5,5

2.4 Prisoner's Dilemma

This game might be the most famous of all. It's the mother of all cooperation games. The story is that two prisoners are interrogated. If both cooperate with the prosecution they get of with 1 year in prison. If both give each other away (defect) they get 3 years in prison each. If one cooperates and the other guy defects, then the cooperating guy is thrown into prison for 10 years, and the defecting guy walks free.

	C	D
C	3,3	−1,4
D	4,−1	0,0

Note, that the best outcome in terms of welfare is if both cooperate. The outcome (D, D) is worst in welfare terms, and is also Pareto dominated by (C, C) because both players can do better. So clearly, (D, D) seems to be a terrible outcome overall.

Some examples of Prisoner's dilemmas are the following:

- Arms races. Two countries engage in an expensive arms race (corresponds to outcome D,D). They both would like to spend their money on (say) healthcare, but if one spends the money on healthcare and the other country engages in arms build-up, the weak country will get invaded.

- Missile defence. The missile defence initiative proposed by the administration is interpreted by some observers as a Prisoner's dilemma. Country 1 (the US) can either not build a missile defence system (strategy C) or build one (strategy D). Country 2 (Russia) can either not build any more missiles (strategy C) or build lots more (strategy D). If the US does not build a missile system, and Russia does not build more missiles then both countries are fairly well off. If Russia builds more missiles and the US has no defence then the US feels very unsafe. If the US builds a missile shield, and Russia does not missiles then the

US is happy but Russia feels unsafe. If the US builds missile defence and Russia builds more missiles then they are equally unsafe as in the (C,C) case, but they are much less well off because they both have to increase their defence budget.

- Driving a big SUV can be a Prisoner's Dilemma. I want my car to be as safe as possible and buy an SUV. However, my neighbors who has a Volkswagen Beetle suddenly is much worse off. If she also buys an SUV she will be again safe but in this case both of us have to drive a big car and buy a lot of gas.

2.5 Cournot Competition

This game has an infinite strategy space. Two firms choose output levels q_i and have cost function $c_i(q_i)$. The products are undifferentiated and market demand determines a price $p(q_1+q_2)$. Note, that this specification assumes that the products of both firms are perfect substitutes, i.e. they are homogenous products.

$$\begin{aligned} D &= \{1,2\} \\ S_1 &= S_2 = R^+ \\ u_1(q_1,q_2) &= q_1 p(q_1+q_2) - c_1(q_1) \\ u_2(q_1,q_2) &= q_2 p(q_1+q_2) - c_2(q_2) \end{aligned}$$

2.6 Bertrand Competition

Bertrand competition is in some ways the opposite of Cournot competition. Firms compete in a homogenous product market but they set prices. Consumers buy from the lowest cost firm.

Remark 1 *It is interesting to compare Bertrand and Cournot competition with perfect competition analyzed in standard micro theory. Under perfect competition firms are price takers i.e. they cannot influence the market. In this case there is not strategic interaction between firms - each firm solves a simple profit maximization problem (decision problem). This is of course not quite true since the auctioneer does determine prices such that demand and supply equalize.*

3 What is a Game?

Before moving on it is useful to discuss two possible interpretations of normal-form games.

1. The normal form game is simply a game played once in time between a set of players.

2. The normal form game is 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. Examples include driving on the right-hand side (a coordination game continuously played between motorists in the US). Random matching is important here: if the stage is played repeatedly with the same player we have a repeated extensive form game (discussed in future lectures) and new strategic considerations arise.

4 Two Brief Experiments

Experiment 1 *(not done this year - reported from 2001 spring semester) Student were asked which strategy they would play in the Prisoner's dilemma. The class was roughly divided in half - we calculated the expected payoff from both strategies if people in the class would be randomly matched against each other. We found that strategy D was better - this is unsurprising as we will see later since strategy C is strictly dominated by strategy D.*

Experiment 2 ***Iterated Deletion Game*** *Class was asked to choose a strategy for player 1 in the game below. No student chose strategy A, 7 students chose B, 11 students chose C and 4 students chose D.*

The reason people gave were interesting:

- One student from the C-group said that that row gave the highest sum of payoff when the row payoffs were added. Note, that this reasoning is correct IF player 2 randomizes across her four strategies.

- In previous years, people often noted that C looks safe - it always gives the highest or second-highest payoff.

- One student chose B because in 2 our of 4 cases it is a best response while A,C,D are best responses in only 1 out of four cases.

	A	B	C	D
A	5,2	2,6	1,4	0,4
B	0,0	3,2	2,1	1,1
C	7,0	2,2	1,5	5,1
D	9,5	1,3	0,2	4,8

In 2001 the results were similar: no student chose strategy A which is weakly dominated by C. 2 students chose B, 9 students chose C because it looked 'safe' and 16 students chose D because of the high payoffs in that row.

It turns out that only (B,B) survives iterated deletion (see below).

5 Iterated Deletion of Dominated Strategies

How do agents play games? We can learn a lot by exploiting the assumption that players are rational and that each player knows that other players are rational. Sometimes this reasoning allows us to 'solve' a game.

5.1 Rational Behavior

Assume that agent i has belief μ_i about the play of her opponents. A belief is a probability distribution over the strategy set S_{-i}.

Definition 2 *Player i is rational with beliefs μ_i if*

$$s_i \in \arg\max_{s_i'} E_{\mu_i(s_{-i})} u_i\left(s_i', s_{-i}\right),$$

or alternatively

$$s_i \textit{maximizes} \sum_{s_{-i}} u_i\left(s_i', s_{-i}\right)\mu_i\left(s_{-i}\right).$$

Note, that player i faces a simple decision problem as soon as she has formed her belief μ_i.

An example illustrates this point: assume that I believe in the New-York game that my friend will come to the Empire state building with 60 percent probability and to Central Park with 40 percent probability. If I go to central park I induce the following lottery L^C over outcomes of the game: with 60 percent probability I will see the outcome (C, E) and with 40 percent (C, C):

$$L^C = 0.6(C, E) \oplus 0.4(C, C) \tag{1}$$

Thanks to our expected utility theorem we can easily evaluate the expected utility of this lottery which is .4! Similarly, we can evaluate that playing E induces a lottery with expected value .6. So I am rational and have the above belief then I should choose E.

Definition 3 *Strategy s_i is strictly dominated for player i if there is some $s_i' \in S_i$ such that*

$$u_i\left(s_i', s_{-i}\right) > u_i\left(s_i, s_{-i}\right)$$

for all $s_{-i} \in S_{-i}$.

Note that the inequality is strict for all s_{-i}. A strategy is weakly dominated if the inequality is weak for all s_{-i} and strict for at least one s_{-i}.

Proposition 1 *If player i is rational he will not play a strictly dominated strategy.*

Proof: If strategy s_i is strictly dominated by strategy s_i' we can deduce that for any belief of player i we have $E_{\mu_i(s_{-i})} u_i\left(s_i', s_{-i}\right) > E_{\mu_i(s_{-i})} u_i\left(s_i, s_{-i}\right)$.

5.2 Iterated Dominance

The hardest task in solving a game 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 Prisoner's Dilemma.

Cooperating is a dominated strategy. A rational player would therefore never cooperate. This solves the game since every player will defect. Notice that I don't have to know anything about the other player. This prediction is interesting because it is the worst outcome in terms of joint surplus and it would be Pareto improving if both players would cooperate. This result highlights the value of commitment in the Prisoner's dilemma - commitment consists of credibly playing strategy C. For example, in the missile defence example the ABM treaty (prohibits missile defence) and the START II agreement (prohibits building of new missiles) effectively restrict both country's strategy sets to strategy C.

Now look at the next game.

	L	M	R
U	2,2	1,1	4,0
D	1,2	4,1	3,5

1. If the column player is rational he shouldn't play M

2. Row player should realize this if he know that the other player is rational. Thus he won't play D.

3. Column player should realize that R knows that C is rational. If he knows that R is rational he knows that R won't play D. Hence he won't play R. This leaves (U,L) as only outcome for rational players.

It's worth while to discuss the level of knowledge required by players. R has to know that C is rational. C has to know that R knows that C is rational. This latter knowledge is a 'higher order' form of knowledge. It's not enough to know that my opponent is rational - I also have to be sure that my opponent knows that I am rational. There are even higher order types of knowledge. I might know that my opponent is rational and that he knows that I am. But maybe he doesn't know that I know that he knows.

The higher the order of knowledge the more often the process of elimination can be repeated. For example, the game of experiment 2 can be solved by the iterated deletion of dominated strategies.

If rationality is *common knowledge* we can repeat iterated deletion of dominated strategies indefinitely - I know that my opponent is rational, that he knows that I am rational, that I know that he knows that I know that I am rational etc.

We will usually assume that rationality is common knowledge and that we can therefore perform iterated deletion of strictly dominated strategies as often as we like.

5.2.1 Other Models of Knowledge

To illustrate the importance of "Rationality is common knowledge" assumption we discuss an alternative model of knowledge in the game above:

- Both players 1 and 2 are rational.
- Player 1 thinks that player 2 is clueless and randomizes across his strategies with equal probability.
- Player 2 thinks that player 1 is rational and that player 1 thinks he is randomizing.

In this case player 1 will optimally choose action D which gives her the highest average payoff. Player 2 will correctly anticipate this and choose action R.

There is nothing wrong with this alternative model. However, there are some potentially troublesome inconsistencies:

- Assume that we adopt the repeated game interpretation of a normal-form game: the above game is one instance in a repeated game between a large population of players 1's and player 2's who are repeatedly matched against each other.
- Assume that player 1's and player 2's have a model of knowledge as described above and play (D,R) all the time.
- A player 1 should realize after a while that player 2's consistently play R. Hence they should update their model of knowledge and conclude that player 2 is not as clueless as he assumed.

- In fact, once player 1 concludes that player 2 consistently plays strategy R he should switch to U which will increase his winnings.

When rationality is common knowledge these types of inconsistences will not emerge. That's one of the reasons why it is so commonly used in game theory.

However, the example of experiment 2 also showed that players are often do not seem to be able to do deleted iteration in their head for more than one or two rounds. It would be interesting to repeat the stage game many times with random matching to see if player 1's will switch to strategy B gradually.

5.3 Formal Definition Of Iterated Dominance

- **Step I:** Define $S_i^0 = S_i$
- **Step II:** Define

$$S_i^1 = \left\{ s_i \in S_i^0 \,\middle|\, \nexists s_i' \in S_i^0 u_i\left(s_i', s_{-i}\right) > u_i\left(s_i, s_{-i}\right) \forall s_{-i} \in S_{-i}^0 \right\}$$

- **Step k+1:** Define

$$S_i^{k+1} = \left\{ s_i \in S_i^k \,\middle|\, \nexists s_i' \in S_i^k u_i\left(s_i', s_{-i}\right) > u_i\left(s_i, s_{-i}\right) \forall s_{-i} \in S_{-i}^k \right\}$$

  S_i^{k+1} is the set still not strictly dominated when you know your opponent uses some strategy in S_{-i}^k.

  Note restrictions S_{-i}^0, S_{-i}^1, ..

  Players know that opponents are rational, know that opponents know that they are rational ..., e.g. rationality is common knowledge.

- **Step ∞:** Let $S_i^\infty = \bigcap_{k=1}^\infty S_i^k$.

  Note, that the process must stop after finitely many steps if the strategy set is finite because the sets can only get smaller after each iteration.

Definition 4 *G is solvable by pure strategy iterated strict dominance if S^∞ contains a single strategy profile.*

Note:

- Most games are not dominance solvable (coordination game, zero sum game).

- We have not specified the order in which strategies are eliminated. You will show in the problem set that the speed and order of elimination does not matter.

  Intuition: Assume that you don't delete all dominated strategies at one stage of the iteration. Will you do so later? Sure you will: a dominated strategy will still be dominated; at most you have deleted a few more of your opponents strategies in the meantime which will make it even 'easier' to dominate the strategy.

- The same is not true for the elimination of *weakly dominated strategies* as the next example shows.

	L	R
T	1,1	0,0
M	1,1	2,1
B	0,0	2,1

We can first 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. So weak elimination does not deliver consistent results and is therefore generally a less attractive option than the deletion of strictly dominated strategies.

Intuition: player 2's strategies R and L give the same payoff when 1 plays M. This can lead to L weakly dominating R or vice versa depending on whether player 1's strategies B or T are deleted first. If L would strictly dominate R (or vice versa) this could not be the case: L would *always* dominate R regardless of how many of 1's strategies are deleted.

Remark 2 *With finite strategy sets the set S^∞ is clearly non-empty because after each stage there must be some dominant strategy left (in fact in a 2-player n by m game the iterative process has to stop after at most $n+m-2$ steps).*

Remark 3 *(for the mathematically inclined only) For infinite strategy sets it is not obvious that the iterative process will result in a non-empty set. There are examples of sequences of nested sets whose intersection is empty:*

$$S^n = \left(0, \left(\frac{1}{2}\right)^n\right) \tag{2}$$

The intersection S^∞ of all these open intervals is the empty set. One way to get a non-empty set S^∞ is make sure that the sets S^k are closed and bounded sets and hence compact (assuming a final-dimensional action space). Typically, this will be the case if the utility functions are continuous in players' strategies (as in Cournot game).

6 Example: Cournot Competition

Cournot competition with two firms can be solved by iterated deletion in some cases. Specifically, we look at 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. It will be usefull to calculate the 'best-response' function $BR(q_j)$ 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 easily show that the best-response function for both firms (there is symmetry!) is

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

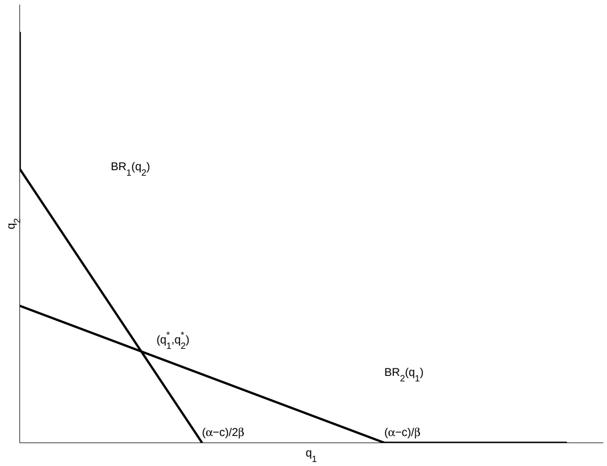

The best-response function is decreasing in my belief of the other firm's action. Note, that for $q_j > \frac{\alpha-c}{\beta}$ firm i makes negative profits even if it chooses the profit maximizing output. It therefore is better off to stay out of the market and choose $q_i = 0$.

- Initially, firms can set any quantity, i.e. $S_1^0 = S_2^0 = \Re^+$. However, the best-responses of each firm to any belief has to lie in the interval $\left[\underline{q}, \overline{q}\right]$ with $\underline{q} = 0$ and $\overline{q} = \frac{\alpha-c}{2\beta}$. All other strategies make negative profits, are therefore dominated by some strategy inside this interval, and eliminated.

- In the second stage only the strategies $S_1^1 = S_2^1 = [BR_1(\overline{q}), BR_1(\underline{q})]$ survive (check for yourself graphically in the picture above!). Because the BR functions are symmetric we can simply write:

$$S_1^1 = S_2^1 = [BR(\overline{q}), BR(\underline{q})] \tag{3}$$

  How do we get this set? These is set of possible best responses to *some* strategy played by the other player in S_i^0. *All other strategies are never best-responses and hence are dominated.*

- In the third stage we get:

$$\begin{aligned} S_1^3 = S_2^3 &= [BR_2(BR_1(\underline{q})), BR_2(BR_1(\overline{q}))] \\ &= [BR(BR(\underline{q})), BR(BR(\overline{q}))] \end{aligned} \tag{4}$$

  (note, that the BR function is decreasing which causes the reversal of bounds in each iteration!).

- Therefore in the $2k+1$th stage only strategies in
$$S_1^{2k+1} = S_2^{2k+1} = \left[BR_2\left(..BR_1\left(\underline{q}\right)\right), BR_2\left(..BR_1\left(\overline{q}\right)\right)\right] \tag{5}$$
survive.

It's easy to show graphically that this interval shrinks in each iteration and that the two limits converge to the intersection $q_1^* = q_2^*$ of both best response functions where $q_2^* = BR_2\left(q_1^*\right)$. Therefore, the Cournot game is solvable through the iterated deletion of strictly dominated strategies.

A precise proof of this claim follows below - this is NOT required material for the class.

- Let's focus on the strategy set $S^{2k+1} = [x_k, y_k]$ where:
$$\begin{aligned} x_k &= BR(BR(x_{k-1})) \\ y_k &= BR(BR(y_{k-1})) \\ x_0 &= 0 \\ y_0 &= \frac{\alpha - c}{2\beta} \end{aligned} \tag{6}$$

- The expression $x_k = BR(BR(x_{k-1}))$ can be calculated
$$x_k = \frac{\alpha - c}{4\beta} + \frac{x_{k-1}}{4} \tag{7}$$

- (x_k) is an increasing sequence because the strategy sets are nested and these are the lower bounds. The sequence is also bounded from above and hence has to converge. So assume $x_k \to x^*$. The following has to hold:
$$\begin{aligned} \lim_{k\to\infty} x_k &= \lim_{k\to\infty}\left[\frac{\alpha - c}{4\beta} + \frac{x_k}{4}\right] \\ x^* &= \frac{\alpha - c}{4\beta} + \frac{x^*}{4} \\ x^* &= \frac{\alpha - c}{3\beta} \end{aligned} \tag{8}$$

The same is true for y_k which converges to precisely the same limit. This means that the process of iterated deletion eventually collapses to the point where both firm set output to x^* which is exactly at the intersection of both firms' best response curves.

Remark 4 *It can be shown that the same game with three firms is NOT dominance solvable. You have to show that on the problem set!*

Lecture IV: Nash Equilibrium

Markus M. Möbius

February 19, 2004

Readings:

- Gibbons, sections 1.1.C and 1.2.B
- Osborne, sections 2.6-2.8 and sections 3.1 and 3.2

Iterated dominance is an attractive solution concept because it only assumes that all players are rational and that it is common knowledge that every player is rational (although this might be too strong an assumption as our experiments showed). It is essentially a constructive concept - the idea is to restrict my assumptions about the strategy choices of other players by eliminating strategies one by one.

For a large class of games iterated deletion of strictly dominated strategies significantly reduces the strategy set. However, only a small class of games are solvable in this way (such as Counot competition with linear demand curve).

Today we introduce the most important concept for solving games: Nash equilibrium. We will later show that all finite games have at least one Nash equilibrium, and that the set of Nash equilibria is a subset of the strategy profiles which survive iterated deletion. In that sense, Nash equilibrium makes stronger predictions than iterated deletion would but it is not excessively strong in the sense that it does not rule out any equilibrium play for some games.

Definition 1 *A strategy profile s^* is a pure strategy Nash equilibrium of G if and only if*

$$u_i\left(s_i^*, s_{-i}^*\right) \geq u_i\left(s_i, s_{-i}^*\right)$$

for all players i and all $s_i \in S_i$.

Definition 2 *A pure strategy NE is strict if*

$$u_i\left(s_i^*, s_{-i}^*\right) > u_i\left(s_i, s_{-i}^*\right)$$

A Nash equilibrium captures the idea of equilibrium. Both players know what strategy the other player is going to choose, and no player has an incentive to deviate from equilibrium play because her strategy is a best response to her belief about the other player's strategy.

1 Games with Unique NE

1.1 Prisoner's Dilemma

	C	D
C	3,3	−1,4
D	4,−1	0,0

This game has the unique Nash equilibrium (D,D). It is easy to check that each player can profitably deviate from every other strategy profile. For, example (C,C) cannot be a NE because player 1 would gain from playing D instead (as would player 2).

1.2 Example II

	L	M	R
U	2,2	1,1	4,0
D	1,2	4,1	3,5

In this game the unique Nash equilibrium is (U,L). It is easy to see that (U,L) is a NE because both players would lose from deviating to any other strategy.

To show hat there are no other Nash equilibria we could check each other strategy profile, or note that $S_1^\infty = \{U\}$ and $S_2^\infty = \{L\}$ and use:

Proposition 1 *If s^* is a pure strategy Nash equilibrium of G then $s^* \in S^\infty$.*

Proof: Suppose not. Then there exists T such that $s^* \in S_1^T \times ... \times S_I^T$ but $s^* \notin S_1^{T+1} \times ... \times S_I^{T+1}$. The definition of ISD implies that there exists $s_i' \in S_i^T \subseteq S_i$ such that $u_i(s_i', s_{-i}) > u_i(s_i^*, s_{-i})$ for all $s_{-i} \in S_{-i}^T$. Therefore there exists a $s_i' \in S_i$ such that $u_i\left(s_i', s_{-i}^*\right) > u_i\left(s_i^*, s_{-i}^*\right)$ which contradicts that s^* was a NE.

1.3 Cournot Competition

Using our new result it is easy to see that the unique Nash equilibrium of the Cournot game with linear demand and constant marginal cost is the intersection of the two BR functions since this was the only profile which survived IDSDS.

A more direct proof notes that any Nash equilibrium has to lie on the best response function of both players by the definition of NE:

Lemma 1 *(q_1^*, q_2^*) is a NE if and only if $q_i^* \in BR_i(q_{-i})$ for all i.*

We have derived the best response functions of both firms in previous lectures (see figure 1).

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

The NE is the solution to $q_1 = BR_1(q_2)$ and $q_2 = BR_2(q_1)$. This system has exactly one solution. This can be shown algebraically or simply by looking at the intersections of the BR graphs in figure 1. Because of symmetry we know that $q_1 = q_2 = q^*$. Hence we obtain:

$$q^* = \frac{\alpha - c}{\beta} - \frac{q^*}{2}$$

This gives us the solution $q^* = \frac{2(\alpha-c)}{3\beta}$.

If both firms are not symmetric you have to solve a system of two equations with two unknowns (q_1 and q_2).

Figure 1: BR functions of two firm Cournot game

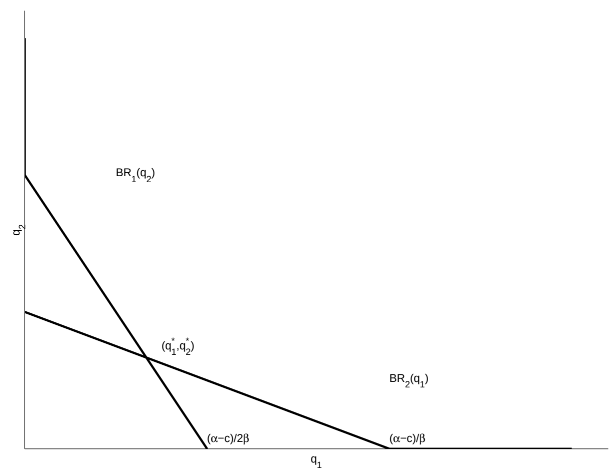

1.4 Bertrand Competition

Recall the Bertrand price setting game between two firms that sell a homogenous product to consumers.

Two firms can simultaneously set any positive price p_i $(i = 1, 2)$ and produce output at constant marginal cost c. They face a downward sloping demand curve $q = D(p)$ and consumers always buy from the lowest price firm (this would not be true if the goods weren't homogenous!). Therefore, each firm faces demand

$$D_i(p_1, p_2) = \begin{cases} D(p_i) & \text{if } p_i < p_j \\ D(p_i)/2 & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases}$$

We also assume that $D(c) > 0$, i.e. firms can sell a positive quantity if they price at marginal cost (otherwise a market would not be viable - firms couldn't sell any output, or would have to accumulate losses to do so).

Lemma 2 *The Bertrand game has the unique NE* $(p_1^*, p_2^*) = (c, c)$.

Proof: First we must show that (c,c) is a NE. It is easy to see that each firm makes zero profits. Deviating to a price below c would cause losses to the deviating firm. If any firm sets a higher price it does not sell any output and also makes zero profits. Therefore, there is no incentive to deviate.

To show uniqueness we must show that any other strategy profile (p_1, p_2) is not a NE. It's easiest to distinguish lots of cases.

Case I: $p_1 < c$ **or** $p_2 < c$ In this case one (or both players) makes negative losses. This player should set a price above his rival's price and cut his losses by not selling any output.

Case II: $c \leq p_1 < p_2$ **or** $c \leq p_2 < p_1$ Assume first that $c < p_1 < p_2$ or $c < p_2 < p_1$. In this case the firm with the higher price makes zero profits. It could profitably deviate by setting a price equal to the rival's price and thus capture at least half of his market, and make strictly positive profits. Now consider the case $c = p_1 < p_2$ or $c = p_2 < p_1$. Now the lower price firm can charge a price slightly above marginal cost (but still below the price of the rival) and make strictly positive profits.

Case III: $c < p_1 = p_2$ Firm 1 could profitably deviate by setting a price $p_1 = p_2 - \epsilon > c$. The firm's profits before and after the deviation are:

$$\pi_B = \frac{D(p_2)}{2}(p_2 - c)$$
$$\pi_A = D(p_2 - \epsilon)(p_2 - \epsilon - c)$$

Note, that the demand function is decreasing. We can therefore deduce:

$$\Delta\pi = \pi_A - \pi_B \geq= \frac{D(p_2)}{2}(p_2 - c) - \epsilon D(p_2)$$

This expression (the gain from deviating) is strictly positive for sufficiently small ϵ. Therefore, (p_1, p_2) cannot be a NE.

Remark 1 *In problem 11 of problem set 1 you had to solve for the unique Nash equilibrium when one firm (say 2) has higher marginal cost $c_2 > c_1$. Intuitively the price in the unique NE should be just below c_2 - this would keep firm 2 out of the market and firm 1 has no incentive to cut prices any further. However, if firms can set any real positive price there is no pure NE. Assume $c_2 = 10$. If firm 1 sets prices at 9.99 it can do better by setting them at 9.999 etc. Therefore, we have to assume that the pricing is discrete, i.e.*

can only be done in multiples of pennies say. In this way, the unique NE has firm 1 setting a price $p_1 = c_2$ *minus one penny.*

Food for Thought: How would you modify the Bertrand game to make it solvable through IDSDS? *Hint: You have to (a) discretize the strategy space, and (b) assume that* $D(p) = 0$ *for some sufficiently high price.*

Lecture IV: Nash Equilibrium II - Multiple Equilibria

Markus M. Möbius

February 24, 2004

- Gibbons, sections 1.1.C and 1.2.B
- Osborne, sections 2.6-2.8 and sections 3.1 and 3.2

1 Multiple Equilibria I - Coordination

Lots of games have multiple Nash equilibria. In this case the problem arises how to select between different equilibria.

1.1 New-York Game

Look at this simple coordination game:

	E	C
E	1,1	0,0
C	0,0	1,1

This game has two Nash equilibria - (E,E) and (C,C). In both cases no player can profitably deviate. (E,C) and (C,E) cannot be NE because both players would have an incentive to deviate.

1.2 Voting Game

Three players simultaneously cast ballots for one of three alternatives A,B or C. If a majority chooses any policy that policy is implemented. If the votes split 1-1-1 we assume that the status quo A is retained. Suppose the preferences are:

$$u_1(A) > u_1(B) > u_1(C)$$
$$u_2(B) > u_2(C) > u_2(A)$$
$$u_3(C) > u_3(A) > u_3(B)$$

Claim 1 *The game has several Nash equilibria including* (A, A, A), (B, B, B), (C, C, C),(A, B, A), *and* (A, C, C).

Informal Proof: In the first three cases no single player can change the outcome. Therefore there is no profitable deviation. In the last two equilibria each of the two A and two C players, respectively, is pivotal but still would not deviate because it would lead to a less desirable result.

1.3 Focal Points

In the New York game there is no sense in which one of the two equilibria is 'better' than the other one.

For certain games Schelling's (1961) concept of a tipping point can be a useful way to select between different Nash equilibria. A focal point is a NE which stands out from the set of NE - in games which are played frequently social norms can develop. In one-shot games strategies which 'stand out' are frequently played. In both cases, players can coordinate by using knowledge and information which is not part of the formal description of our game.

An example of a social norm is the fact that Americans drive on the right hand side of the road. Consider the following game. Tom and Jerry drive in two cars on a two lane road and in opposite directions. They can drive on the right or on the left, but if they mis-coordinate they cause a traffic crash. The game can be represented as follows:

	R	L
R	1,1	0,0
L	0,0	1,1

We expect both drivers to choose (R,R) which is the social norm in this game.

Next, let's conduct a class experiment.

Class Experiment 1 *You have to coordinate on what of the following four actions - coordinating with your partner gives you a joint payoff of 1 Dollar. Otherwise you both get zero Dollars. The actions are* $\{Fiat95, Fiat97, Saab98, Fiat98\}$.

We played the game with four pairs of students - three pairs coordinated on SAAB98, one pair did not coordinate.

This experiment is meant to illustrate that a strategy which looks quite distinct from the set of other available strategies (here, Fiats) can be a focal point in a one-shot game (when no social norm can guide us).

2 Multiple Equilibria II - Battle of the Sexes

The payoffs in the Battle of the Sexes are assumed to be Dollars.

	F	O
F	2,1	0,0
O	0,0	1,2

(F,F) and (O,O) are both Nash equilibria of the game. The Battle of the Sexes is an interesting coordination game because players are not indifferent on which strategy to coordinate. Men want to watch Football, while Women want to go to the Opera.

Class Experiment 2 *You are playing the battle of the sexes. You are player 1. What will you play?*

Last year: We divided the class up into men and women. 18 out of 25 men (i.e. 72 percent) chose the action which in case of coordination would give them the higher payoff. In contrast, only 6 out of 11 women did the same. These results replicate similar experiments by Rubinstein at Princeton and Tel Aviv University. Men are simply more aggressive creatures...

When we aggregate up we found that 24 out of 36 people (66 percent) play the preferred strategy in BoS.

Because there is an element of conflict in the BoS players use the framing of the game in order to infer the strategies of their opponent. In the following experiments the underlying game is always the above BoS. However, in each case the results differ significantly from the basic experiment we just conducted. This tells us that players signal their intention to each other, and that the normal strategic form does not capture this belief formation process.

Class Experiment 3 *You are player 1 in the Battle of the sexes. Player 2 makes the first move and chooses an action. You cannot observe her action until you have chosen your own action. Which action will you choose.*

Last year: Now a significantly higher number of students (17 instead of 12) choose the less desirable action (O). Note, that the game is still the same simultaneous move game as before. However, players seem to believe that player 1 has an advantage by moving first, and they are more likely to 'cave in'.

Class Experiment 4 *You are player 1 in the Battle of the sexes. Before actually playing the game, your opponent (player 2) had an opportunity to make an announcement. Her announcement was "I will play O". You could not make a counter-announcement. What will you play ?*

Now 35 out of 36 students chose the less desirable action. The announcement seems to strengthen beliefs that the other player will choose O.

This kind of communication is called *cheap talk* because this type of message is costless to the sender. For exactly this reason, it should not matter for the analysis of the game. To see that, simply expand the strategy set of player 2. Instead of strategies F and O she now has 4 strategies - Ff, Fo, Of, Oo - where strategy Ff means that player 2 plays F and announces to play f, while Of means that player 2 announces O and plays f. Clearly, the strategies Of and Oo yield exactly the same payoffs when played against any strategy of player 1. Therefore, the game has exactly the same NE as before. However, the announcement seems to have successfully signalled to player 1 that player 2 will choose her preferred strategy.

Class Experiment 5 *Two players are playing the Battle of the Sexes. You are player 1. Before actually playing the game, player 2 (the wife) had an opportunity to make a short announcement. Player 2 choose to remain silent. What is your prediction about the outcome of the game?*

Less than 12 people choose the less desirable action in this case. Apparently, silence is interpreted as weakness.

3 Multiple Equilibria III - Coordination and Risk Dominance

The following symmetric coordination game is given.

	A	B
A	9,9	−15,8
B	8,−15	7,7

Class Experiment 6 *Ask class how many would choose strategy A in this coordination game.*

Observations:

1. This game has the two Nash equilibria, namely (A,A) and (B,B). Coordinating on A Pareto dominates coordination on B. Unlike the New York and the Battle of the Sexes game, one of equilibria is clearly 'better' for both players. We might therefore be tempted to regard (A,A) as the more likely equilibrium.

2. However, lots of people typically choose strategy B in most experimental settings. Playing A seems too 'risky' for many players.

3. Harsanyi-Selten developed the notion of risk-dominance. Assume that you don't know much about the other player and assign 50-50 probability to him playing A or B. Then playing A gives you utility -3 in expectation while playing B gives you 7.5. Therefore, B *risk-dominates* A.

4 Interpretations of NE

IDSDS is a constructive algorithm to predict play and does not assume that players know the strategy choices of other players. In contrast, in a Nash equilibrium players have precise beliefs about the play of other players, and these beliefs are self-fulfilling. However, where do these beliefs come from?

There are several interpretations:

1. **Play Prescription:** Some outside party proposes a prescription of how to play the game. This prescription is stable, i.e. no player has an incentive to deviate from if she thinks the other players follow that prescription.

2. **Preplay communication:** There is a preplay phase in which players can communicate and agree on how to play the game. These agreements are self-enforcing.

3. **Rational Introspection:** A NE seems a reasonable way to play a game because my beliefs of what other players do are consistent with them being rational. This is a good explanation for explaining NE in games with a unique NE. However, it is less compelling for games with multiple NE.

4. **Focal Point:** Social norms, or some distinctive characteristic can induce players to prefer certain strategies over others.

5. **Learning** Agents learn other players' strategies by playing the same game many time over.

6. **Evolution:** Agents are programmed to play a certain strategy and are randomly matched against each other. Assume that agents do not play NE initially. Occasionally 'mutations' are born, i.e. players who deviate from the majority play. If this deviation is profitable, these agents will 'multiply' at a faster rate than other agents and eventually take over. Under certain conditions, this system converges to a state where all agents play a Nash equilibrium, and mutating agents cannot benefit from deviation anymore.

Remark 1 *Each of these interpretations makes different assumptions about the knowledge of players. For a play prescription it is sufficient that every player is rational, and simply trusts the outside party. For rational introspection it has to be common knowledge that players are rational. For evolution players do not even have to be rational.*

Remark 2 *Some interpretations have less problems in dealing with multiplicity of equilibria. If we believe that NE arises because an outside party prescribes play for both players, then we don't have to worry about multiplicity - as long as the outside party suggests some Nash equilibrium, players have no reason to deviate. Rational introspection is much more problematic: each player can rationalize any of the multiple equilibria and therefore has no clear way to choose amongst them.*

Lecture V: Mixed Strategies

Markus M. Möbius

February 24, 2004

- Gibbons, sections 1.3-1.3.A
- Osborne, chapter 4

1 The Advantage of Mixed Strategies

Consider the following Rock-Paper-Scissors game: *Note that RPS is a zero-sum game.*

	R	P	S
R	0,0	−1,1	1,−1
P	1,−1	0,0	−1,1
S	−1,1	1,−1	0,0

This game has no pure-strategy Nash equilibrium. Whatever pure strategy player 1 chooses, player 2 can beat him. A natural solution for player 1 might be to randomize amongst his strategies.

Another example of a game without pure-strategy NE is matching pennies. As in RPS the opponent can exploit his knowledge of the other player's action.

	H	T
H	1,−1	−1,1
T	−1,1	1,−1

Fearing this what might the opponent do? One solution is to randomize and play a mixed strategy. Each player could flip a coin and play H with probability $\frac{1}{2}$ and T with probability $\frac{1}{2}$.

Note that each player cannot be taken advantage of.

Definition 1 *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 finite a mixed strategy is a function $\sigma_i : S_i \rightarrow \Re^+$ such that $\sum_{s_i \in S_i} \sigma_i (s_i) = 1$.*

Several notations are commonly used for describing mixed strategies.

1. Function (measure): $\sigma_1 (H) = \frac{1}{2}$ and $\sigma_1 (T) = \frac{1}{2}$
2. Vector: If the pure strategies are s_{i1},..s_{iN_i} write $(\sigma_i (s_{i1}), .., \sigma_i (s_{iN_i}))$ e.g. $\left(\frac{1}{2}, \frac{1}{2}\right)$.
3. $\frac{1}{2}H + \frac{1}{2}T$

Class Experiment 1 *Three groups of two people. Play RPS with each other 30 times. Calculate frequency with which each strategy is being played.*

- Players are indifferent between strategies if opponent mixes equally between all three strategies.
- In games such as matching pennies, poker bluffing, football run/pass etc you want to make the opponent guess and you worry about being found out.

2 Mixed Strategy Nash Equilibrium

Write Σ_i (also $\Delta(S_i)$) for the set of probability distributions on S_i.

Write Σ for $\Sigma_1 \times .. \times \Sigma_I$. A mixed strategy profile $\sigma \in \Sigma$ is an I-tuple $(\sigma_1, .., \sigma_I)$ with $\sigma_i \in \Sigma_i$.

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 play as in σ_{-i}.

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_i, s_{-i} u_i(s_i, s_{-i}) \sigma_i(s_i) \sigma_{-i}(s_i)} \tag{1}$$

Remark 1 *For the definition of a mixed strategy payoff we have to assume that the utility function fulfills the VNM axioms. Mixed strategies induce lotteries over the outcomes (strategy profiles) and the expected utility of a lottery allows a consistent ranking only if the preference relation satisfies these axioms.*

Definition 2 *A mixed strategy NE of G is a mixed profile $\sigma^* \in \Sigma$ such that*

$$u_i\left(\sigma_i^*, \sigma_{-i}^*\right) \geq u_i\left(\sigma_i, \sigma_{-i}^*\right)$$

for all i and all $\sigma_i \in \Sigma_i$.

3 Testing for MSNE

The definition of MSNE makes it cumbersome to check that a mixed profile is a NE. The next result shows that it is sufficient to check against pure strategy alternatives.

Proposition 1 *σ^* is a Nash equilibrium if and only if*

$$u_i\left(\sigma_i^*, \sigma_{-i}^*\right) \geq u_i\left(s_i, \sigma_{-i}^*\right)$$

for all i and $s_i \in S_i$.

Example 1 *The strategy profile $\sigma_1^* = \sigma_2^* = \frac{1}{2}H + \frac{1}{2}T$ is a NE of Matching Pennies.*

Because of symmetry is it sufficient to check that player 1 would not deviate. If he plays his mixed strategy he gets expected payoff 0. Playing his two pure strategies gives him payoff 0 as well. Therefore, there is no incentive to deviate.

Note: Mixed strategies can help us to find MSNE when no pure strategy NE exist.

4 Finding Mixed Strategy Equilibria I

Definition 3 *In a finite game, the support of a mixed strategy σ_i, $supp(\sigma_i)$, is the set of pure strategies to which σ_i assigns positive probability*

$$supp(\sigma_i) = \{s_i \in S_i | \sigma_i(s_i) > 0\}$$

Proposition 2 *If σ^* is a mixed strategy Nash equilibrium and $s_i', s_i'' \in supp(\sigma_i^*)$ then*

$$u_i\left(s_i', \sigma_{-i}^*\right) = u_i\left(s_i'', \sigma_{-i}^*\right)$$

Proof: Suppose not. Assume WLOG that

$$u_i\left(s_i', \sigma_{-i}^*\right) > u_i\left(s_i'', \sigma_{-i}^*\right)$$

with $s_i', s_i'' \in supp(\sigma_i^*)$.

Define a new mixed strategy $\hat{\sigma}_i$ for player i by

$$\hat{\sigma}_i(s_i) = \begin{cases} \sigma_i^*(s_i') + \sigma_i^*(s_i'') & \text{if } s_i = s_i' \\ 0 & \text{if } s_i = s_i'' \\ \sigma_i^*(s_i) & \text{otherwise} \end{cases}$$

We can calculate the gain from playing the modified stratgey:

$$\begin{aligned} u_i\left(\hat{\sigma}_i, \sigma_{-i}^*\right) &- u_i\left(\sigma_i^*, \sigma_{-i}^*\right) \\ &= \sum_{s_i \in S_i} u_i\left(s_i, \sigma_{-i}^*\right) \hat{\sigma}_i(s_i) - \sum_{s_i \in S_i} u_i\left(s_i, \sigma_{-i}^*\right) \sigma_i^*(s_i) \\ &= \sum_{s_i \in S_i} u_i\left(s_i, \sigma_{-i}^*\right) [\hat{\sigma}_i(s_i) - \sigma_i^*(s_i)] \\ &= u_i\left(s_i', \sigma_{-i}^*\right) \sigma_i^*(s_i'') - u_i\left(s_i'', \sigma_{-i}^*\right) \sigma_i^*(s_i'') \\ &> 0 \end{aligned}$$

Note that a mixed strategy NE is never strict.
The proposition suggests a process of finding MSNE.

1. Look at all possible supports for mixed equilibria.

2. Solve for probabilities and check if it works.

Example 2 *Find all the Nash equilibria of the game below.*

	L	R
U	1,1	0,4
D	0,2	2,1

It is easy to see that this game has no pure strategy Nash equilibria. For a mixed strategy Nash equilibrium to exist player 1 has to be indifferent between strategies U and D and player 2 has to be indifferent between L and R. Assume player 1 plays U with probability α and player 2 plays L with probability β.

$$\begin{aligned}
u_1(U, \sigma_2^*) &= u_1(D, \sigma_2^*) \\
\beta &= 2(1-\beta) \\
u_2(L, \sigma_1^*) &= u_2(R, \sigma_1^*) \\
\alpha + 2(1-\alpha) &= 4\alpha + (1-\alpha)
\end{aligned}$$

We can deduce that $\alpha = \frac{1}{4}$ and $\beta = \frac{2}{3}$. There is a unique mixed Nash equilibrium with $\sigma_1^* = \frac{1}{4}U + \frac{3}{4}D$ and $\sigma_2^* = \frac{2}{3}L + \frac{1}{3}R$

Remark 2 *Recall the Battle of the Sexes experiments from last class. It can be shown that the game has a mixed NE where each agent plays her favorite strategy with probability $\frac{2}{3}$. This was not quite the proportion of people playing it in class (but pretty close to the proportion of people choosing it in the previous year)! In many instances of this experiment one finds that men and women differed in their 'aggressiveness'. Does that imply that they were irrational? No. In a mixed NE players are indifferent between their strategies. As long as men and women are matched completely randomly (i.e. woman-woman and man-man pairs are also possible) it only matters how players mix in the aggregate! It does NOT matter if subgroups (i.e. 'men' and 'women') mix at different states, although it would matter if they* *would play only against players within their subgroup. Interestingly, that suggests that letting women 'segregate' into their own communities should make them more aggressive, and men less aggressive. The term 'aggressive' is a bit misleading because it does not result in bigger payoffs. However, you could come up with a story of lexicographic preferences - people care first of all about payoffs, but everything else equal they want to fit gender stereotypes - so playing 'football' is good for men's ego.*

5 Finding Mixed Strategy Equilibria II

Finding mixed NE in 2 by 2 games is relatively easy. It becomes harder if players have more than two strategies because we have to start worrying about supports. In many cases it is useful to exploit iterated deletion in order to narrow down possible supports.

Proposition 3 *Let σ^* be a NE of G and suppose that $\sigma^*(s_i) > 0$ then $s_i \in S_i^\infty$, i.e. strategy s_i is not removed by ISD.*

Proof: see problem set 2

Having introduced mixed strategies we can even define a tighter notion of IDSDS. Consider the next game. No player has a strategy which is strictly dominated by another pure strategy. However, C for player 2 is strictly dominated by $\frac{1}{2}L + \frac{1}{2}R$. Thus we would think that C won't be used.

	L	C	R
U	1,1	0,2	0,4
M	0,2	5,0	1,6
D	0,2	1,1	2,1

After we delete C we note that M is dominated by $\frac{2}{3}D+\frac{1}{3}U$. Using the above proposition we can conclude that the only Nash equilibria are the NE of the 2 by 2 game analyzed in the previous section. Since that game had a unique mixed strategy equilibrium we can conclude that the only NE of the 3 by 3 game is $\sigma_1^* = \frac{1}{4}U + \frac{3}{4}D$ and $\sigma_2^* = \frac{2}{3}L + \frac{1}{3}R$.

It is useful to adjust the formal definition of IDSDS and allow for mixed strategy domination:

Definition 4 *The set of strategy profiles surviving iterated strict dominance is $S^\infty = S_1^\infty \times S_2^\infty \times .. \times S_I^\infty$ where*

$$\begin{aligned} S_i^\infty &= \bigcap_{k=1}^{\infty} S_i^k \\ S_i^0 &= S_i \\ S_i^{k+1} &= \left\{ s_i \in S_i^k | \ \nexists \sigma_i \in \Delta\left(S_i^k\right) | u_i\left(\sigma_i, s_{-i}\right) > u_i\left(s_i, s_{-i}\right) \forall s_{-i} \in S_{-i}^k \right\} \end{aligned}$$

Remark 3 *Recall the above 3 by 3 game. If we would look for possible mixed NE with supports (U,M) and (L,C) respectively, we would find a potential NE $\frac{2}{3}U + \frac{1}{3}M, \frac{5}{6}L + \frac{1}{6}C$. However, this is NOT a NE because player 2 would play R instead.*

Remark 4 *In the RPS game we cannot reduce the set of strategies through IDSDS. Therefore we have to check all possible supports and check if it works.*

6 Finding Mixed Strategy Equilibria III

Definition 5 *A correspondence $F : A \to B$ is a mapping which associates to every element of $a \in A$ a subset $F(a) \subset B$.*

The mixed strategy best response correspondence for player i $BR_i : \Sigma_{-i} \to \Sigma_i$ is defined by

$$BR_i(\sigma_{-i}) = \arg\max_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \sigma_{-i})$$

Proposition 4 *σ^* is a Nash equilibrium if and only if $\sigma_i^* \in BR_i\left(\sigma_{-i}^*\right)$ for all i.*

In the 2 by 2 game we have:

$$BR_1(\beta L + (1-\beta) R) = \begin{cases} U & \text{if } \beta > \frac{2}{3} \\ \{\alpha U + (1-\alpha) D | \alpha \in [0,1]\} & \text{if } \beta = \frac{2}{3} \\ D & \text{if } \beta < \frac{2}{3} \end{cases}$$

$$BR_2(\alpha U + (1-\alpha) D) = \begin{cases} L & \text{if } \alpha < \frac{1}{4} \\ \{\beta L + (1-\beta) R | \beta \in [0,1]\} & \text{if } \alpha = \frac{1}{4} \\ R & \text{if } \alpha > \frac{1}{4} \end{cases}$$

We can graph both correspondences to find the set of Nash equilibria.

7 Interpretation of Mixed NE

1. Sometimes players explicitly flip coins. That fits games like Poker, soccer etc., where players have to randomize credibly.

2. Large populations of players with each player playing a fixed strategy and random matching. That's very related to the social norm explanation of pure Nash equilibrium.

3. Payoff uncertainty (Harsanyi, purification). Roughly their argument goes as follows in the matching penny game. There are two types of players - those who get slightly higher payoff from playing heads, and those who get higher payoff from getting tails (their preferences are almost the same - think of one guy getting 1 dollar and the other guy getting $1 + \epsilon = 1.01$ dollars from playing head). Also, we assume that there is an equal probability that my opponent is of type 1 or type 2. In this circumstances no player loses from just playing her favorite strategy (i.e. a pure strategy) because it will do best on average. To show that this is a NE we have to introduce the notion of incomplete information which we don't do just yet. Harsanyi-Selten then let the payoff uncertainty ϵ go to 0 such that the game in the limit approaches the standard matching pennies. Players' 'average' strategies converge to the mixed equilibrium.

Lecture VI: Existence of Nash equilibrium

Markus M. Möbius

February 26, 2004

- Gibbons, sections 1.3B
- Osborne, chapter 4

1 Nash's Existence Theorem

When we introduced the notion of Nash equilibrium the idea was to come up with a solution concept which is stronger than IDSDS. Today we show that NE is not too strong in the sense that it guarantees the existence of at least one mixed Nash equilibrium in most games (for sure in all finite games). This is reassuring because it tells that there is at least one way to play most games.[1]

Let's start by stating the main theorem we will prove:

Theorem 1 ***(Nash Existence)*** *Every finite strategic-form game has a mixed-strategy Nash equilibrium.*

Many game theorists therefore regard the set of NE for this reason as the lower bound for the set of reasonably solution concept. A lot of research has gone into refining the notion of NE in order to retain the existence result but get more precise predictions in games with multiple equilibria (such as coordination games).

However, we have already discussed games which are solvable by IDSDS and hence have a unique Nash equilibrium as well (for example, the two thirds of the average game), but subjects in an experiment will not follow those equilibrium prescription. Therefore, if we want to *describe* and *predict*

[1]Note, that a pure Nash equilibrium is a (degenerate) mixed equilibrium, too.

the behavior of real-world people rather than come up with an explanation of how they *should* play a game, then the notion of NE and even even IDSDS can be too restricting.

Behavioral game theory has tried to weaken the joint assumptions of rationality and common knowledge in order to come up with better theories of how real people play real games. Anyone interested should take David Laibson's course next year.

Despite these reservation about Nash equilibrium it is still a very useful benchmark and a starting point for any game analysis.

In the following we will go through three proofs of the Existence Theorem using various levels of mathematical sophistication:

- existence in 2×2 games using elementary techniques
- existence in 2×2 games using a fixed point approach
- general existence theorem in finite games

You are only required to understand the simplest approach. The rest is for the intellectually curious.

2 Nash Existence in 2×2 Games

Let us consider the simple 2×2 game which we discussed in the previous lecture on mixed Nash equilibria:

	L	R
U	1,1	0,4
D	0,2	2,1

We next draw the best-response curves of both players. Recall that player 1's strategy can be represented by a single number α such that $\sigma_1 = \alpha U + (1 - \alpha)D$ while player 2's strategy is $\sigma_2 = \beta L + (1 - \beta)R$.

Let's find the best-response of player 2 to player 1 playing strategy α:

$$\begin{aligned} u_2(L, \alpha U + (1-\alpha)D) &= 2 - \alpha \\ u_2(R, \alpha U + (1-\alpha)D) &= 1 + 3\alpha \end{aligned} \tag{1}$$

Therefore, player 2 will strictly prefer strategy L iff $2 - \alpha > 1 + 3\alpha$ which implies $\alpha < \frac{1}{4}$. The best-response correspondence of player 2 is therefore:

$$BR_2(\alpha) = \begin{cases} 1 & \text{if } \alpha < \frac{1}{4} \\ [0,1] & \text{if } \alpha = \frac{1}{4} \\ 0 & \text{if } \alpha > \frac{1}{4} \end{cases} \tag{2}$$

We can similarly find the best-response correspondence of player 1:

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

We draw both best-response correspondences in a single graph (the graph is in color - so looking at it on the computer screen might help you):

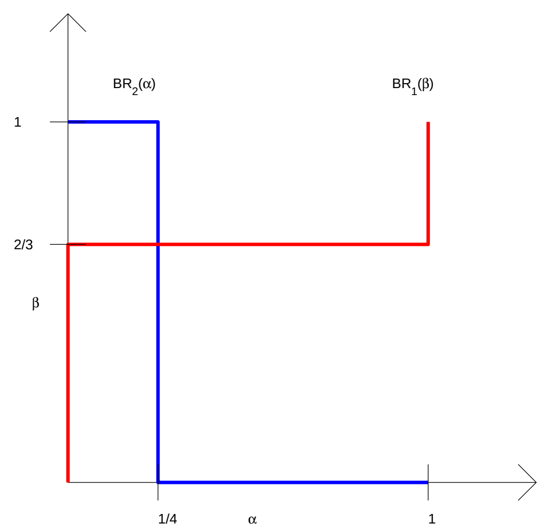

We immediately see, that both correspondences intersect in the single point $\alpha = \frac{1}{4}$ and $\beta = \frac{2}{3}$ which is therefore the unique (mixed) Nash equilibrium of the game.

What's useful about this approach is that it generalizes to a proof that any two by two game has at least one Nash equilibriu, i.e. its two best response correspondences have to intersect in at least one point.

An informal argument runs as follows:

1. The best response correspondence for player 2 maps each α into at least one β. The graph of the correspondence connects the left and right side of the square $[0,1] \times [0,1]$. This connection is continuous - the only discontinuity could happen when player 2's best response switches from L to R or vice versa at some α^*. But at this switching point player 2 has to be exactly indifferent between both strategies - hence the graph has the value $BR_2(\alpha^*) = [0,1]$ at this point and there cannot be a discontinuity. Note, that this is precisely why we need mixed strategies - with pure strategies the BR graph would generally be discontinuous at some point.

2. By an analogous argument the BR graph of player 1 connects the upper and lower side of the square $[0,1] \times [0,1]$.

3. Two lines which connect the left/right side and the upper/lower side of the square respectively have to intersect in at least one point. Hence each 2 by 2 game has a mixed Nash equilibrium.

3 Nash Existence in 2×2 Games using Fixed Point Argument

There is a different way to prove existence of NE on 2×2 games. The advantage of this new approach is that it generalizes easily to general finite games.

Consider any strategy profile $(\alpha U + (1-\alpha)D, \beta L + (1-\beta)R)$ represented by the point (α, β) inside the square $[0,1] \times [0,1]$. Now imagine the following: player 1 assumes that player 2 follows strategy β and player 2 assumes that player 1 follows strategy α. What should they do? They should play their BR to their beliefs - i.e. player 1 should play $BR_1(\beta)$ and player 2 should play $BR_2(\alpha)$. So we can imagine that the strategy profile (α, β) is mapped onto $(BR_1(\beta), BR_2(\alpha))$. This would describe the actual play of both players if their beliefs would be summarizes by (α, β). We can therefore define a giant correspondence $BR : [0,1] \times [0,1] \to [0,1] \times [0,1]$ in the following way:

$$BR(\alpha, \beta) = BR_1(\beta) \times BR_2(\alpha) \tag{4}$$

The following figure illustrates the properties of the combined best-response map BR:

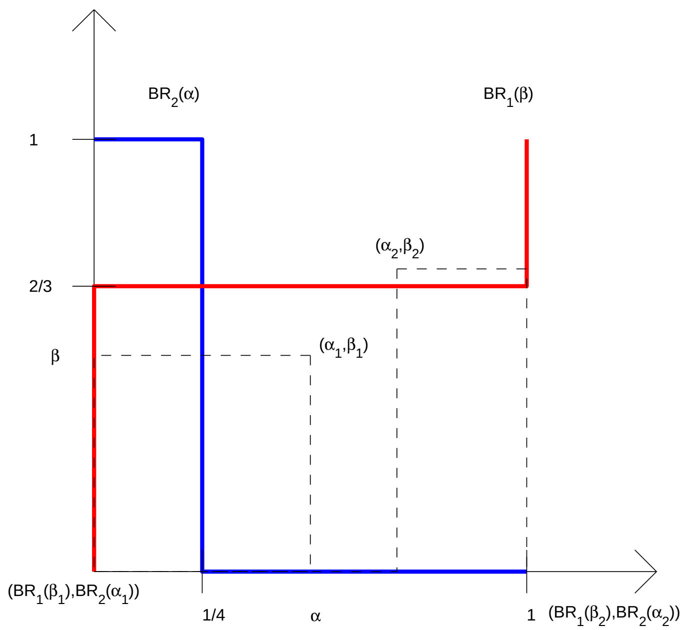

The neat fact about BR is that the Nash equilibria σ^* are precisely the fixed points of BR, i.e. $\sigma^* \in BR(\sigma^*)$. In other words, if players have beliefs σ^* then σ^* should also be a best response by them. The next lemma follows directly from the definition of mixed Nash equilibrium:

Lemma 1 *A mixed strategy profile σ^* is a Nash equilibrium if and only if it is a fixed point of the BR correspondence, i.e. $\sigma^* \in BR(\sigma^*)$.*

We therefore look precisely for the fixed points of the correspondence BR which maps the square $[0,1] \times [0,1]$ onto itself. There is well developed mathematical theory for these types of maps which we utilize to prove Nash existence (i.e. that BR has at least one fixed point).

3.1 Kakutani's Fixed Point Theorem

The key result we need is Kakutani's fixed point theorem. You might have used Brower's fixed point theorem in some mathematics class. This is not sufficient for proving the existence of nash equilibria because it only applies to functions but not to correspondences.

Theorem 2 ***Kakutani*** *A correspondence* $r : X \rightarrow X$ *has a fixed point* $x \in X$ *such that* $x \in r(x)$ *if*

1. X *is a compact, convex and non-empty subset of* $\Re^n$.

2. $r(x)$ *is non-empty for all* x.

3. $r(x)$ *is convex for all* x.

4. r *has a closed graph.*

There are a few concepts in this definition which have to be defined:

Convex Set: A set $A \subseteq \Re^n$ is convex if for any two points $x, y \in A$ the straight line connecting these two points lies inside the set as well. Formally, $\lambda x + (1 - \lambda) y \in A$ for all $\lambda \in [0, 1]$.

Closed Set: A set $A \subseteq \Re^n$ is closed if for any converging sequence $\{x_n\}_{n=1}^{\infty}$ with $x_n \rightarrow x^*$ as $n \rightarrow \infty$ we have $x^* \in A$. Closed intervals such as $[0, 1]$ are closed sets but open or half-open intervals are not. For example $(0, 1]$ cannot be closed because the sequence $\frac{1}{n}$ converges to 0 which is not in the set.

Compact Set: A set $A \subseteq \Re^n$ is compact if it is both closed and bounded. For example, the set $[0, 1]$ is compact but the set $[0, \infty)$ is only closed but unbounded, and hence not compact.

Graph: The graph of a correspondence $r : X \rightarrow Y$ is the set $\{(x, y) \, | y \in r(x)\}$. If r is a real function the graph is simply the plot of the function.

Closed Graph: A correspondence has a closed graph if the graph of the correspondence is a closed set. Formally, this implies that for a sequence of point on the graph $\{(x_n, y_n)\}_{n=1}^{\infty}$ such that $x_n \rightarrow x^*$ and $y_n \rightarrow y^*$ as $n \rightarrow \infty$ we have $y^* \in r(x^*)$.[2]

It is useful to understand exactly why we need each of the conditions in Kakutani's fixed point theorem to be fulfilled. We discuss the conditions by looking correspondences on the real line, i.e. $r : \Re \rightarrow \Re$. In this case, a fixed point simply lies on the intersection between the graph of the correspondence and the diagonal $y = x$. Hence Kakutani's fixed point theorem tells us that

[2]If the correspondence is a function then the closed graph requirement is equivalent to assuming that the function is continuous. It's easy to see that a continuous function has a closed graph. For the reverse, you'll need Baire's category theorem.

a correspondence $r : [0,1] \to [0,1]$ which fulfills the conditions above always intersects with the diagonal.

3.1.1 Kakutani Condition I: X is compact, convex and non-empty.

Assume X is not compact because it is not closed - for example $X = (0,1)$. Now consider the correspondence $r(x) = x^2$ which maps X into X. However, it has no fixed point. Now consider X non-compact because it is unbounded such as $X = [0,\infty)$ and consider the correspondence $r(x) = 1 + x$ which maps X into X but has again no fixed point.

If X is empty there is clearly no fixed point. For convexity of X look at the example $X = [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$ which is not convex because the set has a hole. Now consider the following correspondence (see figure below):

$$r(x) = \begin{cases} \frac{3}{4} & \text{if } x \in [0, \frac{1}{3}] \\ \frac{1}{4} & \text{if } x \in [\frac{2}{3}, 1] \end{cases} \tag{5}$$

This correspondence maps X into X but has no fixed point again.

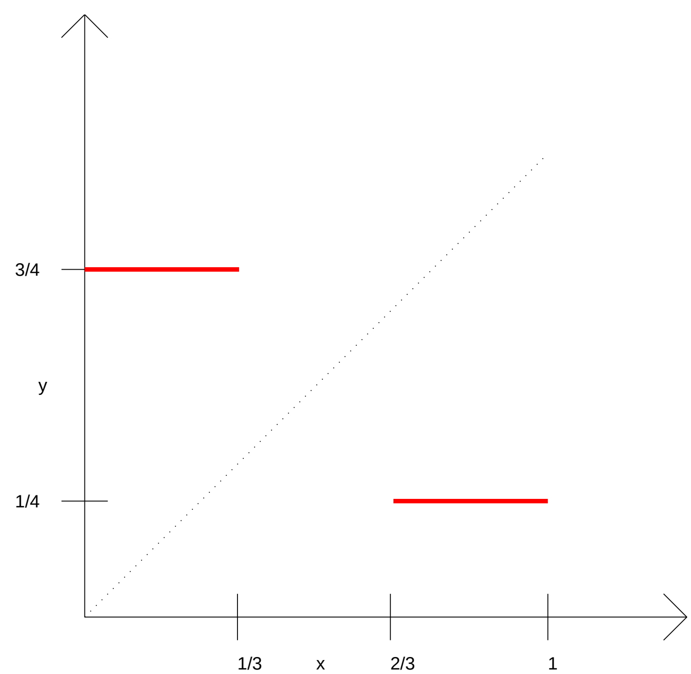

From now on we focus on correspondences $r : [0,1] \to [0,1]$ - note that $[0,1]$ is closed and bounded and hence compact, and is also convex.

3.1.2 Kakutani Condition II: $r(x)$ is non-empty.

If $r(x)$ could be empty we could define a correspondence $r : [0,1] \longrightarrow [0,1]$ such as the following:

$$r(x) = \begin{cases} \frac{3}{4} & \text{if } x \in [0, \frac{1}{3}] \\ \emptyset & \text{if } x \in [\frac{1}{3}, \frac{2}{3}] \\ \frac{1}{4} & \text{if } x \in [\frac{2}{3}, 1] \end{cases} \tag{6}$$

As before, this correspondence has no fixed point because of the hole in the middle.

3.1.3 Kakutani Condition III: $r(x)$ is convex.

If $r(x)$ is not convex, then the graph does not have to have a fixed point as the following example of a correspondence $r : [0,1] \longrightarrow [0,1]$ shows:

$$r(x) = \begin{cases} 1 & \text{if } x < \frac{1}{2} \\ \left[0, \frac{1}{3}\right] \cup \left[\frac{2}{3}, 1\right] & \text{if } x = \frac{1}{2} \\ 0 & \text{if } x > \frac{1}{2} \end{cases} \tag{7}$$

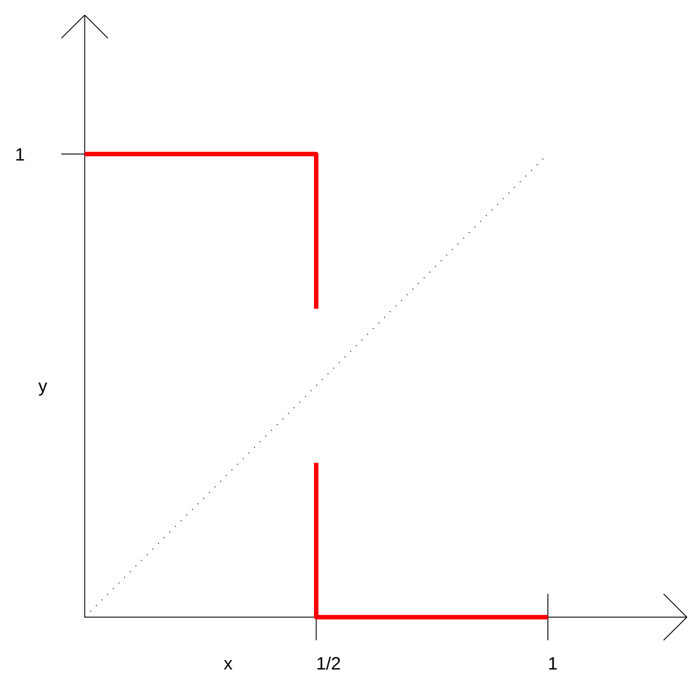

The graph is non-convex because $r(\frac{1}{2})$ is not convex. It also does not have a fixed point.

3.1.4 Kakutani Condition IV: $r(x)$ has a closed graph.

This condition ensures that the graph cannot have holes. Consider the following correspondence $r : [0,1] \to [0,1]$ which fulfills all conditions of Kakutani except (4):

$$r(x) = \begin{cases} \frac{1}{2} & \text{if } x < \frac{1}{2} \\ \left[\frac{1}{4}, \frac{1}{2}\right) & \text{if } x = \frac{1}{2} \\ \frac{1}{4} & \text{if } x > \frac{1}{2} \end{cases} \tag{8}$$

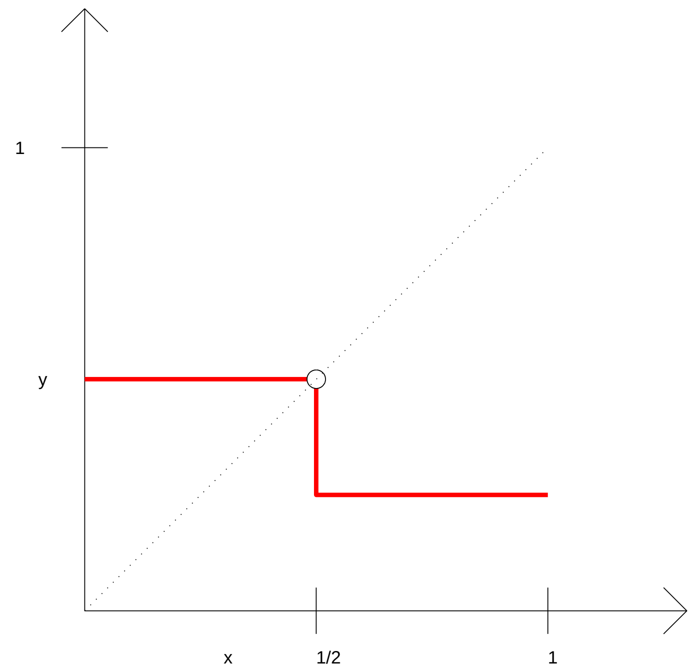

Note, that $r(\frac{1}{2})$ is the convex set $\left[\frac{1}{4}, \frac{1}{2}\right)$ but that this set is not closed. Hence the graph is not closed. For example, consider the sequence $x_n = \frac{1}{2}$ and $y_n = \frac{1}{2} - \frac{1}{n+2}$ for $n \geq 1$. Clearly, we have $y_n \in r(x_n)$. However, $x_n \to x^* = \frac{1}{2}$ and $y_n \to y^* = \frac{1}{2}$ but $y^* \notin r(x^*)$. Hence the graph is not closed.

3.2 Applying Kakutani

We now apply Kakutani to prove that 2×2 games have a Nash equilibrium, i.e. the giant best-response correspondence BR has a fixed point. We denote the strategies of player 1 with U and D and the strategies of player 2 with L and R.

We have to check (a) that BR is a map from some compact and convex set X into itself, and (b) conditions (1) to (4) of Kakutani.

- First note, that $BR : [0,1] \times [0,1] \to [0,1] \times [0,1]$. The square $X = [0,1] \times [0,1]$ is convex and compact because it is bounded and closed.

- Now check condition (2) of Kakutani - $BR(\sigma)$ is non-empty. This is true if $BR_1(\sigma_2)$ and $BR_2(\sigma_1)$ are non-empty. Let's prove it for BR_1 - the proof for BR_2 is analogous. Player 1 will get the following payoff $u_{1,\beta}(\alpha)$ from playing strategy α if the other player plays β:

$$\begin{aligned} u_{1,\beta}(\alpha) &= \alpha\beta u_1(U,L) + \alpha(1-\beta)u_1(U,R) + \\ &+ (1-\alpha)\beta u_1(D,L) + (1-\alpha)(1-\beta)u_1(D,R) \end{aligned} \quad (9)$$

  The function $u_{1,\beta}$ is continuous in α. We also know that $\alpha \in [0,1]$ which is a closed interval. Therefore, we know that the continuous function $u_{1,\beta}$ reaches its maximum over that interval (standard min-max result from real analysis - continuous functions reach their minimum and maximum over closed intervals). Hence there is at least one best response α^* which maximizes player 1's payoff.

- Condition (3) requires that if player 1 has tow best responses $\alpha_1^* U + (1-\alpha_1^*)D$ and $\alpha_2^* U + (1-\alpha_2^*)D$ to player 2 playing $\beta L + (1-\beta)R$ then the strategy where player 1 chooses U with probability $\lambda\alpha_1^* + (1-\lambda)\alpha_2^*$ for some $0 < \lambda < 1$ is also a best response (i.e. $BR_1(\beta)$ is convex). But since both the α_1 and the α_2 strategy are best responses of player 1 to the same β strategy of player 2 they also have to provide the same payoffs to player 1. But this implies that if player 1 plays strategy α_1 with probability λ and α_2 with probability $1-\lambda$ she will get exactly the same payoff as well. Hence the strategy where she plays U with probability $\lambda\alpha_1^* + (1-\lambda)\alpha_2^*$ is also a best response and her best response set $BR_1(\beta)$ is convex.

- The final condition (4) requires that BR has a closed graph. To show this consider a sequence $\sigma^n = (\alpha^n, \beta^n)$ of (mixed) strategy profiles and $\tilde{\sigma}^n = (\tilde{\alpha}^n, \tilde{\beta}^n) \in BR(\sigma^n)$. Both sequences are assumed to converge to $\sigma^* = (\alpha^*, \beta^*)$ and $\tilde{\sigma}^* = (\tilde{\alpha}^*, \tilde{\beta}^*)$, respectively. We now want to show that $\tilde{\sigma} \in BR(\sigma)$ to prove that BR has a closed graph.

  We know that for player 1, for example, we have

$$u_1(\tilde{\alpha}^n, \beta^n) \geq u_1(\alpha', \beta^n)$$

for any $\alpha' \in [0,1]$. Note, that the utility function is continuous in both arguments because it is linear in α and β. Therefore, we can take the limit on both sides while preserving the inequality sign:

$$u_1(\tilde{\alpha}^*, \beta^*) \geq u_2(\alpha', \beta)$$

for all $\alpha' \in [0,1]$. This shows that $\tilde{\alpha^*} \in BR_1(\beta)$ and therefore $\tilde{\sigma}^* \in BR(\sigma^*)$. Hence the graph of the BR correspondence is closed.

Therefore, all four Kakutani conditions apply and the giant best-response correspondence BR has a fixed point, and each 2×2 game has a Nash equilibrium.

4 Nash Existence Proof for General Finite Case

Using the fixed point method it is now relatively easy to extend the proof for the 2×2 case to general finite games.The biggest difference is that we cannot represent a mixed strategy any longer with a single number such as α. If player 1 has three pure strategies A_1,A_2 and A_3, for example, then his set of mixed strategies is represented by two probabilities - for example, (α_1, α_2) which are the probabilities that A_1 and A_2 are chosen. The set of admissible α_1 and α_2 is described by:

$$\Sigma_1 = \{(\alpha_1, \alpha_2) | 0 \leq \alpha_1, \alpha_2 \leq 1 \quad \text{and} \quad \alpha_1 + \alpha_2 \leq 1\} \tag{10}$$

The definition of the set of mixed strategies can be straightforwardly extended to games where player 1 has a strategy set consisting of n pure strategies A_1,..,A_n. Then we need $n-1$ probabilities α_1,..,α_{n-1} such that:

$$\Sigma_1 = \{(\alpha_1, .., \alpha_{n-1}) | 0 \leq \alpha_1, .., \alpha_{n-1} \leq 1 \quad \text{and} \quad \alpha_1 + .. + \alpha_{n-1} \leq 1\} \tag{11}$$

So instead of representing strategies on the unit interval $[0,1]$ we have to represent as elements of the simplex Σ_1.

Lemma 2 *The set Σ_1 is compact and convex.*

Proof: It is clearly convex - if you mix between two mixed strategies you get another mixed strategy. The set is also compact because it is bounded (all $|\alpha_i| \leq 1$) and closed. To see closedness take a sequence $(\alpha_1^i, .., \alpha_{n-1}^i)$ of elements of Σ_1 which converges to $(\alpha_1^*, ..)$. Then we have $\alpha_i^* \geq 0$ and $\sum_{i=1}^{n-1} \alpha_i^* \leq 1$ because the limit preserves weak inequalities. QED

We can now check that all conditions of Kakutani are fulfilled in the general finite case. Checking them is almost 1-1 identical to checking Kakutani's condition for 2×2 games.

Condition 1: The individual mixed strategy sets Σ_i are clearly non-empty because every player has at least one strategy. Since Σ_i is compact $\Sigma = \Sigma_1 \times ... \times \Sigma_I$ is also compact. Hence the BR correspondence $BR : \Sigma \to \Sigma$ acts on a compact and convex non-empty set.

Condition 2: For each player i we can calculate his utiltiy $u_{i,\sigma_{-i}}(\sigma_i)$ for $\sigma_i \in \Sigma_i$. Since Σ_i is compact and $u_{i,\sigma_{-i}}$ is continuous the set of payoffs is also compact and hence has a maximum. Therefore, $BR_i(\Sigma_i)$ is non-empty.

Condition 3: Assume that σ_i^1 and σ_i^2 are both BR of player i to σ_{-i}. Both strategies have to give player i equal payoffs then and any linear combination of these two strategies has to be a BR for player i, too.

Condition 4: Assume that σ^n is a sequence of strategy profiles and $\tilde{\sigma}^n \in BR(\sigma^n)$. Both sequences converge to σ^* and $\tilde{\sigma}^*$, respectively. We know that for each player i we have

$$u_i\left(\tilde{\sigma}_i^n, \sigma_{-i}^n\right) \geq u_i\left(\sigma_i', \sigma_{-i}^n\right)$$

for all $\sigma_i' \in \Sigma_i$. Note, that the utility function is continuous in both arguments because it is linear.[3] Therefore, we can take the limit on both sides while preserving the inequality sign:

$$u_i\left(\tilde{\sigma}_i^*, \sigma_{-i}^*\right) \geq u_i\left(\sigma_i', \sigma_{-i}^*\right)$$

for all $\sigma_i' \in \Sigma_i$. This shows that $\tilde{\sigma}_i^* \in BR_i\left(\sigma_{-i}^*\right)$ and therefore $\tilde{\sigma}^* \in BR(\sigma^*)$. Hence the graph of the BR correspondence is closed.

So Kakutani's theorem applies and the giant best-response map BR has a fixed point.

[3] It is crucial here that the set of pure strategies is finite.

Lecture VII: Common Knowledge

Markus M. Möbius

March 4, 2004

This is the one of the two advanced topics (the other is learning) which is not discussed in the two main texts. I tried to make the lecture notes self-contained.

- Osborne and Rubinstein, sections 5.1,5.2,5.4

Today we formally introduce the notion of common knowledge and discuss the assumptions underlying players' knowledge in the two solution concepts we discussed so far - IDSDS and Nash equilibrium.

1 A Model of Knowledge

There is a set of states of nature $\Omega = \{\omega_1, \omega_2, .., \omega_n\}$ which represent the uncertainty which an agent faces when making a decision.

Example 1 *Agents* $1, 2$ *have a prior over the states of nature*

$$\begin{aligned} \Omega &= \{\omega_1 = \textit{It will rain today}, \omega_2 = \textit{It will be cloudy today}, \\ \omega_3 &= \textit{It will be sunny today}\,\} \end{aligned}$$

where each of the three events is equally likely ex ante.

The knowledge of every agent i is represented by an information partition H_i of the set Ω.

Definition 1 *An information partition* H_i *is a collection* $\{h_i(\omega) \mid \omega \in \Omega\}$ *of disjoint subsets of* Ω *such that*

- *(P1)* $\omega \in h_i(\omega)$,
- *(P2) If* $\omega' \in h_i(\omega)$ *then* $h_i(\omega') = h_i(\omega)$.

Note, that the subsets $h_i(\omega)$ span Ω. We can think of $h_i(\omega)$ as the knowledge of agent i if the state of nature is in fact ω. Property P1 ensures that the true state of nature ω is an element of an agent's information set (or knowledge) - this is called the axiom of knowledge. Property P2 is a consistency criterion. Assume for example, that $\omega' \in h_i(\omega)$ and that there is a state $\omega'' \in h_i(\omega')$ but $\omega'' \notin h_i(\omega)$. Then in the state of nature is ω the decision-maker could argue that because ω'' is inconsistent with his information the true state can not be ω'.

Example 1 *(cont.) Agent* 1 *has the information partition*

$$H_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$$

So the agent has good information if the weather is going to be sunny but cannot distinguish between bad weather.

We next define a knowledge function K.

Definition 2 *For any event* E *(a subset of* Ω*) we have*

$$K(E) = \{\omega \in \Omega | h_i(\omega) \subseteq E\}.$$

So the set $K(E)$ is the collection of all states in which the decision maker knows E.

We are now ready to define common knowledge (for simplicity we only consider two players).

Definition 3 *Let* K_1 *and* K_2 *be the knowledge functions of both players. An event* $E \subseteq \Omega$ *is common knowledge between 1 and 2 in the state* $\omega \in \Omega$ *if* ω *if a member of every set in the infinite sequence* $K_1(E)$, $K_2(E)$, $K_1(K_2(E))$, $K_2(K_1(E))$ *and so on.*

This definition implies that player 1 and 2 knows E, they know that the other player knows it, and so on.

There is an equivalent definition of common knowledge which is frequently easier to work with.

Definition 4 *An event $F \subseteq \Omega$ is self-evident between both players if for all $\omega \in F$ we have $h_i(\omega) \subseteq F$ for $i = 1, 2$. An event $E \subseteq \Omega$ is common knowledge between both players in the state $\omega \in \Omega$ if there is a self-evident event F for which $\omega \in F \subseteq E$.*

Example 1 *(cont.) Agent 2 has information function*

$$H_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}\}$$

In this case the event $E = \{\omega_1, \omega_2\}$ is common knowledge if the state of nature is ω_1 or ω_2. Both definition can be applied - E survives iterated deletion, but is also self-evident.

We finally show that both definitions of common knowledge are equivalent. We need the next proposition first.

Proposition 1 *The following are equivalent:*

1. *$K_i(E) = E$ for $i = 1, 2$*
2. *E is self evident between 1 and 2.*
3. *E is a union of members of the partition induced by H_i for $i = 1, 2$.*

Proof: Assume a). Then for every $\omega \in E$ we have $h_i(\omega) \subseteq E$ and b) follows. c) follows because immediately. c) implies a).

We can now prove the following theorem.

Theorem 1 *Definitions 3 and 4 are equivalent.*

Proof: Assume that the event E is common knowledge in state ω according to definition 3. First note, that

$$E \supseteq K_i(E) \supseteq K_j(K_i(E)) \ldots$$

Because Ω is finite and ω is a member of those subsets the infinite regression must eventually produce a set F such that $K_i(F) = F$. Therefore, F is self-evident and we are done.

Next assume that the event E is common knowledge in state ω according to definition 4. Then $F \subseteq E$, $K_i(F) = F \subseteq K_i(E)$ etc, and F is a member of every of the regressive subsets $K_i(K_j(\ldots E \ldots))$ and so is ω. This proves the theorem.

2 Dirty Faces

We have played the game in class. Now we are going to analyze it by using the mathematical language we just developed. Recall, that any agent can only see the faces of all $n-1$ other agents but not her own. Furthermore, at least one face is dirty.

First of all we define the states of nature Ω. If there are n players there are $2^n - 1$ possible states (since all faces clean cannot be a state of nature). It's convenient to denote the states by the n-tuples $\omega = (C, C, D, .., C)$. We also denote the number of dirty faces with $|\omega|$ and note that $|\omega| \geq 1$ by assumption (there is at least one dirty face).

The initial information set of each agent in some state of nature ω has at most two elements. The agent knows the faces of all other agents $\omega(-i)$ but does not know if she has a clean or dirty face, i.e. $h_i(\omega) = \{(C, \omega(-i)), (D, \omega(-i))\}$. Initially, all agents information set has two elements except in the case $|\omega| = 1$ and one agent sees only clean faces - then she know for sure that she has a dirty face because all clean is excluded. You can easily show that the event "There is at least one dirty face" is common knowledge as you would expect.

The game ends when at least player knows the state of the world for sure, i.e. her knowledge partition $h_i(\omega)$ consists of a single element. In the first this will only be the case if the state of world is such that only a single player has a dirt face.

What happens if no agent raises her hand in the first period? All agents update their information partition and exclude all states of nature with just one dirty face. All agents who see just one dirty face now know for sure the state of nature (they have a dirty face!). They raise their hand and the game is over. Otherwise, all agents can exclude states of nature with at most two dirty faces. Agents who see two dirty faces now know the state of nature for sure (they have a dirty face!). etc.

The state of nature with k dirty faces therefore gets revealed at stage k of the game. At that point all guys with dirty faces know the state of nature for sure.

Question: What happens if it is common knowledge that neither all faces are clean, nor all faces are dirty?

Remark 1 *The game crucially depends on the fact that it is common knowledge that at least one agent has a dirty face. Assume no such information would be known - so the state of the world where all faces are clean would* *be a possible outcome. Then no agent in the first round would ever know for sure if she had a dirty face. Hence the information partition would never get refined after any number of rounds.*

3 Coordinated Attack

This story shows that 'almost common knowledge' can be very different from common knowledge.

Two divisions of an army are camped on two hilltops overlooking a common valley. In the valley waits the enemy. If both divisions attack simultaneously they will win the battle, whereas if only one division attacks it will be defeated. Neither general will attack unless he is sure that other will attack with him.

Commander A is in peace negotiations with the enemy. The generals agreed that if the negotiations fail, commander A will send a message to commander B with an attack plan. However, there is a small probability ϵ that the messenger gets intercepted and the message does not arrive. The messenger takes one hour normally. How long will it take to coordinate on the attack?

The answer is: never! Once commander B receives the message he has to confirm it - otherwise A is not sure that he received it and will not attack. But B cannot be sure that A receives his confirmation and will not attack until he receives another confirmation from A. etc. The messenger can run back and forth countless times before he is intercepted but the generals can never coordinate with certainty.

Let's define the state of nature to be (n, m) if commander A sent n messages and received $n-1$ confirmation from commander B, and commander B sent m messages and received $m-1$ confirmations. We also introduce the state of nature $(0, 0)$. In that state the peace negotiations have succeeded, and no attack is scheduled.[1]

The information partition of commander A is

$$H_A = \{\{(0,0)\}, \{(1,0),(1,1)\}, \{(2,1),(2,2)\}, ..\}. \quad (1)$$

[1] As was pointed out by an alert student in class this state is necessary to make this exercise interesting. Otherwise, the generals could agree on an attack plan in advance, and no communication would be necessary at all - the attack would be common knowledge already.

The information partition of commander B is

$$H_B = \{\{(0,0),(1,0)\},\{(1,1),(2,1)\},..\}. \tag{2}$$

Both commanders only attack in some state of the world ω if it is common knowledge that commander B has sent a message, i.e. $n \geq 1$ (the negotiations have failed and an attack should occur). However, this event can never be common knowledge for any state of nature (i.e. after any sequence of messages) because there is no self-evident set F contained in the event E. This is easy to verify: take the union of any collection of information sets of commander A (only those can be candidates for a self-evident F). Then ask yourself whether such a set can be also the union of a collection of information sets of commander B. The answer is no - there will always some information set of B which 'stick out' at either 'end' of the candidate set F.

Lecture VIII: Learning

Markus M. Möbius

March 10, 2004

Learning and evolution are the second set of topics which are not discussed in the two main texts. I tried to make the lecture notes self-contained.

- Fudenberg and Levine (1998), The Theory of Learning in Games, Chapter 1 and 2

1 Introduction

What are the problems with Nash equilibrium? It has been argued that Nash equilibrium are a reasonable minimum requirement for how people should play games (although this is debatable as some of our experiments have shown). It has been suggested that players should be able to figure out Nash equilibria starting from the assumption that the rules of the game, the players' rationality and the payoff functions are all common knowledge.[1] As Fudenberg and Levine (1998) have pointed out, there are some important conceptual and empirical problems associated with this line of reasoning:

1. If there are multiple equilibria it is not clear how agents can coordinate their beliefs about each other's play by pure introspection.

2. Common knowledge of rationality and about the game itself can be difficult to establish.

[1] We haven't discussed the connection between knowledge and Nash equilibrium. Assume that there is a Nash equilibrium σ^* in a two player game and that each player's best-response is unique. In this case player 1 knows that player 2 will play σ_2^* in response to σ_1^*, player 2 knows that player 1 knows this etc. Common knowledge is important for the same reason that it matters in the coordinated attack game we discussed earlier. Each player might be unwilling to play her prescribed strategy if she is not absolutely certain that the other play will do the same.

3. Equilibrium theory does a poor job explaining play in early rounds of most experiments, although it does much better in later rounds. This shift from non-equilibrium to equilibrium play is difficult to reconcile with a purely introspective theory.

1.1 Learning or Evolution?

There are two main ways to model the processes according to which players change their strategies they are using to play a game. A *learning model* is any model that specifies the learning rules used by individual players and examines their interaction when the game (or games) is played repeatedly. These types of models will be the subject of today's lecture.

Learning models quickly become very complex when there are many players involved. *Evolutionary models* do not specifically model the learning process at the individual level. The basic assumption there is that some unspecified process at the individual level leads the population as a whole to adopt strategies that yield improved payoffs. These type of models will the subject of the next few lectures.

1.2 Population Size and Matching

The natural starting point for any learning (or evolutionary) model is the case of fixed players. Typically, we will only look at 2 by 2 games which are played repeatedly between these two fixed players. Each player faces the task of inferring future play from the past behavior of agents.

There is a serious drawback from working with fixed agents. Due to the repeated interaction in every game players might have an incentive to influence the future play of their opponent. For example, in most learning models players will defect in a Prisoner's dilemma because cooperation is strictly dominated for any beliefs I might hold about my opponent. However, if I interact frequently with the same opponent, I might try to cooperate in order to 'teach' the opponent that I am a cooperator. We will see in a future lecture that such behavior can be in deed a Nash equilibrium in a *repeated game*.

There are several ways in which repeated play considerations can be assumed away.

1. We can imagine that players are locked into their actions for quite

a while (they invest infrequently, can't build a new factory overnight etc.) and that their discount factors (the factor by which they weight the future) is small compared that lock-in length. It them makes sense to treat agents as approximately myopic when making their decisions.

2. An alternative is to dispense with the fixed player assumption, and instead assume that agents are drawn from a large population and are randomly matched against each other to play games. In this case, it is very unlikely that I encounter a recent opponent in a round in the near future. This breaks the strategic links between the rounds and allows us to treat agents as approximately myopic again (i.e. they maximize their short-term payoffs).

2 Cournot Adjustment

In the Cournot adjustment model two fixed players move sequentially and choose a best response to the play of their opponent in the last period.

The model was originally developed to explain learning in the Cournot model. Firms start from some initial output combination (q_1^0, q_2^0). In the first round both firms adapt their output to be the best response to q_2^0. They therefore play $(BR_1(q_2^0), BR_2(q_1^0))$.

This process is repeated and it can be easily seen that in the case of linear demand and constant marginal costs the process converges to the unique Nash equilibrium. If there are several Nash equilibria the initial conditions will determine which equilibrium is selected.

2.1 Problem with Cournot Learning

There are two main problems:

- Firms are pretty dim-witted. They adjust their strategies today as if they expect firms to play the same strategy as yesterday.
- In each period play can actually change quite a lot. Intelligent firms should anticipate their opponents play in the future and react accordingly. Intuitively, this should speed up the adjustment process.

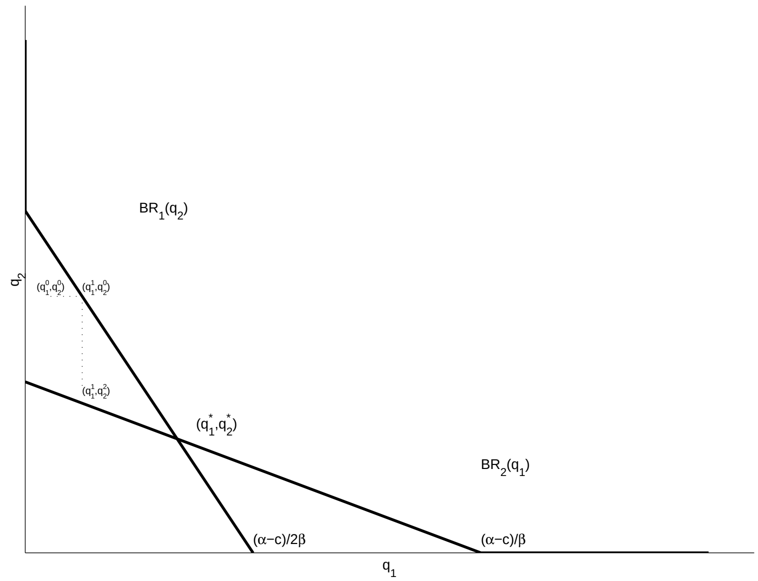

Cournot adjustment can be made more realistic by assuming that firms are 'locked in' for some time and that they move alternately. Firms 1 moves in period 1,3,5,... and firm 2 moves in periods 2,4,6,.. Starting from some initial play (q_1^0, q_2^0), firms will play (q_1^1, q_2^0) in round 1 and (q_1^1, q_2^2) in round 2. Clearly, the Cournot dynamics with alternate moves has the same long-run behavior as the Cournot dynamics with simultaneous moves.

Cournot adjustment will be approximately optimal for firms if the lock-in period is large compared to the discount rate of firms. The less locked-in firms are the smaller the discount rate (the discount rate is the weight on next period's profits).

Of course, the problem with the lock-in interpretation is the fact that it is not really a model of learning anymore. Learning is irrelevant because firms choose their optimal action in each period.

3 Fictitious Play

In the process of fictitious play players assume that their opponents strategies are drawn from some stationary but unknown distribution. As in the Cournot adjustment model we restrict attention to a fixed two-player setting. We also assume that the strategy sets of both players are finite.

In fictitious play players choose the best response to their assessment of their opponent's strategy. Each player has some exogenously given weighting function $\kappa_i^0 : S_{-i} \to \Re^+$. After each period the weights are updated by adding 1 to each opponent strategy each time is has been played:

$$\kappa_i^t\left(s_{-i}\right) = \begin{cases} \kappa_i^{t-1}\left(s_{-i}\right) & \text{if } s_{-i} \neq s_{-i}^{t-1} \\ \kappa_i^{t-1}\left(s_{-i}\right) + 1 & \text{if } s_{-i} = s_{-i}^{t-1} \end{cases}$$

Player i assigns probability $\gamma_i^t\left(s_{-i}\right)$ to strategy profile s_{-i}:

$$\gamma_i^t\left(s_{-i}\right) = \frac{\kappa_i^t\left(s_{-i}\right)}{\sum_{\tilde{s}_{-i} \in S_{-i}} \kappa_i^t\left(\tilde{s}_{-i}\right)}$$

The player then chooses a pure strategy which is a best response to his assessment of other players' strategy profiles. Note that there is not necessarily a unique best-response to every assessment - hence fictitious play is not always unique.

We also define the empirical distribution $d_i^t\left(s_i\right)$ of each player's strategies as

$$d_i^t\left(s_i\right) = \frac{\sum_{\tilde{t}=0}^{t} I^{\tilde{t}}\left(s_i\right)}{t}$$

The indicator function is set to 1 if the strategy has been played in period $\tilde{t}$ and 0 otherwise. Note, that as $t \to \infty$ the empirical distribution d_j^t of player j's strategies approximate the weighting function κ_i^t (since in a two player game we have $j = -i$).

Remark 1 *The updating of the weighting function looks intuitive but also somewhat arbitrary. It can be made more rigorous in the following way. Assume, that there are n strategy profiles in S_{-i} and that each profile is played by player i's opponents' with probability $p\left(s_{-i}\right)$. Agent i has a prior belief according to which these probabilities are distributed. This prior is a Dirichlet distribution whose parameters depend on the weighting function. After each round agents update their prior: it can be shown that the posterior belief is again Dirichlet and the parameters of the posterior depend now on the updated weighting function.*

3.1 Asymptotic Behavior

Will fictitious play converge to a Nash equilibrium? The next proposition gives a partial answer.

Proposition 1 *If s is a strict Nash equilibrium, and s is played at date t in the process of fictitious play, s is played at all subsequent dates. That is, strict Nash equilibria are absorbing for the process of fictitious play. Furthermore, any pure-strategy steady state of fictitious play must be a Nash equilibrium.*

Proof : Assume that $s = (s_i, s_j)$ is played at time t. This implies that s_i is a best-response to player i's assessment at time t. But his next period assessment will put higher relative weight on strategy s_j. Because s_i is a BR to s_j and the old assessment it will be also a best-response to the updated assessment. Conversely, if fictitious play gets stuck in some pure steady state then players' assessment converge to the empirical distribution. If the steady state is not Nash players would eventually deviate.

A corollary of the above result is that fictitious play cannot converge to a pure steady state in a game which has only mixed Nash equilibria such as matching pennies.

	H	T
H	1,−1	−1,1
T	−1,1	1,−1

Assume that both players have initial weights weights (1.5,2) and (2,1.5). Then fictitious play cycles as follows: In the first period, 1 and 2 play T, so the weights the next period are (1.5,3) and (2, 2.5). Then 1 plays T and 2 plays H for the next two periods, after which 1's weights are (3.5,3) and 2's are (2,4.5). At this point 1 switches to H, and both players play H for the next three periods, at which point 2 switches to T, and so on. It may not be obvious, but although the actual play in this example cycles, the empirical distribution over each player's strategies are converging to $\left(\frac{1}{2}, \frac{1}{2}\right)$ - this is precisely the unique mixed Nash equilibrium.

This observation leads to a general result.

Proposition 2 *Under fictitious play, if the empirical distributions over each player's choices converge, the strategy profile corresponding to the product of these distributions is a Nash equilibrium.*

Proof : Assume that there is a profitable deviation. Then in the limit at least one player should deviate - but this contradicts the assumption that strategies converge.

These results don't tell us when fictitious play converges. The next theorem does precisely that.

Theorem 1 *Under fictitious play the empirical distributions converge if the stage has generic payoffs and is 2 2, or zero sum, or is solvable by iterated strict dominance.*

We won't prove this theorem in this lecture. However, it is intuitively clear why fictitious play observes IDSDS. A strictly dominated strategy can never be a best response. Therefore, in the limit fictitious play should put zero relative weight on it. But then all strategies deleted in the second step can never be best responses and should have zero weight as well etc.

3.2 Non-Convergence is Possible

Fictitious play does not have to converge at all. An example for that is due to Shapley.

	L	M	R
T	0,0	1,0	0,1
M	0,1	0,0	1,0
D	1,0	0,1	0,0

The unique mixed NE of the game is $s_1 = s_2 = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

If the initial play is (T, M) then the sequence becomes (T, M), (T, R), (M, R), (M, L), (D, L), (D, M), (T, M).

One can show that the number of time each profile is played increases at a fast enough rate such that the play never converges. Also note, that the diagonal entries are never played.

3.3 Pathological Convergence

Convergence in the empirical distribution of strategy choices can be misleading even though the process converges to a Nash equilibrium. Take the following game:

	A	B
A	0,0	1,1
B	1,1	0,0

Assume that the initial weights are $\left(1, \sqrt{2}\right)$ for both players. In the first period both players think the other will play B, so both play A. The next period the weights are $\left(2, \sqrt{2}\right)$, and both play B; the outcome is the alternating sequence (B, B), (A, A), (B, B), and so on. The empirical frequencies of each player's choices converge to 1/2, 1/2, which is the Nash equilibrium. The realized play is always on the diagonal, however, and both players receive payoff 0 each period. Another way of putting this is that the empirical joint distribution on pairs of actions does not equal the product of the two marginal distributions, so the empirical joint distribution corresponds to correlated as opposed to independent play.

This type of correlation is very appealing. In particular, agents don't seem to be smart enough to recognize cycles which they could exploit. Hence the attractive property of convergence to a Nash equilibrium can be misleading if the equilibrium is mixed.

Lecture IX: Evolution

Markus M. Möbius

March 10, 2004

Learning and evolution are the second set of topics which are not discussed in the two main texts. I tried to make the lecture notes self-contained.

- Fudenberg and Levine (1998), The Theory of Learning in Games, Chapter 1 and 2

1 Introduction

For models of learning we typically assume a fixed number of players who find out about each other's intentions over time. In evolutionary models the process of learning is not explicitly modeled. Instead, we assume that strategies which do better on average are played more often in the population over time. The biological explanation for this is that individuals are genetically programmed to play one strategy and their reproduction rate depends on their fitness, i.e. the average payoff they obtain in the game. The economic explanation is that there is social learning going on in the background - people find out gradually which strategies do better and adjust accordingly. However, that adjustment process is slower than the rate at which agents play the game.

We will focus initially on models with random matching: there are N agents who are randomly matched against each other over time to play a certain game. Frequently, we assume that N as infinite. We have discussed in the last lecture that random matching gives rise to myopic play because there are no repeated game concerns (I'm unlikely to ever encounter my current opponent again).

We will focus on symmetric n by n games for the purpose of this course. In a symmetric game each player has the same strategy set and the payoff matrix satisfies $u_i(s_i, s_j) = u_j(s_j, s_i)$ for each player i and j and strategies $s_i, s_j \in S_i = S_j = \{s_1, .., s_n\}$. Many games we encountered so far in the course are symmetric such as the Prisoner's Dilemma, Battle of the Sexes, Chicken and all coordination games. In symmetric games both players face exactly the same problem and their optimal strategies do not depend on whether they play the role of player 1 or player 2.

An important assumption in evolutionary models is that each agent plays a fixed pure strategy until she dies, or has an opportunity to learn and about her belief. The game is fully specified if we know the fraction of agents who play strategy $s_1, s_2, .., s_n$ which we denote with $x_1, x_2, .., x_n$ such that $\sum_{i=1}^{n} x_i = 1$.

2 Mutations and Selection

Every model of evolution relies on two key concepts - a *mutation mechanism* and a *selection mechanism.* We have already discussed selection - strategies spread if they give above average payoffs. This captures social learning in a reduced form.

Mutations are important to add 'noise' to the system (i.e. ensure that $x_i > 0$ at all times) and prevent it from getting 'stuck'. For example, in a world where players are randomly matched to play a Prisoner's Dilemma mutations make sure that it will never be the case that all agents cooperate or all agents defect because there will be random mutations pushing the system away from the two extremes.[1]

3 Replicator Dynamics

The replicator dynamics is one particular selection mechanism which captures the notion that strategies with above average payoff should spread in the population. Typically, the replicator dynamics is modelled without allowing for mutations - the dynamics therefore becomes deterministic.

Since the stage game is symmetric we know that $u(s_i, s_j) = u_1(s_i, s_j) = u_2(s_j, s_i)$. The average payoff of strategy s_i for a player is $u(s_i, x)$ since he

[1] If we start from an all-cooperating state, mutating agents will defect and do better. Hence they spread, and finally take over. Without mutations the system might be 'stuck' in the all-cooperating state.

is randomly matched with probability x_j against agents playing s_j ($x = (x_1, x_2, .., x_n)$). The average payoff of all strategies is $\sum_{i=1}^{n} x_i u(s_i, x) = u(x, x)$.

Strategy s_i does better than the average strategy if $u(s_i, x) > u(x, x)$, and worse otherwise. A minimum requirement for a selection mechanism is that $sgn(\dot{x}_i(t)) = sgn[u(s_i, x) - u(x, x)]$. The share x_i increases over time if and only if strategy s_i does better than average. The replicator dynamics is one particular example:

Definition 1 *In the replicator dynamics the share x_i of the population playing strategy s_i evolves over time according to:*

$$\frac{\dot{x}_i}{x_i} = u(s_i, x) - u(x, x)$$

If $x_i = 0$ at time 0 then we have $x_i = 0$ at all subsequent time periods: if nobody plays strategy s_i then the share of population playing it can neither decrease not increase.

The next proposition makes sure that the definition is consistent (i.e. population share always sum up to 1).

Proposition 1 *The population shares always sum up to* 1*.*

Proof: We can write:

$$\sum \dot{x}_i = \sum x_i u(s_i, x) - \sum x_i u(x, x) = u(x, x) - u(x, x) = 0$$

This establishes that $\sum x_i = const$. The constant has to be 1 because the population shares sum up to 1 initially.

3.1 Steady States and Nash Equilibria

Definition 2 *The strategy σ is a steady state if for $x_i = \sigma_i(s_i)$ we have* $\frac{dx}{dt} = 0$.

Proposition 2 *If σ is the strategy played by each player in a symmatric mixed NE then it is a steady state.*

Proof: In a NE each player has to be indifferent between the strategies in her support. Therefore, we have $u(s_i, x) = u(x, x)$.

Note, that the reverse is NOT true. If all players cooperate in a Prisoner's Dilemma this will be a steady state (since there are no mutations).

3.1.1 Example I

In 2 by 2 games the replicator dynamics is easily understood. Look at the following game:

	A	B
A	0,0	1,1
B	1,1	0,0

There are only two types in the population and $x = (x_A, x_B)$. It's enough to keep track of x_A.

$$\frac{\dot{x}_A}{x_A} = x_B - 2x_A x_B = (1 - x_A)(1 - 2x_A) \tag{1}$$

It's easy to see that $\dot{x}_A > 0$ for $0 < x_A < \frac{1}{2}$ and $\dot{x}_A < 0$ for $1 > x_A > \frac{1}{2}$. This makes $x_A = 1/2$ a 'stable' equilibrium (see below for a precise definition).

3.1.2 Example II

Now look at the New York game.

	E	C
E	1,1	0,0
C	0,0	1,1

We can show:

$$\frac{\dot{x}_E}{x_E} = x_E - (x_E x_E + x_C x_C) = (1 - x_E)(2x_E - 1) \quad (2)$$

Now the steady state $x_E = \frac{1}{2}$ is 'unstable'.

3.2 Stability

Definition 3 *A steady state σ is stable if for all $\epsilon > 0$ there exists $\delta > 0$ such that if the process starts a distance δ away from the steady state it will never get further away than ϵ.*

The mixed equilibrium is stable in example I and unstable in example II.

Definition 4 *A steady state σ is asymptotically stable if there exists some $\delta > 0$ such that the process converges to σ if it starts from a distance at most δ away from the steady state.*

The mixed equilibrium in example I is asymptotically stable.

Definition 5 *A steady state σ is globally stable if the process converges to the steady state from any initial state where $x_i > 0$ for all s_i.*

Stability can be interpreted as a form of equilibrium selection.

Theorem 1 *If the steady state σ is stable then it is a NE.*

Proof: Assume not. Then there exists a profitable deviation s_i such that $u(s_i, \sigma) - u(\sigma, \sigma) = b > 0$. Because the linear utility function is uniformly continuous there exists some ϵ such that for all x a distance less than ϵ away we have $|u(s_i, x) - u(s_i, \sigma)| < \frac{b}{4}$ and $|u(x, x) - u(\sigma, \sigma)| < \frac{b}{4}$. This implies that $|u(s_i, x') - u(x, x)| > \frac{b}{2}$. Because σ is stable we know that for x close enough to σ (less than a distance δ the dynamics converges to σ. So take $x(0) = \left(1 - \frac{1}{2}\delta\right) x + \frac{1}{2}\delta s_i$. Then we have $\frac{dx_i}{dt} = x_i (u(s_i, x) - u(x, x)) \geq x_i \frac{b}{2} \geq \frac{\frac{\delta}{2} b}{2}$. But this implies that $x(t) \rightarrow \infty$ which is a contradiction.

3.2.1 Example III

This part is harder and is NOT required for the exam.

The mixed equilibrium in the RPS game is stable but not asymptotically stable. The RPS game is harder to analyze because each player has three possible strategies. This implies that there are two differential equations to keep track of.

	R	P	S
R	0,0	−1,1	1,−1
P	1,−1	0,0	−1,1
S	−1,1	1,−1	0,0

We can concentrate on x_R and x_P. The average payoff is

$$u(x, x) = x_R[-x_P + x_S] + x_P[x_R - x_S] + x_S[-x_R + x_P] \quad (3)$$
$$= 0 \quad (4)$$

We then get:

$$\frac{\dot{x}_R}{x_R} = -x_P + x_S \quad (5)$$
$$\frac{\dot{x}_P}{x_P} = x_R - x_S \quad (6)$$

The unique mixed equilibrium $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ is the only steady state of the system.

To see whether it is stable we have to linearize the system around the steady state. We define $\tilde{x}_i = x_i - \frac{1}{3}$ and obtain:

$$3\dot{\tilde{x}}_R = -\tilde{x}_R - 2\tilde{x}_P \quad (7)$$
$$3\dot{\tilde{x}}_P = 2\tilde{x}_R + \tilde{x}_P \quad (8)$$

We have to find the eigenvalues of the system which are $\lambda = \pm\sqrt{3}i$. This implies that the population shares 'circle' around the steady state.

4 ESS-Evolutionary Stable States

The ESS concept concentrates on the role of mutations. Intuitively, the stability notion in the replicator dynamics is already a selection criterion because if the system is disturbed a little bit, it moves away from unstable steady states. The ESS concept expands on this intuition.

Definition 6 *The state x is ESS if for all $y \neq x$ there exists $\bar{\epsilon}$ such that $u(x, (1-\epsilon)x + \epsilon y) > u(y, (1-\epsilon)x + \epsilon y)$ for all $0 < \epsilon < \bar{\epsilon}$.*

This definition captures the idea that a stable state is impervious to mutations since they do worse than the original strategy.

Proposition 3 *x is ESS iff for all $y \neq x$ either (a) $u(x, x) > u(y, x)$ or (b) $u(x, x) = u(y, x)$ and $u(x, y) > u(y, y)$.*

The following results establish the link between Nash equilibrium and ESS.

Proposition 4 *If x is ESS then it is a NE.*

Proposition 5 *If x is a strict NE then it is an ESS.*

Proposition 6 *If x is a totally mixed ESS then it is the unique ESS.*

The next theorem establishes the link between ESS and stability under the replicator dynamics.

Theorem 2 *If x is an ESS then it is asymptotically stable under the Replicator Dynamics.*

5 Stochastic Stability

Stochastic stability combines both mutations and a selection mechanism. It provides a much more powerful selection mechanism for Nash equilibria than ESS and the replicator dynamics.

The ESS concept gives us some guidance as to which Nash equilibria are stable under perturbations or 'experimentation' by agents. The replicator dynamics tells us how a group of agents can converge to some steady state starting from initial conditions. However, selection relies in many cases on the initial conditions (take the coordination game, for example). In particular, we cannot select between multiple *strict* equilibria.

For this section we concentrate on generic coordination games:

	A	B
A	a,a	b,c
B	c,b	d,d

We assume, that $a > c$ and $d > b$ such that both (A, A) and (B, B) are NE of the game. We assume that $a + b > c + d$ such that A is the risk-dominant strategy for each player. Note, that the risk-dominant NE is not necessarily the Pareto optimal NE.

We have seen in experiments that agents tend to choose A. Can we justify this with an evolutionary story?

YES!

Assume that there is a finite number n of agents who are randomly matched against each other in each round. Assume that agents choose the best response to whatever strategy did better in the population in the last period. This is called the BR dynamics. Clearly, all agents will choose A if $x_A > q^* = \frac{d-b}{a-b+d-c}$. Because A is risk-dominant we have $q^* < \frac{1}{2}$.

There are also mutations in each period. Specifically, with probability ϵ each agent randomizes between both strategies.

We define the basin of attraction of $x_A = 1$ (everybody plays A) to be $B_A = [q^*, 1]$ and the basin of attraction of $x_A = 0$ to be $[0, q^*]$. Whenever the initial state of the system is within the basin of attraction it converges to all A/ all B for sure *if there are no mutations*. We define the radius of the basin B_A to be the number of mutations it takes to 'jump out' of the state all A. We get $R_A = (1 - q^*)\, n$. Similarly, we define the co-radius CR_A as the number of mutations it takes at most to 'jump into' the basin B_A. We get $CR_A = q^* n$.

Theorem 3 *If $CR_A < R_A$ then the state 'all A' is stochastically stable. The waiting time to reach the state 'all A' is of the order ϵ^{CR_A}.*

Therefore, the risk-dominant equilibrium is stochastically stable as $q^* < \frac{1}{2}$.

5.1 The Power Local Interaction

Local interaction can significantly speed up the evolution of the system. Assume, that agents are located on the circle and play the BR to average play of their direct neighbors in the previous period. It can be shown that $CR_A = 2$ and $CR_A = \frac{n}{2}$. The convergence is a lot faster than under global interaction.

Lecture X: Extensive Form Games

Markus M. Möbius

March 17, 2004

- Gibbons, chapter 2
- Osborne, sections 5.1, 5.2 and chapter 6

1 Introduction

While models presented so far are fairly general in some ways it should be noted that they have one main limitation as far as accuracy of modeling goes - in each game each player moves once and moves simultaneously.

This misses common features both of many classic games (bridge, chess) and of many economic models.[1] A few examples are:

- auctions (sealed bid versus oral)
- executive compensation (contract signed; then executive works)
- patent race (firms continually update resources based on where opponent are)
- price competition: firms repeatedly charge prices
- monetary authority and firms (continually observe and learn actions)

Topic today is how to represent and analyze such games.

[1] Gibbons is a good reference for economic applications of extensive form games.

2 Extensive Form Games

The extensive form of a game is a complete description of

1. The 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 players know when they move.

2.1 Example I: Model of Entry

Currently firm 1 is an incumbent monopolist. A second firm 2 has the opportunity to enter. After firm 2 makes the decision to enter, firm 1 will have the chance to choose a pricing strategy. It can choose either to *fight* the entrant or to *accommodate* it with higher prices.

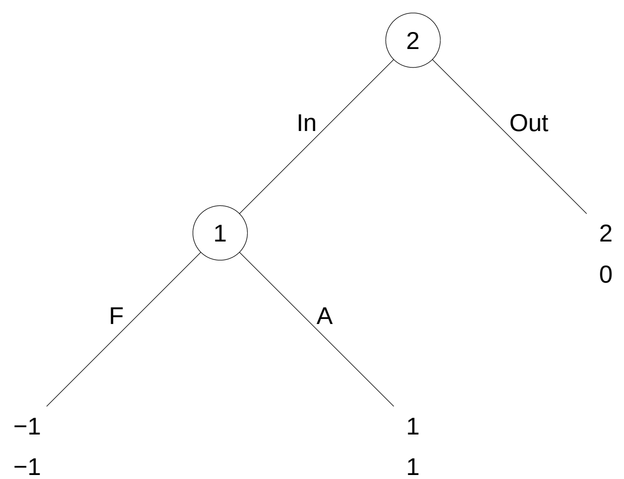

2.2 Example II: Stackelberg Competition

Suppose firm 1 develops a new technology before firm 2 and as a result has the opportunity to build a factory and commit to an output level q_1 before firm 2 starts. Firm 2 then observes firm 1 before picking its output level q_2. For concreteness suppose $q_i \in \{0, 1, 2\}$ and market demand is $p(Q) = 3 - Q$. The marginal cost of production is 0.

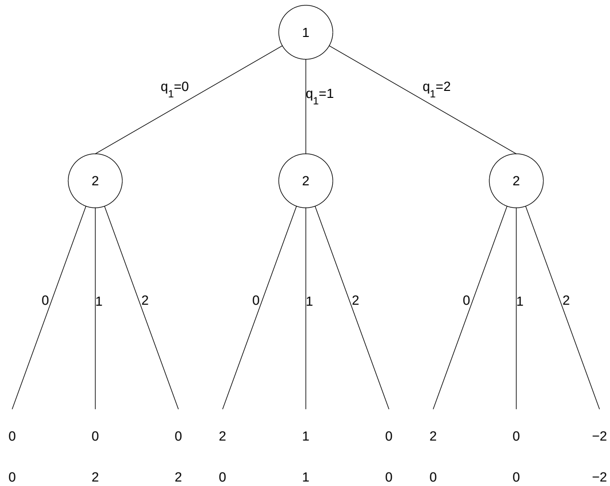

2.3 Example III: Matching Pennies

So far we assumed that players can observe all previous moves. In order to model the standard matching pennies game in extensive form we have to assume that the second player cannot observe the first player's move.

Sequential matching pennies is represented as follows:

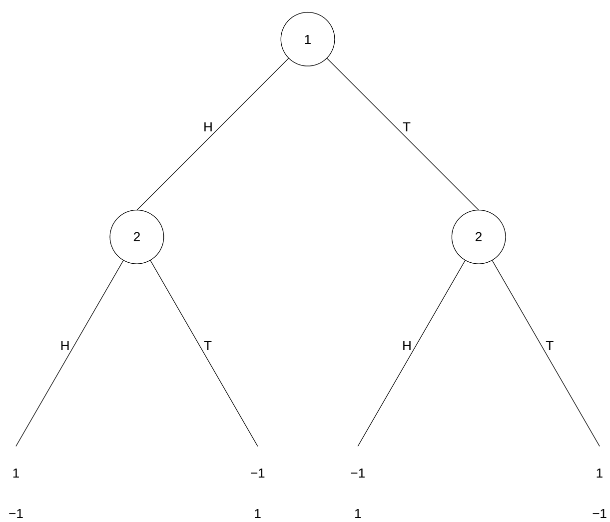

If we want to indicate that player 2 cannot observe the move of player 1 we depict the game as follows:

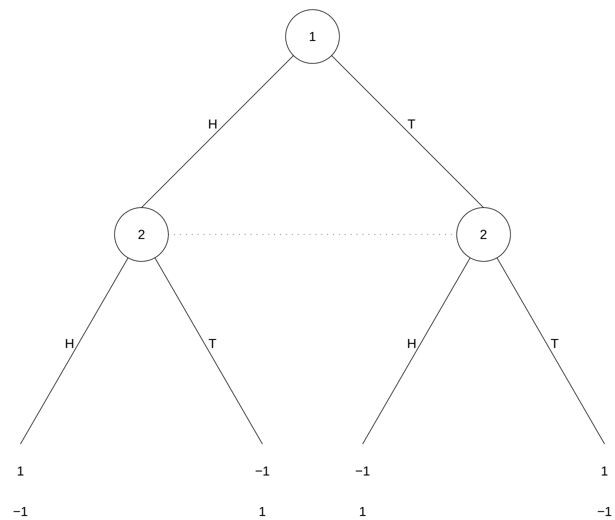

The extensive form representation allows that players can 'forget' information. For example we can assume that in a game with 4 rounds player 2 can observe player 1's move in round 1, but in round 4 he has forgotten the move of player 1. In most cases, we assume *perfect recall* which rules out that players have such 'bad memory'.[2]

3 Definition of an Extensive Form Game

Formally a finite extensive form game consists of

1. A finite set of players.
2. A finite set T of nodes which form a tree along with functions giving for each non-terminal node $t \notin Z$ (Z is the set of terminal nodes)
    - the player $i(t)$ who moves
    - the set of possible actions $A(t)$
    - the successor node resulting from each possible action $N(t, a)$

[2]It becomes difficult to think of a solution concept of a game where players are forgetful. Forgetfulness and rational behavior don't go well together, and concepts like Nash equilibrium assume that players are rational.

3. Payoff functions $u_i : Z \to \Re$ giving the players payoffs as a function of the terminal node reached (the terminal nodes are the outcomes of the game).

4. An information partition: for each node t, $h(t)$ is the set of nodes which are possible given what player $i(x)$ knows. This partition must satisfy

$$t' \in h(x) \Rightarrow i(t') = i(t), A(t') = A(t), \text{and} \quad h(t') = h(t)$$

We sometimes write $i(h)$ and $A(h)$ since the action set is the same for each node in the same information set.

It is useful to go over the definition in detail in the matching pennies game where player 2 can't observe player 1's move. Let's number the non-terminal nodes t_1, t_2 and t_3 (top to bottom).

1. There are two players.

2. $S_1 = S_2 = \{H, T\}$ at each node.

3. The tree defines clearly the terminal nodes, and shows that t_2 and t_3 are successors to t_1.

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

4 Normal Form Analysis

In an extensive form game write H_i for the set of information sets at which player i moves.

$$H_i = \{S \subset T | S = h(t) \text{ for some } t \in T \text{ with } i(t) = i\}$$

Write A_i for the set of actions available to player i at any of his information sets.

Definition 1 *A pure strategy for player i in an extensive form game is a function $s_i : H_i \to A_i$ such that $s_i(h) \in A(h)$ for all $h \in H_i$.*

Note that a strategy is a **complete contingent plan** explaining what a player will do in *any* situation that arises. At first, a strategy looks overspecified: earlier action might make it impossible to reach certain sections of a tree. Why do we have to specify how players would play at nodes which can never be reached if a player follows his strategy early on? The reason is that play off the equilibrium path is crucial to determine if a set of strategies form a Nash equilibrium. Off-equilibrium threats are crucial. This will become clearer shortly.

Definition 2 *A mixed behavior strategy for player* i *is a function* $\sigma_i : H_i \rightarrow \Delta(A_i)$ *such that* $supp(\sigma_i(h)) \subset A(h)$ *for all* $h \in A_i$.

Note that we specify an independent randomization at each information set![3]

4.1 Example I: Entry Game

We can find the pure strategy sets $S_1 = \{\text{Fight}, \text{Accomodate}\}$ and $S_2 = \{\text{Out}, \text{In}\}$. We can represent the game in normal form as:

	Out	In
F	2,0	−1,−1
A	2,0	1,1

[3] You might think that a more natural definition would simply define mixed strategies as a randomization over a set of pure strategies (just as in simultaneous move games). It can be shown that for games with perfect recall this definition is equivalent to the one given here, i.e. a mixed strategy is a mixed behavior strategy and vice versa. In games without perfect recall this is no longer true - it is instructive to convince yourself that in such games each mixed behavior strategy is a mixed strategy but not vice versa.

4.2 Example II: Stackelberg Competition

Firm 1 chooses q_1 and firm 2 chooses a quantity $q_2(q_1)$. With three possible output levels, firm 1 has three strategies, while firm 2 has $3^3 = 9$ different strategies because it can choose three strategies at its three information sets.

4.3 Example III: Sequential Matching Pennies

We have $S_1 = \{H, T\}$. Firm 2 has four strategies as it can choose two actions at two information sets. Strategy HH implies that firm 2 chooses H at both nodes, while HT implies that it chooses H in the left node (after having observed H) and T in the right node (after having observed T).

	HH	HT	TH	TT
H	1,−1	1,−1	−1,1	−1,1
T	−1,1	1,−1	−1,1	1,−1

4.4 Example IV

Look at the following extensive form game:

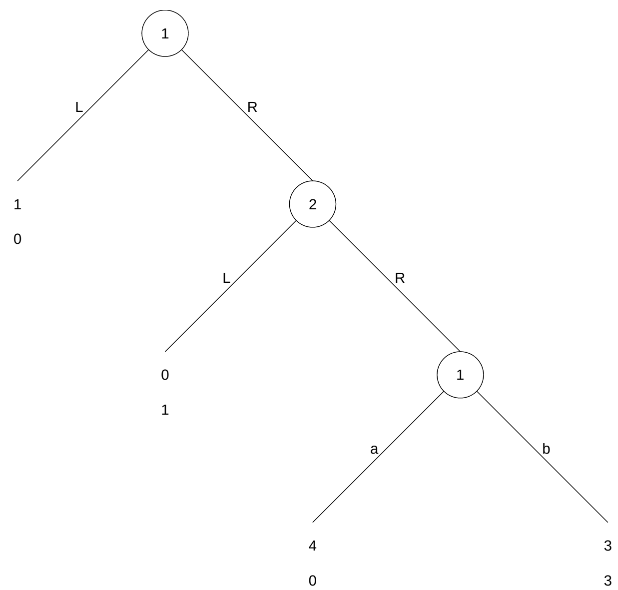

One might be tempted to say that player 1 has three strategies because there are only three terminal nodes which can be reached. However, there are 4 because La and Lb are two distinct strategies. After player 1 plays L it is irrelevant for the final outcome what he would play in the bottom node. However, this off equilibrium pay is important for player 2's decision process which in turn makes 1 decide whether to play L or R.

	L	R
La	1,0	1,0
Lb	1,0	1,0
Ra	0,1	4,0
Rb	0,1	3,3

5 Nash Equilibrium in Extensive Form Games

We can apply NE in extensive form games simply by looking at the normal form representation. It turns out that this is not an appealing solution concept because it allows for too many profiles to be equilibria.

Look at the entry game. There are two pure strategy equilibria: (A,In) and (F,Out) as well as mixed equilibria $(\alpha F + (1-\alpha)\,A, \text{Out})$ for $\alpha \geq \frac{1}{2}$.

Why is (F,Out) a Nash equilibrium? Firm 2 stays out because he thinks that player 2 will fight entry. In other words, the threat to fight entry is sufficient to keep firm 2 out. Note, that in equilibrium this threat is never played since firm 2 stays out in the first place.

The problem with this equilibrium is that firm 2 could call firm 1's bluff and enter. Once firm 2 has entered it is in the interest of firm 1 to accommodate. Therefore, firm 1's threat is *not credible*. This suggests that only (A,In) is a reasonable equilibrium for the game since it does not rely on non-credible threats. The concept of subgame perfection which we will introduce in the next lecture rules out non-credible threats.

5.1 Example II: Stackelberg

We next look at a Stackelberg game where each firm can choose $q_i \in [0,1]$ and $p = 1 - q_1 - q_2$ and $c = 0$.

Claim: For any $q_1' \in [0,1]$ the game has a NE in which firm 1 produces q_1'.

Consider the following strategies:

$$\begin{aligned} s_1 &= q_1' \\ s_2 &= \begin{cases} \frac{1-q_1'}{2} & \text{if } q_1 = q_1' \\ 1 - q_1' & \text{if } q_1 \neq q_1' \end{cases} \end{aligned}$$

In words: firm 2 floods the market such that the price drops to zero if firm 1 does not choose q_1'. It is easy to see that these strategies form a NE. Firm 1 can only do worse by deviating since profits are zero if firm 2 floods the market. Firm 2 plays a BR to q_1' and therefore won't deviate either.

Note, that in this game things are even worse. Unlike the Cournot game where we got a unique equilibrium we now have a continuum of equilibria. Second, we have even more disturbing non-credible threats. For instance, in the equilibrium where $q_1' = 1$ firm 2's threat is "if you don't flood the market and destroy the market I will". Not only won't the threat be carried out - it's also hard to see why it would be made in the first place.

Lecture XI: Subgame Perfect Equilibrium

Markus M. Möbius

April 3, 2004

- Gibbons, chapter 2.1.A,2.1.B,2.2.A
- Osborne, sections 5.4, 5.5

1 Introduction

Last time we discussed extensive form representation and showed that there are typically lots of Nash equilibria. Many of them look unreasonable because they are based on out of equilibrium threats. For example, in the entry game the incumbent can deter entry by threatening to flood the market. In equilibrium this threat is never carried out. However, it seems unreasonable because the incumbent would do better accommodating the entrant if entry in fact occurs. In other words, the entrant can call the incumbent's bluff by entering anyway.

Subgame perfection is a refinement of Nash equilibrium. It rules out non-credible threats.

2 Subgames

Definition 1 *A subgame G' of an extensive form game G consists of*

1. *A subset T' of the nodes of G consisting of a single node x and all of its successors which has the property that $t \in T'$, $t' \in h(t)$ then $t' \in T'$.*
2. *Information sets, feasible moves and payoffs at terminal nodes as in G.*

2.1 Example I: Entry Game

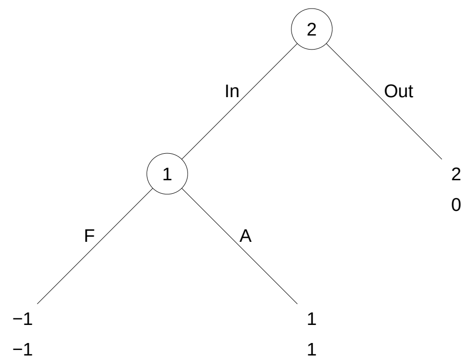

This game has two subgames. The entire game (which is *always* a subgame) and the subgame which is played after player 2 has entered the market:

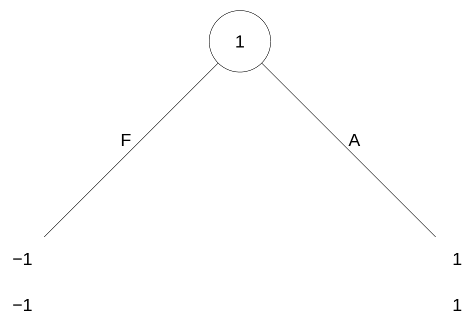

2.2 Example II

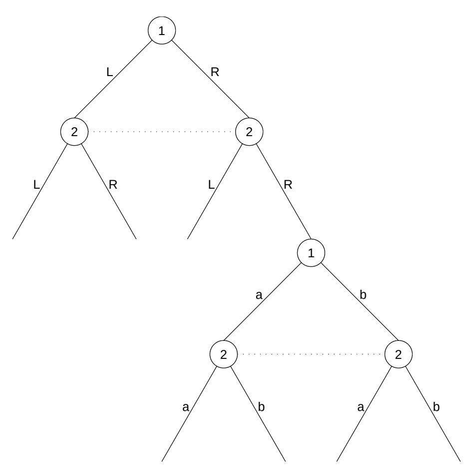

This game has also two subgames. The entire game and the subgame (a simultaneous move game) played after round 2:

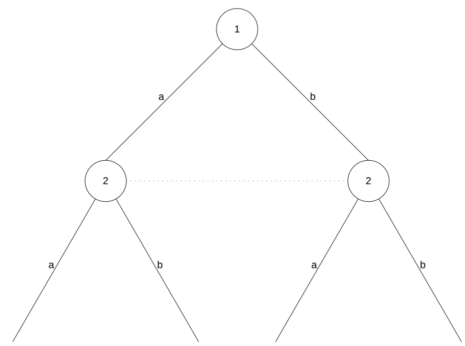

This subgame has no further subgames: otherwise the information set of player 2 would be separated which is not allowed under our definition.

3 Subgame Perfect Equilibrium

Definition 2 *A strategy profile s^* is a subgame perfect equilibrium of G if it is a Nash equilibrium of every subgame of G.*

Note, that a SPE is also a NE because the game itself is a (degenerate) subgame of the entire game.

Look at the entry game again. We can show that $s_1 = A$ and $s_2 = \text{Entry}$ is the unique SPE. Accomodation is the unique best response in the subgame after entry has occurred. Knowing that, firm 2's best response is to enter.

3.1 Example: Stackelberg

We next continue the Stackelberg example from the last lecture. We claim that the unique SPE is $q_2^* = \frac{1}{2}$ and $q_1^*(q_2) = \frac{1-q_2}{2}$.

The proof is as follows. A SPE must be a NE in the subgame after firm 1 has chosen q_1. This is a one player game so NE is equivalent to firm 1 maximizing its payoff, i.e. $q_1^*(q_1) \in \arg\max q_1 [1 - (q_1 + q_2)]$. This implies that $q_1^*(q_2) = \frac{1-q_2}{2}$. Equivalently, firm 1 plays on its BR curve.

A SPE must also be a NE in the whole game, so q_2^* is a BR to q_1^*:

$$u_2(q_1, q_2^*) = q_2 (1 - (q_2 + q_1^*(q_2))) = q_1 \frac{1 - q_1}{2}$$

The FOC for maximizing u_2 is $q_2^* = \frac{1}{2}$.

Note that firm 2 (the first mover) produces more than in Cournot.

There are many games which fit the Stackelberg paradigm such as monetary policy setting by the central bank, performance pay for managers etc. We will discuss general results for this class of games in the next lecture.

4 Backward Induction

The previous example illustrates the most common technique for finding and verifying that you have found the SPE of a game. Start at the end of the game and work your way from the start.

We will focus for the moment on extensive form games where each information set is a single node (i.e. players can perfectly observe all previous moves).

Definition 3 *An extensive form is said to have perfect information if each information set contains a single node.*

Proposition 1 *Any finite game of perfect information has a pure strategy SPE. For generic payoffs in a finite extensive form game with perfect information the SPE is unique.*

What does generic mean? With generic payoffs players are never indifferent between two strategies. If payoffs are randomly selected at the terminal nodes then indifference between two actions is a zero probability event. More mathematically, we can say that the results holds for almost all games.

Proof: I did it in class, and I do a more general proof in the next section for games with imperfect information. Intuitively, you solve the last rounds of the game, then replace these subgames with the (unique) outcome of the NE and repeat the procedure.

What happens if players are indifferent between two strategies at some point? Then there is more than one SPE, and you have to complete the backward induction for each possible outcome of the subgame.

Consider the following game:

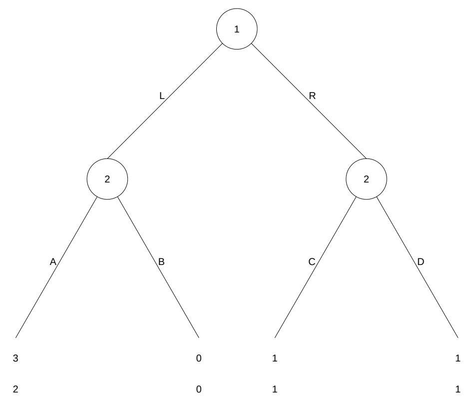

After player 1 played R player 2 is indifferent between C and D. It is easy to see that there are infinitely many SPE such that $s_1^* = L$ and $s_2^* = (A, \alpha C + (1-\alpha) D)$ for $0 \leq \alpha \leq 1$.

Note however, that each of these SPE yields the same equilibrium outcome in which the left terminal node is reached. Hence equilibrium play is identical but off equilibrium pay differs. There are several SPE in this perfect information game because it is not generic.

5 Existence of SPE

The next theorem shows that the logic of backward induction can be extended to games with imperfect information.

Theorem 1 *Every finite extensive form game has a SPE.*

This theorem is the equivalent of the Nash existence theorem for extensive form games. It establishes that SPE is not too strong in the sense that a SPE exists for each extensive form game. We have seen that NE is too weak in extensive form games because there are too many equilibria.

The proof of the theorem is a generalization of backward induction. In backward induction we solve the game from the back by solving node after node. Now we solve it backwards subgame for subgame.

Formally, define the set Γ of subgames of the game G. Γ is never empty because G itself is a member. We can define a partial order on the set Γ such that for two subgames G_1 and G_2 we have $G_1 \geq G_2$ if G_2 is a subgame of G_1.

Look at the following example.

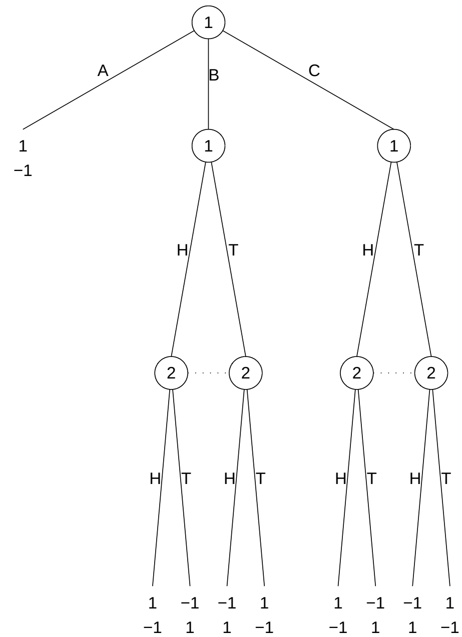

This game has three subgames: the whole game G_1 and two matching pennies subgames G_2 and G_3. We have $G_1 \geq G_2$ and $G_1 \geq G_3$ but G_2 and G_3 are not comparable.

Step I *Identify the terminal subgames.* Terminal subgames are those which do not dominate another subgame (G' is terminal if there is no G'' such that $G' > G''$.

Step II *Solve the terminal subgames* These subgames have no further subgames. They have a Nash equilibrium by the Nash existence result (they are finite!).

In our example the matching pennies subgames have the unique NE $\frac{1}{2}H + \frac{1}{2}T$ for each player.

Step III *Calculate the Nash payoffs of the terminal subgames and replace these subgames with the Nash payoffs.*

In our example the matching pennies payoffs are 0 for each player. We get:

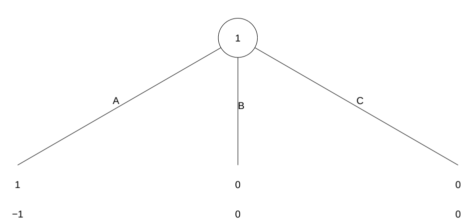

Step IV *Goto step I.* Repeat this procedure until all subgames have been exhausted. In this way we construct 'from the back' a SPE. In many cases the procedure does not produce a unique SPE if a subgame has multiple NE.

In our example we are lucky because matching pennies just has one NE. In the reduced game player 1 plays A which is his unique BR. The unique SPE is therefore $s_1^* = \left(A, \frac{1}{2}H + \frac{1}{2}T, \frac{1}{2}H + \frac{1}{2}T\right)$ and $s_2^* = \left(\frac{1}{2}H + \frac{1}{2}T, \frac{1}{2}H + \frac{1}{2}T\right)$.

6 Application of SPE to Behavioral Economics

In the first two lectures of the course we analyzed decision problems and later contrasted them to proper games where agents have to think strategically. Our decision problems were essentially static - people one action out of a number of alternatives.

In reality many decision problems involve taking decision over time: retirement decision, savings, when to finish a paper for class etc. are standard examples. Intertemporal decision making is different from static decision making because agents might want to revise past decisions in the future (they never have a chance to do so under static decision making). If agents revise decisions we say that they are *time-inconsistent*. In most economics classes you will never hear about such behavior - the decision making process of agents is assumed to be time-consistent. This reduces intertemporal decision making essentially to static decision making.

6.1 Time-Consistent Preferences

How do we model intertemporal decision making? Economists assume that the future counts less than the present - agents discount. Typically we assume that the utility in different time periods can be added up. So getting x_1 now and x_2 in the next period has total utility:

$$U = u(x_1) + \delta u(x_2) \tag{1}$$

For concreteness assume that agents want to spend 100 Dollars over two time periods. Their discount factor for next period's utility is 0.5. Their utility function is the square root function $u(x) = \sqrt{x}$. You can check that the agent would allocate 80 Dollar to today's consumption and the rest to tomorrow's consumption.

The agent is necessarily time-consistent in a two-period model because she cannot reallocate any resources in period 2. However this is no longer true in a 3-period model.

Let's first look at *exponential discounting* where consumption in period t is discounted by δ^t:

$$U_1 = u(x_1) + \delta u(x_2) + \delta^2 u(x_3) \tag{2}$$

Let's assume the same parameters as above. It's easy to check that the agent would want to allocate $\frac{2100}{16} \approx 76$ Dollars to the first period, 19 Dollars to the second and 5 to the third.

Now the agent could potentially change her allocation in period 2. Would she do so? The answer is no. Her decision problem in period 2 can be written as maximizing U_2 given that $x_2 + x_3 = 100 - 76$:

$$U_2 = u(x_2) + \delta u(x_3) \tag{3}$$

Note, that:

$$U_1 = u(x_1) + \delta U_2 \tag{4}$$

Therefore, if an agent would just a different consumption plan in period 2 she would have done so in period 1 as well.

We say that there is no conflict between different selves in games with exponential discounting. Agents are time-consistent. Time-consistent preferences are assumed in most of micro and macro economics.

6.2 Time-Inconsistent Preferences

Let's now look at a difference discounting rule for future consumption which we refer to as *hyperbolic discounting*. Agents discount at rate δ between all future time periods. However, they use an additional discount factor $\beta < 1$ to discount future versus present consumption. The idea here is, that consumers discount more strongly between period 1 and 2 than between period 2 and 3:

$$U = u(x_1) + \beta\delta u(x_2) + \beta\delta^2 u(x_3) \tag{5}$$

For simplicity we assume $\delta = 1$ from now on.

Let's assume $\beta = \frac{1}{2}$. In this case a period 1 agent would allocate 50 Dollars to today and 25 Dollars to both tomorrow and the day after tomorrow.

What would the period 2 agent do with her remaining 50 Dollars? Her decision problem would look as follows:

$$U = u(x_2) + \beta u(x_3) \tag{6}$$

So she would allocate 40 Dollars to period 2 and only 10 Dollars to the third period.

Therefore there is conflict between agent 1's and agent 2's preferences! Agent 1 would like agent 2 to save more, but agent 2 can't help herself and splurges!

6.3 Naive and Sophisticated Agents

There are two ways to deal with the self-control problem of agents. First, agents might not be aware of their future self's self-control problem - we say that they are *naive*. In this case you solve a different decision problem in each period and the consumption plans of agents get continuously revised.

If agents **are** aware of their self-control problem we call them *sophisticated*. Sophisticated agents play a game with their future self, are aware that they do so, and use SPE to solve for a consumption plan.

Let's return to our previous problem. A period 2 agent would always spend four times as much on this period than on period 3 (sophisticated or naive). Period 1 agent realizes this behavior of agent 2 and therefore takes the constraint $x_2 = 4x_3$ into account when allocating her own consumption. She maximizes:

$$U_1 = \sqrt{x_1} + \frac{1}{2}\left[\sqrt{\frac{4}{5}(1-x_1)} + \sqrt{\frac{1}{5}(1-x_1)}\right] \tag{7}$$

She would now spend 68 Dollars in the first period and predict that her future self spends 24 Dollars in the second period such that there are 6 left in the last period.

What has happened? Effectively, self 1 has taken away resources from self 2 by spending more in period 1. Self 1 predicts that self 2 would splurge - so self 1 might as well splurge immediately.

6.4 The Value of Commitment

If agents are time-inconsistent they can benefit from commitment devices which effectively constrain future selves. For example, a sophisticated hyperbolic agent could invest 50 Dollars in a 401k from which he can't withdraw more than 25 Dollars in the second period. While a time-consistent agent would never enter such a bargain (unexpected things might happen - so why should he constrain his future choices), a time-inconsistent agent might benefit.

7 Doing it Now or Later (Matt Rabin, AER, 1999)

This is a nice little paper which analyzes procrastination.

7.1 Salient Costs

Example 1 *Assume you go to the cinema on Saturdays. The schedule consists of a mediocre movie this week, a good movie next week, a great movie in two weeks and (best of all) a Johnny Depp movie in three weeks. Also assume that you must complete a report during the next four weeks so that you have to skip one of the movies. The benefit of writing the report is the same in each period (call it* $(\overline{v}, \overline{v}, \overline{v}, \overline{v})$*. The cost of not seeing a movie is* $(c_1, c_2, c_3, c_4) = (3, 5, 8, 13)$*. When do you write the report?*

Let's assume that there are three types of agents. Time consistent agents (TC) have $\delta = 1$ and $\beta = 1$. Naive agents have $\beta = \frac{1}{2}$ and sophisticated agents are aware of their self-control problem.

TC agents will write the report immediately and skip the mediocre movie. Generally, TC agents will maximize $\overline{v} - c$.

Naive agents will write the report in the last period. They believe that they will write the report in the second period. In the second period, they assume to write it in the third period (cost 4 versus 5 now). In the third period they again procrastinate.

Sophisticated agents use backward induction. They know that period 3 agent would procrastinate. Period 2 agent would predict period 3's procrastination and write the report. Period 1 agent knows that period 2 agent will write the report and can therefore safely procrastinate.

This example captures the idea that sophistication can somehow help to overcome procrastination because agents are aware of their future selfs tendencies to procrastinate.

7.2 Salient Rewards

Example 2 *Assume you can go to the cinema on Saturdays. The schedule consists of a mediocre movie this week, a good movie next week, a great movie in two weeks and (best of all) a Johnny Depp movie in three weeks. You have no money but a coupon to spend on exactly one movie. The benefit of seeing the movies are* $(v_1, v_2, v_3, v_4) = (3, 5, 8, 13)$. *Which movie do you see?*

The TC agent would wait and see the Depp movie which gives the highest benefit.

The naive agent would see the third movie. He would not the mediocre one because he would expect to see either the Depp movie later. He would also not see the week 2 movie for the same reason. But in week 3 he caves in to his impatience and spends the coupon.

The sophisticated agent would see the mediocre movie! She would expect that period 3 self caves in to her desires. Period 2 self would then go to the movies expecting period 3 self to cave in. But then period 1 self should go immediately because $3 > 2.5$.

The result is the opposite of the result we got for the procrastination example. Sophistication hurts agents! The intuition is that naive agents can pretend that they see the Depp movie and therefore are willing to wait. Sophisticated agents know about their weakness and therefore don't have the same time horizon. This makes them cave in even earlier.

A sophisticated agent would not be able to withstand a jar of cookies because she knows that she would cave in too early anyway, so she might as well cave in right away.

7.3 Choice and Procrastination: The Perfect as the Enemy of the Good

In a related paper (Choice and Procrastination, QJE 2002, forthcoming) Rabin points out that greater choice can lead to more procrastination. In the previous example, there was just procrastination on a single action. Now assume, that agents have the choice between two actions.

Example 3 *Assume you are a naive hyperbolic discounter* ($\beta = \frac{1}{2}$). *You can invest 1,000 Dollar in your 401k plan which gives you a yearly return of 5 percent. Once the money is invested it is out of reach for the next 30 years. Just when you want to sign the forms your friend tells you that he has invested* *his 1,000 Dollars at a rate of 6 percent. You make ca quick calculation and decide that you should do more research before signing on because research only causes you a disutility of 30 Dollars and the compounded interest gain over 30 years far exceed this amount. What will you do?*

You won't do anything and still have your 1000 Dollars after 30 years. Waiting a year has a cost of 50 Dollars (lost interest) which is discounted by $\beta = \frac{1}{2}$ and thus is below the salient cost of doing research. So you wait. and wait. and wait.

This is a nice example of why the perfect can be the enemy of the good: more choice can lead to procrastination. For a naive decision maker choices are determined by long-term benefits (just as for TC decision maker). However, procrastination is caused by small period to period costs of waiting which exceed salient costs of doing research.

Lecture XII: Analysis of Infinite Games

Markus M. Möbius

April 7, 2004

- Gibbons, chapter 2.1.A,2.1.B,2.2.A
- Osborne, sections 14.1-14.4, 16
- Oxborne and Rubinstein, sections 6.5, 8.1 and 8.2

1 Introduction - Critique of SPE

The SPE concept eliminates non-credible threats but it's worth to step back for am minute and ask whether we think SPE is reasonable or in throwing out threats we have been overzealous.

Practically, for this course the answer will be that SPE restrictions are OK and we'll always use them in extensive form games. However, it's worth looking at situations where it has been criticized. Some of the worst anomalies disappear in infinite horizon games which we study next.

1.1 Rationality off the Equilibrium Path

Is it reasonable to play NE off the equilibrium path? After all, if a player does not follow the equilibrium he is probably as stupid as a broomstick. Why should we trust him to play NE in the subgame? Let's look at the following game to illustrate that concern:

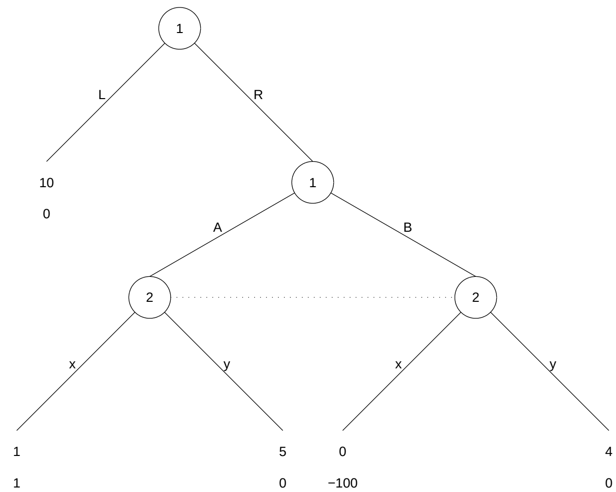

Here (L, A, x) is the unique SPE. However, player 2 has to put a lot of trust into player 1's rationality in order to play x. He must believe that player 1 is smart enough to figure out that A is a dominant strategy in the subgame following R. However, player 2 might have serious doubts about player 1's marbles after the guy has just foregone 5 utils by not playing L.[1]

1.2 Multi-Stage Games

Lemma 1 *The unique SPE of the finitely repeated Prisoner's Dilemma game in which players get the sum of their payoffs from each stage game has every player defecting at each information set.*

The proof proceeds by analyzing the last stage game where we would see defection for sure. But then we would see defection in the second to last stage game etc. In the same way we can show that the finitely repeated Bertrand game results in pricing at marginal cost all the time.

Remark 1 *The main reason for the breakdown of cooperation in the finitely repeated Prisoner's Dilemma is not so much SPE by itself by the fact that there is a final period in which agents would certainly defect. This raises the question whether an infinitely repeated PD game would allow us to cooperate. Essentially, we could cooperate as long as the other player does, and if there*

[1]After all any strategy in which L is played strictly dominates any strategy in which R is played in the normal form.

is defection, we defect from then on. This still looks like a SPE - in any subgame in which I have defected before, I might as well defect forever. If I haven't defected yet, I can jeopardize cooperation by defection, and therefore should not do it as long as I care about the future sufficiently.

These results should make us suspicious. Axelrod's experiments (see future lecture) showed that in the finitely repeated Prisoner's Dilemma people tend to cooperate until the last few periods when the 'endgame effect' kicks in. Similarly, there are indications that rival firms can learn to collude if they interact repeatedly and set prices above marginal cost.

This criticisms of SPE is reminiscent of our criticism of IDSDS. In both cases we use an iterative procedure to find equilibrium. We might have doubts where real-world subjects are able (and inclined) to do this calculation.

A famous example for the perils of backward induction is Rosenthal's centipede game:

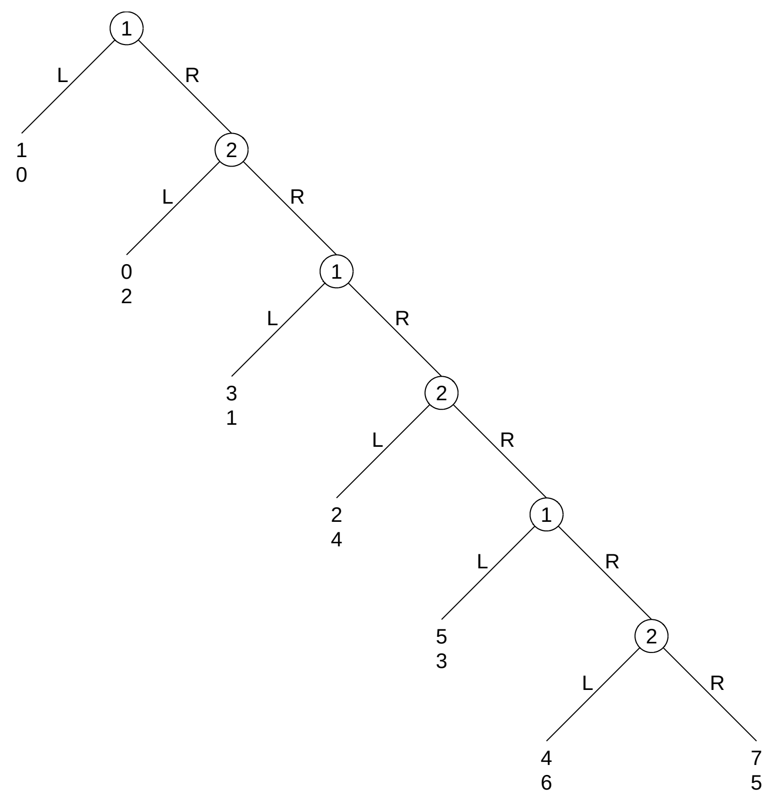

The game can be extended to even more periods. The unique SPE of the game is to drop out immediately (play L) at each stage. However, in experiments people typically continue until almost the last period before they drop out.

2 Infinite Horizon Games

Of the criticism of SPE the one we will take most seriously is that long finite horizon models do not give reasonable answers. Recall that the problem was that the backward induction procedure tended to unravel 'reasonably' looking strategies from the end. It turns out that many of the anomalies go away once we model these games as infinite games because there is not endgame to be played.

The prototypical model is what Fudenberg and Tirole call an infinite horizon multistage game with observed actions.

- At times $t = 0, 1, 2, ..$ some subset of the set of players simultaneously chooses actions.
- All players observe the period t actions before choosing period $t + 1$ actions.
- Players payoffs maybe any function of the infinite sequence of actions (play does not end in terminal nodes necessarily any longer)

2.1 Infinite Games with Discounting

Often we assume that player i's payoffs are of the form:

$$u_i(s_i, s_{-i}) = u_{i0}(s_i, s_{-i}) + \delta u_{i1}(s_i, s_{-i}) + \delta^2 u_{i2}(s_i, s_{-i}) + .. \qquad (1)$$

where $u_{it}(s_i, s_{-i})$ is a payoff received at t when the strategies are followed.

2.1.1 Interpretation of δ

1. Interest rate $\delta = \frac{1}{1+r}$. Having two dollars today or two dollars tomorrow makes a difference to you: your two dollars today are worth more, because you can take them to the bank and get $2(1+r)$ Dollars tomorrow where r is the interest rate. By discounting future payoffs with $\delta = \frac{1}{1+r}$ we correct for the fact that future payoffs are worth less to us than present payoffs.

2. Probabilistic end of game: suppose the game is really finite, but that the end of the game is not deterministic. Instead given that stage t is reached there is probability δ that the game continues and probability $1-\delta$ that the game ends after this stage. Note, that the expected number of periods is then $\frac{1}{1-\delta}$ and finite. However, we can't apply backward induction directly because we can never be sure that any round is the last one. The probabilistic interpretation is particularly attractive for interpreting bargaining games with many rounds.

2.2 Example I: Repeated Games

Let G be a simultaneous move game with finite action spaces A_1,..,A_I. The infinitely repeated game G^∞ is the game where in periods $t = 0, 1, 2, ..$ the players simultaneously choose actions $(a_1^t, .., a_I^t)$ after observing all previous actions. We define payoffs in this game by

$$u_i(s_i, s_{-i}) = \sum_{t=0}^{\infty} \delta^t \tilde{u}_i\left(a_1^t, .., a_I^t\right) \tag{2}$$

where $(a_1^t, .., a_I^t)$ is the action profile taken in period t when players follow strategies $s_1, .., s_I$, and $\tilde{u}_i$ are the utilities of players in each stage game. For example, in the infinitely repeated Prisoner's Dilemma game the $\tilde{u}_i$ are simply the payoffs in the 'boxes' of the normal form representation.

2.3 Example II: Bargaining

Suppose there is a one Dollar to be divided up between two players. The following alternate offer procedure is used:

I. In periods 0,2,4,... player 1 offers the division $(x_1, 1-x_1)$. Player 2 then accepts and the game ends, or he rejects and play continues.

II. In period 1,3,5,... player 2 offers the division $(1-x_2, x_2)$. Player 1 then accepts or rejects.

Assume that if the division $(y, 1-y)$ is agreed to in period t then the payoffs are $\delta^t y$ and $\delta^t(1-y)$.[2]

[2]Note that this is not a repeated game. First of all, the stage games are not identical (alternate players make offers). Second, there is no per period payoff. Instead, players only get payoffs when one of them has agreed to an offer. Waiting to divide the pie is costly.

Remark 2 *There are finite versions of this game in which play end after period T. One has to make some assumption what happens if there is no agreement at time T - typically, one assumes that the pie simply disappears. If $T = 1$ then we get the simple ultimatum game.*

3 Continuity at Infinity

None of the tools we've discussed so far are easy to apply for infinite games. First, backward induction isn't feasible because there is no end to work backward from. Second, using the definition of SPE alone isn't very easy. There are infinitely many subgames and uncountably many strategies that might do better.

We will discuss a theorem which makes the analysis quite tractable in most infinite horizon games. To do so, we must first discuss what continuity at infinity means.

Definition 1 *An infinite extensive form game G is continuous at ∞ if*

$$\lim_{T\to\infty} \sup_{i,\sigma,\sigma' \text{ s. th. } \sigma(h) = \sigma'(h) \text{ for all } h \text{ in periods } t \leq T} |u_i(\sigma) - u_i(\sigma')| = 0$$

In words: compare the payoffs of two strategies which are identical for all information sets up to time T and might differ thereafter. As T becomes large the maximal difference between any two such strategies becomes arbitrarily small. Essentially, this means that distant future events have a very small payoff effect.

3.1 Example I: Repeated Games

If σ and σ' agree in the first T periods then:

$$\begin{aligned} |u_i(\sigma) - u_i(\sigma')| &= \left| \sum_{t=T}^{\infty} \delta^t \left(\tilde{u}_i\left(a^t\right) - \tilde{u}_i\left(a'_t\right)\right) \right| \\ &\leq \sum_{t=T}^{\infty} \delta^t \max_{i,a,a'} |\tilde{u}_i(a) - \tilde{u}_i(a')| \end{aligned}$$

For finite stage games we know that $M = \max_{i,a,a'} |\tilde{u}_i(a) - \tilde{u}_i(a')|$ is finite. This implies that

$$\lim_{T\to\infty} |u_i(\sigma) - u_i(\sigma')| \leq \lim_{T\to\infty} \sum_{t=T}^{\infty} \delta^t M = \lim_{T\to\infty} \frac{\delta^T}{1-\delta} M = 0.$$

3.2 Example II: Bargaining

It's easy to check that the bargaining game is also continuous.

3.3 Example III: Non-discounted war of attrition

This is an example for an infinite game which is NOT continuous at infinity. Players 1 and 2 choose $a_i^t \in \{\text{Out}, \text{Fight}\}$ at time $t = 0, 1, 2, ...$ The game ends whenever one player quits with the other being the 'winner'. Assume the payoffs are

$$u_i(s_i, s_{-i}) = \begin{cases} 1 - ct & \text{if player } i \text{ 'wins' in period } t \\ -ct & \text{if player } i \text{ quits in period } t \end{cases}$$

Note, that players in this game can win a price of 1 by staying in the game longer than the other player. However, staying in the game is costly for both players. Each player wants the game to finish as quickly as possible, but also wants the other player to drop out first.

This game is not continuous at ∞. Let

$$\begin{aligned} \sigma^T &= \text{Both fight for } T \text{ periods and then 1 quits} \\ \sigma'^T &= \text{Both fight for } T \text{ periods and then 2 quits.} \end{aligned}$$

Then we have

$$\left| u_i\left(\sigma^T\right) - u_i\left(\sigma'^T\right) \right| = 1.$$

This expression does not go to zero as $T \to \infty$.

4 The Single-Period Deviation Principle

The next theorem makes the analysis of infinite games which are continuous at ∞ possible.

Theorem 1 *Let G be an infinite horizon multistage game with observed actions which is continuous at ∞. A strategy profile $(\sigma_1, .., \sigma_I)$ is a SPE if and only if there is no player i and strategy $\hat{\sigma}_i$ that agrees with σ_i except at a single information set h_i^t and which gives a higher payoff to player i conditional on h_i^t being reached.*

We write $u_i(\sigma_i, \sigma_{-i}|x)$ for the payoff conditional on x being reached. For example, in the entry game below we have $u_2(\text{Accomodate}, \text{Out}|\text{node 1 is reached}) = 1$.

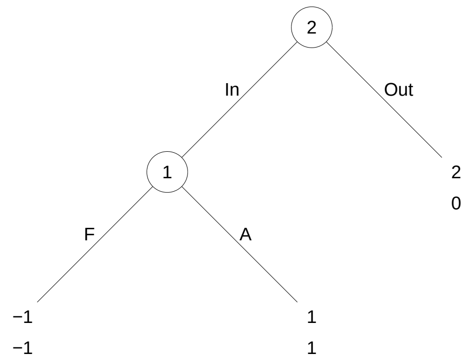

Recall, that we can condition on nodes which are not on the equilibrium path because the strategy of each player defines play at each node.

4.1 Proof of SPDP for Finite Games

I start by proving the result for finite-horizon games with observed actions.

Step I: By the definition of SPE there cannot be a profitable deviation for any player at some information set in games with observed actions.[3]

Step II: The reverse is a bit harder to show. We want to show that $u_i(\hat{\sigma}_i, \sigma_{-i}|x_t) \leq u_i(\sigma_i, \sigma_{-i}|x_t)$ for all initial nodes x_t of a subgame (subgame at some round t).

We prove this by induction on T which is the number of periods in which σ_i and $\hat{\sigma}_i$ differ.

[3]We have to very careful at this point. We have defined SPE as NE in every subgame. Subgames can only originate at nodes and not information sets. However, in games with observed actions all players play simultaneous move games in each round t. Therefore any deviation by a player at an information set at round t which is not a singleton is on the equilibrium path of some subgame at round t.

T=1: In this case the result is clear. Suppose σ_i and $\hat{\sigma}_i$ differ only in the information set in period t'. If $t > t'$ it is clear that $u_i(\hat{\sigma}_i, \sigma_{-i}|x_t) = u_i(\sigma_i, \sigma_{-i}|x_t)$ because the two strategies are identical at all the relevant information sets. If $t \leq t'$ then:

$$\begin{aligned} u_i(\hat{\sigma}_i, \sigma_{-i}|x_t) &= \sum_{h_{it'}} u_i(\hat{\sigma}_i, \sigma_{-i}|h_{it'})\, Prob\{h_{it'}|\hat{\sigma}_i, \sigma_{-i}, x_t\} \\ &\leq \sum_{h_{it'}} \underbrace{u_i(\sigma_i, \sigma_{-i}|h_{it'})}_{\text{follows from one stage deviation criterion}}\ \underbrace{Prob\{h_{it'}|\sigma_i, \sigma_{-i}, x_t\}}_{\text{follows from } \sigma_i \text{ and } \hat{\sigma}_i \text{ having same play between } t \text{ and } t'} \\ &= u_i(\sigma_i, \sigma_{-i}|x_t) \end{aligned}$$

T → T+1: Assuming that the result holds for T let $\hat{\sigma}_i$ be any strategy differing from σ_i in $T+1$ periods. Let t' be the *last* period at which they differ and define $\tilde{\sigma}_i$ by:

$$\tilde{\sigma}_i(h_{it}) = \begin{cases} \hat{\sigma}_i(h_{it}) & \text{if } t < t' \\ \sigma_i(h_{it}) & \text{if } t \geq t' \end{cases}$$

In other words, $\tilde{\sigma}_i$ differs from σ_i only at T periods. Therefore we have for any x_t

$$u_i(\tilde{\sigma}_i, \sigma_{-i}|x_t) \leq u_i(\sigma_i, \sigma_{-i}|x_t)$$

by the inductive hypothesis since we assumed that the claim holds for T.

However, we also know that $\tilde{\sigma}$ and $\hat{\sigma}$ only differ in a single deviation at round t'. Therefore, we can use exactly the same argument as in the previous step to show that

$$u_i(\hat{\sigma}_i, \sigma_{-i}|x_t) \leq u_i(\tilde{\sigma}_i, \sigma_{-i}|x_t)$$

for any x_t.[4]

Combining both inequalities we get the desires result:

$$u_i(\hat{\sigma}_i, \sigma_{-i}|x_t) \leq u_i(\sigma_i, \sigma_{-i}|x_t)$$

This proves the result for finite games with observed actions.

[4] It is important that we have defined $\tilde{\sigma}_i$ in differing only in the last period deviation. Therefore, after time t' strategy $\tilde{\sigma}_i$ follows σ_i. This allows us to use the SPDP.

4.2 Proof for infinite horizon games

Note, that the proof for finite-horizon games also establishes that we for a profile σ which satisfies SPDP player i cannot improve on σ_i in some subgame x_t by considering a new strategy $\hat{\sigma}_i$ with finitely many deviations from σ_i. However, it is still possible that deviations at infinitely many periods might be an improvement for player i.

Assume this would be the case for some $\hat{\sigma}_i$. Let's denote the gain from using this strategy with ϵ:

$$\epsilon = u_i\left(\hat{\sigma}_i, \sigma_{-i}|x_t\right) - u_i\left(\sigma_i, \sigma_{-i}|x_t\right)$$

Because the game is continuous at ∞ this implies that if we choose T large enough we can define some strategy $\tilde{\sigma}_i$ which agrees with $\hat{\sigma}_i$ up to period T and then follows strategy σ_i such that:

$$\left|u_i\left(\hat{\sigma}_i, \sigma_{-i}|x_t\right) - u_i\left(\tilde{\sigma}_i, \sigma_{-i}x_t\right)\right| < \frac{\epsilon}{2}$$

This implies that the new strategy $\tilde{\sigma}_i$ gives player i strictly more utility than σ_i. However, it can only differ from σ_i for at most T periods. But this is a contradiction as we have shown above. QED

Remark 3 *Games which are at continuous at ∞ are in some sense 'the next best thing' to finite games.*

5 Analysis of Rubinstein's Bargaining Game

To illustrate the use of the single period deviation principle and to show the power of SPE in one interesting model we now return to the Rubinstein bargaining game introduced before.

First, note that the game has many Nash equilibria. For example, player 2 can implement any division by adopting a strategy in which he only accepts and proposes a share x_2 and rejects anything else.

Proposition 1 *The bargaining game has a unique SPE. In each period of the SPE the player i who proposes picks $x_i = \frac{1}{1+\delta}$ and the other player accepts any division giving him at least $\frac{\delta}{1+\delta}$ and rejects any offer giving him less.*

Several observations are in order:

1. We get a roughly even split for δ close to 1 (little discounting). The proposer can increase his share the more impatient the other player is.

2. Agreement is immediate and bargaining is therefore efficient. Players can perfectly predict play in the next period and will therefore choose a division which makes the other player just indifferent between accepting and making her own proposal. There is no reason to delay agreement because it just shrinks the pie. Immediate agreement is in fact not observed in most experiments - last year in a two stage bargaining game in class we observed that only 2 out of 14 bargains ended in agreement after the first period. There are extensions of the Rubinstein model which do not give immediate agreement.[5]

3. The division becomes less equal for finite bargaining games. Essentially, the last proposer at period T can take everything for himself. Therefore, he will tend to get the greatest share of the pie in period 1 as well - otherwise he would continue to reject and take everything in the last period. In our two-period experiments we have in deed observed greater payoffs to the last proposer (ratio of 2 to 1). However, half of all bargains resulted in disagreement after the second period and so zero payoffs for everyone. Apparently, people care about fairness as well as payoffs which makes one wonder whether monetary payoffs are the right way to describe the utility of players in this game.

5.1 Useful Shortcuts

The hard part is to show that the Rubinstein game has a unique SPE. If we know that it is much easier to calculate the actual strategies.

5.1.1 Backward Induction in infinite game

You can solve the game by assuming that random payoff after rejection of offer at period T. The game then becomes a finite game of perfect information which can be solved through backward induction. It turns out that as $T \to \infty$

[5]For example, when there is uncertainty about a players discount factor the proposer might start with a low offer in order to weed out player 2 types with low discount factor. Players with high δ will reject low offers and therefore agreement is not immediate. To analyze these extension we have to first develop the notion of incomplete information games.

the backward solution converges to the Rubinstein solution. This technique also provides an alternative proof for uniqueness.

5.1.2 Using the Recursive Structure of Game

The game has a recursive structure. Each player faces essentially the same game, just with interchanged roles. Therefore, in the unique SPE player 1 should proposes some $(x, 1-x)$ and player 2 should propose $(1-x, x)$. Player 1's proposal to 2 should make 2 indifferent between accepting immediately or waiting to make her own offer (otherwise player 1 would bid higher or lower). This implies:

$$\underbrace{1-x}_{\text{2's payoff at period 1}} = \underbrace{\delta x}_{\text{2's discounted payoff from period 2}}$$

$$x = \frac{1}{1+\delta}$$

5.2 Proof of Rubinstein's Solution

5.2.1 Existence

We show that there is no profitable single history deviation.

Proposer: If he conforms at period t the continuation payoff is $\delta^t \frac{1}{1+\delta}$. If he deviates and asks for more he gets $\delta^{t+1} \frac{\delta}{1+\delta}$. If he deviates and asks for less he gets less. Either way he loses.

Recipient: Look at payoffs conditional on y being proposes in period t.

- If he rejects he gets $\delta^{t+1} \frac{1}{1+\delta}$.
- If he accepts he gets $\delta^t y$.
- If $y \geq \frac{\delta}{1+\delta}$ the strategy says accept and this is better than rejecting.
- If $y < \frac{\delta}{1+\delta}$ the strategy says reject and accepting is not a profitable deviation.

5.2.2 Uniqueness

This proof illustrates a number of techniques which enable us to prove properties about equilibrium without actually constructing it.

Let $\overline{v}$ and $\underline{v}$ be the highest and lowest payoffs received by a proposer in any SPE. We first observe:

$$1-\overline{v} \geq \delta \underline{v} \tag{3}$$

If this equation would not be true then no proposer could propose $\overline{v}$ because the recipient could always get more in any subgame. We also find:

$$\underline{v} \geq 1-\delta\overline{v} \tag{4}$$

If not then no proposer would propose $\underline{v}$ - she would rather wait for the other player to make her proposal because she would get a higher payoff this way.

We can use both inequalities to derive bounds on $\underline{v}$ and $\overline{v}$:

$$\underline{v} \geq 1-\delta\overline{v} \geq 1-\delta\left(1-\delta\underline{v}\right) \tag{5}$$

$$\left(1-\delta^2\right)\underline{v} \geq 1-\delta \tag{6}$$

$$\underline{v} \geq \frac{1}{1+\delta} \tag{7}$$

Similarly, we find:

$$1-\overline{v} \geq \delta\underline{v} \geq \delta\left(1-\delta\overline{v}\right) \tag{8}$$

$$\overline{v} \leq \frac{1}{1+\delta} \tag{9}$$

Hence $\overline{v} = \underline{v} = \frac{1}{1+\delta}$. Clearly, no other strategies can generate this payoff in every subgame.[6]

[6] While being a nice result it does not necessarily hold anymore when we change the game. For example, if both players make simultaneous proposals then any division is a SPE. Also, it no longer holds when there are several players.

Lecture XIII: Repeated Games

Markus M. Möbius

April 19, 2004

- Gibbons, chapter 2.3.B,2.3.C
- Osborne, chapter 14
- Osborne and Rubinstein, sections 8.3-8.5

1 Introduction

So far one might get a somewhat misleading impression about SPE. When we first introduced dynamic games we noted that they often have a large number of (unreasonable) Nash equilibria. In the models we've looked at so far SPE has 'solved' this problem and given us a unique NE. In fact, this is not really the norm. We'll see today that many dynamic games still have a very large number of SPE.

2 Credible Threats

We introduced SPE to rule out non-credible threats. In many finite horizon games though credible threats are common and cause a multiplicity of SPE.

Consider the following game:

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

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

Suppose that the players play the game twice and observe first period actions before choosing the second period actions. Now one way to get a SPE is to play any of the three profiles above followed by another of them (or same one). We can also, however, use credible threats to get other actions played in period 1, such as:

- Play (B,R) in period 1.
- If player 1 plays B in period 1 play (T,L) in period 2 - otherwise play (M,C) in period 2.

It is easy to see that no single period deviation helps here. In period 2 a NE is played so obviously doesn't help.

- In period 1 player 1 gets $4 + 3$ if he follows strategy and at most $5 + 1$ if he doesn't.
- Player 2 gets $4 + 1$ if he follows and at most 2+1 if he doesn't.

Therefore switching to the (M,C) equilibrium in period 2 is a *credible threat.*

3 Repeated Prisoner's Dilemma

Note, that the PD doesn't have multiple NE so in a finite horizon we don't have the same easy threats to use. Therefore, the finitely repeated PD has a unique SPE in which every player defects in each period.

	C	D
C	1,1	−1,2
D	2,−1	0,0

In infinite horizon, however, we do get many SPE because other types of threats are credible.

Proposition 1 *In the infinitely repeated PD with $\delta \geq \frac{1}{2}$ there exists a SPE in which the outcome is that both players cooperate in every period.*

Proof: Consider the following symmetric profile:

$$s_i(h_t) = \begin{cases} C & \text{If both players have played C in every period along the path leading to } h_t. \\ D & \text{If either player has ever played D.} \end{cases}$$

To see that there is no profitable single deviation note that at any h_t such that $s_i(h_t) = D$ player i gets:

$$0 + \delta 0 + \delta^2 0 + ..$$

if he follows his strategy and

$$-1 + \delta 0 + \delta^2 0 + ..$$

if he plays C instead and then follows s_i.

At any h_t such that $s_i(h_t) = C$ player i gets:

$$1 + \delta 1 + \delta^2 1 + .. = \frac{1}{1-\delta}$$

if he follows his strategy and

$$2 + \delta 0 + \delta^2 0 + .. = 2$$

if he plays D instead and then follows s_i.

Neither of these deviations is worth while if $\delta \geq \frac{1}{2}$. QED

Remark 1 *While people sometimes tend to think of this as showing that people will cooperate in they repeatedly interact if does not show this. All it shows is that there is one SPE in which they do. The correct moral to draw is that there many possible outcomes.*

3.1 Other SPE of repeated PD

1. For any δ it is a SPE to play D every period.

2. For $\delta \geq \frac{1}{2}$ there is a SPE where the players play D in the first period and then C in all future periods.

3. For $\delta > \frac{1}{\sqrt{2}}$ there is a SPE where the players play D in every even period and C in every odd period.

4. For $\delta \geq \frac{1}{2}$ there is a SPE where the players play (C,D) in every even period and (D,C) in every odd period.

3.2 Recipe for Checking for SPE

Whenever you are supposed to check that a strategy profile is an SPE you should first try to classify all histories (i.e. all information sets) on and off the equilibrium path. Then you have to apply the SPDP for each class separately.

- Assume you want to check that the cooperation with grim trigger punishment is SPE. There are two types of histories you have to check. Along the equilibrium path there is just one history: everybody cooperated so far. Off the equilibrium path, there is again only one class: one person has defected.

- Assume you want to check that cooperating in even periods and defecting in odd periods plus grim trigger punishment in case of deviation by any player from above pattern is SPE. There are three types of histories: even and odd periods along the equilibrium path, and off the equilibrium path histories.

- Assume you want to check that TFT ('Tit for Tat') is SPE (which it isn't - see next lecture). Then you have you check four histories: only the play of both players in the last period matters for future play - so there are four relevant histories such as player 1 and 2 both cooperated in the last period, player 1 defected and player 2 cooperated etc.1

Sometimes the following result comes in handy.

Lemma 1 *If players play Nash equilibria of the stage game in each period in such a way that the particular equilibrium being played in a period is a function of time only and does not depend on previous play, then this strategy is a Nash equilibrium.*

The proof is immediate: we check for the SPDP. Assume that there is a profitable deviation. Such a deviation will not affect future play by assumption: if the stage game has two NE, for example, and NE1 is played in even periods and NE2 in odd periods, then a deviation will not affect future play.[1] Therefore, the deviation has to be profitable in the current stage game - but since a NE is being played no such profitable deviation can exist.

Corollary 1 *A strategy which has players play the same NE in each period is always SPE.*

In particular, the grim trigger strategy is SPE if the punishment strategy in each stage game is a NE (as is the case in the PD).

4 Folk Theorem

The examples in 3.1 suggest that the repeated PD has a tremendous number of equilibria when δ is large. Essentially, this means that game theory tells us we can't really tell what is going to happen. This turns out to be an accurate description of most infinitely repeated games.

[1] If a deviation triggers a switch to only NE1 this statement would no longer be true.

Let G be a simultaneous move game with action sets A_1, A_2, .., A_I and mixed strategy spaces Σ_1, Σ_2,.., Σ_I and payoff function $\tilde{u}_i$.

Definition 1 *A payoff vector $v = (v_1, v_2, .., v_I) \subset \Re^I$ is feasible if there exists action profiles a^1, a^2, ..,$a^k \in A$ and non-negative weights ω_1,..,ω_I which sum up to* 1 *such that*

$$v_i = \omega_1 \tilde{u}_i \left(a^1\right) + \omega_2 \tilde{u}_i \left(a^2\right) + .. + \omega_k \tilde{u}_i \left(a^k\right) +$$

Definition 2 *A payoff vector v is strictly individually rational if*

$$v_i > \underline{v}_i = \min_{\sigma_{-i} \in \Sigma_{-i}} \max_{\sigma_i(\sigma_{-i}) \in \Sigma_i} \tilde{u}_i \left(\sigma_i \left(\sigma_{-i}\right), \sigma_{-i}\right) \tag{1}$$

We can think of this as the lowest payoff a rational player could ever get in equilibrium if he anticipates his opponents' (possibly non-rational) play.

Intuitively, the minmax payoff $\underline{v}_i$ is the payoff player 1 can guarantee to herself even if the other players try to punish her as badly as they can. The minmax payoff is a measure of the punishment other players can inflict.

Theorem 1 ***Folk Theorem.*** *Suppose that the set of feasible payoffs of G is I-dimensional. Then for any feasible and strictly individually rational payoff vector v there exists $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ there exists a SPE x^* of G^∞ such that the average payoff to s^* is v, i.e.*

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

The normalized (or average) payoff is defined as $P = (1 - \delta) u_i (s^*)$. It is the payoff which a stage game would have to generate in each period such that we are indifferent between that payoff stream and the one generates by s^*:

$$P + \delta P + \delta^2 P + ... = u_i \left(s^*\right)$$

4.1 Example: Prisoner's Dilemma

- The feasible payoff set is the diamond bounded by (0,0), (2,-1), (-1,2) and (1,1). Every point inside can be generated as a convex combinations of these payoff vectors.

- The minmax payoff for each player is 0 as you can easily check. The other player can at most punish his rival by defecting, and each player can secure herself 0 in this case.

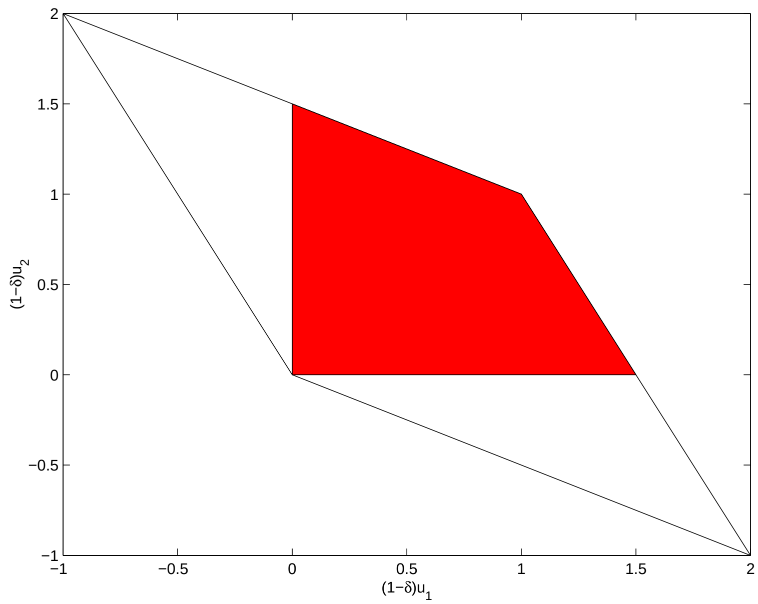

Hence the theorem says that anything in this trapezoid is possible. Note, that the equilibria I showed before generate payoffs inside this area.

4.2 Example: BOS

Consider the Battle of the Sexes game instead.

	F	O
F	2,1	0,0
O	0,0	1,2

Here each player can guarantee herself at least payoff $\frac{2}{3}$ which is the payoff from playing the mixed strategy Nash equilibrium. You can check that whenever player 2 mixes with different probabilities, player 1 can guarantee herself more than this payoff by playing either F or O all the time.

4.3 Idea behind the Proof

1. Have players on the equilibrium path play an action with payoff v (or alternate if necessary to generate this payoff).[2]

2. If some player deviates punish him by having the other players for T periods choose σ_{-i} such that player i gets $\underline{v}_i$.

3. After the end of the punishment phase reward all players (other than i) for having carried out the punishment by switching to an action profile where player i gets $v_i^P < v_i$ and all other players get $v_j^P + \epsilon$.

[2]For example, in the BoS it is not possible to generate $\left(\frac{3}{2}, \frac{3}{2}\right)$ in the stage game even with mixing. However, if players alternate and play (O,F) and then (F,O) the players can get arbitrarily close for large δ.

Lecture XIV: Applications of Repeated Games

Markus M. Möbius

April 28, 2004

- Gibbons, chapter 2.3.D,2.3.E
- Osborne, chapter 14

1 Introduction

We have quite thoroughly discussed the theory of repeated games. In this lecture we discuss applications. The selection of problems is quite eclectic and include:

- collusion of firms
- efficiency wages
- monetary policy
- theory of gift giving

For all these applications we analyze equilibria which are similar to the basic grim-trigger strategy equilibrium which we first studied in the repeated Prisoner's Dilemma context.

2 Collusion of Firms

2.1 Some Notes on Industrial Organization

Industrial organization is a subfield of economics which concerns itself with the study of particular industries. The standard setting is one where there are many consumers but only a few firms (usually more than one). This environment is called an *oligopoly* (in contrast to a monopoly).

Classical economics usually assumes perfect competition where there are many firms and consumers and all of them are price-takers. This simplifies analysis a lot because it gets rid of strategic interactions - a firm which is small relative to the market and hence has an infinitesimal impact on price does not have to worry about the reactions of other firms to its own actions.

Oligopolies are more intriguing environments because strategic interactions do matter now. Therefore, game theory has been applied extremely successfully in IO during the past 30 years.

2.2 The Bertrand Paradox

We have already discussed static Bertrand duopoly where firms set prices equal to marginal cost (as long as they have symmetric costs). Many economists think that this result is counterintuitive - if it would literally hold then firms would not be able to recoup any fixed costs (such as R&D, building a factory etc.) and would not develop products in the first place.

One solution is to assume that firms engage instead in Cournot competition. However, the assumption that firms set prices rather than quantities is appropriate for industries without significant capacity constraints (such as joghurts versus airplanes).

Repeated games provide an alternative resolution of the Bertrand paradox: firms can 'cooperate' and set prices above marginal cost in each period. If a firm defects they both revert to static Nash pricing at marginal cost. This equilibrium is an example of *tacit collusion*.[1]

2.3 Static Bertrand

Let's recall the static Bertrand game:

- Two firms have marginal cost of production $c > 0$.
- They face a downward sloping demand curve $q = D(p)$.
- Firms can set any price $p \geq 0$.

The unique NE of the game is $p_1 = p_2 = c$.

[1] It's called tacit collusion because it is a Nash equilibrium and self-enforcing.

Remark 1 *The finitely repeated Bertrand game still has a unique SPE in which firms set price equal to marginal cost in each period. This follows from the fact the NE of the static game is unique - hence there are no credible threats to enforce any other equilibrium.*

2.4 Infinitely Repeated Bertrand game

Assume firms have a discount factor δ.[2] Assume that the monopoly price is $p^m > 0$ and the monopoly profit is Π^M. The monopoly profit is the maximum of a single firm's profit function:

$$\Pi^M = \max_{p \geq 0} (p - c) D(p) \tag{1}$$

The monopoly price p^m is the maximizer of this function. Finally, we assume that if two firms set equal price they divide sales equally amongst them.

Then the following will be a SPE which ensures that both firms together will make the same profit as a monopolist (which is the largest profit attainable in the industry and hence the 'best possible' equilibrium for the firms).

1. Each firm sets its price equal to p^m in period $t = 1$ and each subsequent period as long as no firm has deviated from this price in previous periods.

2. After a deviation firm sets price equal to c.

This is essentially another version of the familiar grim-trigger strategy equilibrium. There are two types of histories which we have to check - those on the equilibrium path and those off. As usual we use the SPDP.

1. The best deviation on the equilibrium path is to undercut the other firm slightly and capture the entire market for one period. Afterwards, firms set price equal to marginal cost and hence make zero profit. The payoff of this deviation is Π^M. The profit from following the equilibrium strategy instead is:

$$\frac{\Pi^M}{2} + \delta \frac{\Pi^M}{2} + \delta^2 \frac{\Pi^M}{2} + \delta^3 \frac{\Pi^M}{2} + .. = \frac{\frac{\Pi^M}{2}}{1 - \delta} \tag{2}$$

[2]Firms usually discount by the interest rate of the economy such that $\delta = \frac{1}{1+r}$.

For the original strategy to be an SPE the deviation has to be non-profitable:

$$\Pi^M \leq \frac{\frac{\Pi^M}{2}}{1-\delta} \tag{3}$$

This will be satisfied as long as $\delta \geq \frac{1}{2}$.

2. Off the equilibrium path there is no profitable deviation. If a firm sets price above c it will still make zero profit. If it sets price below c it will make negative profits because it will sell at a loss. Note, that as usual checking for SPDP in the trigger phase is easy.

2.5 Rotemberg-Saloner counter-cyclical Pricing

Rotemberg and Saloner modified the basic repeated Bertrand model to allow for booms and recessions. The setup is the following:

- In each period demand can be either high or low. Firms know whether demand is high or low. The profit-maximizing prices in both states are p_L^m and p_H^m and corresponding profits are $\Pi_L^M < \Pi_H^M$.
- The 'best' equilibrium would be again one in which firms set the monopoly price in each period - they would then imitate a monopoly firm. However, we now have to check two types of histories on the equilibrium path - those in which demand is high and those where it is low.
- In the high-demand information set the following has to hold:

$$\Pi_H^M \leq \frac{\Pi_H^M}{2} + \frac{\delta}{1-\delta} \underbrace{\left[p\frac{\Pi_H^M}{2} + (1-p)\frac{\Pi_L^M}{2}\right]}_{\text{average profit per period}} \tag{4}$$

In the low-demand information this condition holds:

$$\Pi_L^M \leq \frac{\Pi_L^M}{2} + \frac{\delta}{1-\delta} \underbrace{\left[p\frac{\Pi_H^M}{2} + (1-p)\frac{\Pi_L^M}{2}\right]}_{\text{average profit per period}} \tag{5}$$

- It's easy to see that only the first condition implies the second one: the incentive to deviate and undercut the rival firm is strongest in high-demand. From the first condition we can calculate the cutoff value δ^* such that both inequalities hold for any $\delta > \delta^*$:

$$\Pi_H^M = \frac{\delta^*}{1-\delta^*}\left[p\Pi_H^M + (1-p)\Pi_L^M\right] \tag{6}$$

- It is also easy to check that the cutoff value δ^* is greater than $\frac{1}{2}$ - in other words, collusion is harder than in a world where there are no booms and can only be sustained for $\delta > \delta^* > \frac{1}{2}$.

- What if $\delta^* > \delta > \frac{1}{2}$?

For these intermediate values of δ some form of collusion can still be maintained as long as firms set prices below the monopoly price in the high-demand state. Intuitively, this reduces the incentive to defect in the high-demand state.

The easiest way to see this is to assume that firms reduce price in the high-demand state to such an extent that they make the same profit as in the low-demand state. In that case we are back to the previous equilibrium which can be sustained for $\delta > \frac{1}{2}$.

However, firms will be able to sustain higher profits than Π_L^M in the high-demand state for $\delta^* > \delta > \frac{1}{2}$. The profit Π_H^δ which can (just) be sustained in the high-demand state satisfies:

$$\Pi_H^\delta = \frac{\delta}{1-\delta}\left[p\Pi_H^\delta + (1-p)\Pi_L^M\right] \tag{7}$$

Note that $\Pi_L^\delta \to \Pi_L^M$ as $\delta \to \frac{1}{2}$.

Remark 2 *One can show that no collusion can be sustained for $\delta < \frac{1}{2}$.*

Remark 3 *This model implies counter-cyclical pricing - firms have to cut prices in booms below the monopoly price in order to prevent collusion from breaking down (by making defection too attractive). Note, that this does not literally mean that prices decrease in booms - it just means that the profit markup is less than the markup during recessions. Also note, that with perfect Bertrand competition in each period we would see prices equal to marginal costs in both booms and recessions.*

3 Efficiency Wages (Stiglitz)

Stiglitz studied the question why firms pay workers often more than have to pay in order to prevent the workers from leaving the firm.

3.1 One-Stage Shirking Model

- A firm and a worker play a two period game.
- In the first period the firm sets a wage w.
- In the second period the worker observes the wage and decides whether to accept or reject the job. If she rejects she has an outside option w_0. If she accepts she can exert effort or exert no effort. If she exerts no effort she will produce output $y > 0$ with probability p and 0 otherwise. If she exerts effort she will produce y for sure (this implies that output is 0 is a sure sign of shirking). Exerting effort has cost e to the firm.
- Clearly, the firm cannot enforce effort - there will be shirking all the time.
- The firm only has to pay the worker w_0 in order to employ the worker - paying a higher wage makes no sense because the worker shirks anyway.

Assumption 1 *We assume $y - e > w_0 > py$. This makes exerting effort socially efficient and shirking less efficient than the outside option. It also means, that the firm would not employ the worker in the first place because the minimum wage would exceed the expected output of the worker!*

3.2 Repeated Interactions in Stiglitz' Model

How can the firm prevent the worker from shirking? The following grim-trigger type strategy accomplishes the trick:

1. First of all, the firm has to pay the worker a higher wage $w^* > w_0$ - otherwise the worker has nothing to lose from shirking.
2. Second, the firm has to fire the worker is it detects shirking. Since the worker is paid above his outside he will perceive this as the threat of punishment.

This is why this model can explain efficiency wages (above a worker's outside option).

We use SPDP to check that this is indeed an SPE for sufficiently patient workers and firms.

A worker who exerts effort gets surplus $V^e = \frac{w^*-e}{1-\delta}$ (discounted over time) where w^* is his per period wage. If the worker shirks for one period he gets the following payoff V^s:

$$V^s = w^* + \delta \left[pV_e + (1-p)\frac{w_0}{1-\delta} \right] \quad (8)$$

In equilibrium we need $V^e \geq V^s$. This implies:

$$w^* \geq w_0 + \frac{1-\delta p}{\delta(1-p)} e \quad (9)$$

The best strategy of the firm is clearly to set w^* equal to this cutoff value.

Remark 4 *The markup $\frac{1-\delta p}{\delta(1-p)}e$ is the efficiency premium which the worker gets to make detection costly to him (because he loses the excess rents he enjoys in his job).*

4 Barro-Gordon Model of Monetary Policy

This model nicely captures how short-sighted policy makers succumb to the sweet poison of inflation to kick-start an ailing economy. The repeated version of the Barro-Gordon model illustrates the usefulness of the repeated game paradigm in applied economics.

4.1 The Static Barro-Gordon Model of Inflation

- There are two periods - in the first period firms choose their expectations π^e of inflation. In the second period the central bank observed π^e and chooses actual inflation π. The timing reflects the idea that the central bank observes more and better data than firms.

- We think of firms as a single player (say because they communicate with each other and form expectations collectively). Firms maximize their utility function $-(\pi - \pi^e)^2$ and hence optimally choose expected inflation to be actual inflation π. The reason is simple: firms base their pricing and wages on expected inflation and they are locked into these decision for a certain period of time (this is intuitive with wages since firms negotiate only periodically with unions and workers). If inflation is higher their prices tend to be too low. If inflation is lower than expected then their prices can be too high.

- The central bank faces the following utility function W:

$$W = -c\pi^2 - (y - y^*)^2 \tag{10}$$

Here, y is the output/GDP of the economy and y^* is the full employment output. Policy makers prefer that $\pi = 0$ and $y = y^*$. However, they can only control inflation (through interest rates) and they face the following inflation/output tradeoff:

$$y = by^* + d(\pi - \pi^e) \tag{11}$$

where $b < 1$: the idea is that because of various frictions in the economy (monopoly power etc.) the economy operates below full employment (since $b < 1$). Policy makers can boost GDP by creating *surprise inflation* - i.e. by setting π above π^e. The idea is that the central bank can make more money available to banks which can lend it out to companies. This boosts output but it also increases inflation because companies will respond to the boom partly by raising prices. Rising inflation is disliked by policy makers because it devalues pensions and savings.

The model captures nicely the dilemma of the central bank. It has just one instrument with which it tries to control to objectives - maximizing output and minimizing inflation.

What inflation rate will the central bank set? We can substitute for y and get:

$$W(\pi, \pi^e) = -c\pi^2 - [(b-1)y^* + d(\pi - \pi^e)]^2 \tag{12}$$

We can maximize with respect to π and find that the best response of the central bank to inflationary expectations π^e is:

$$\pi^*(\pi^e) = \frac{d}{c + d^2}\left[(1-b)y^* + d\pi^e\right] \tag{13}$$

Note:

- Although the government likes zero inflation it will not be able to implement it even if the private sector expects zero inflation. The government cannot resist to boost inflation a little bit in order to increase output.

In the SPE the government will play its best response to inflationary expectations. The private sector can predict the government's response and will set inflationary expectation such that:

$$\pi^S = \pi^e = \pi^*(\pi^e) \tag{14}$$

We can solve this fixed point and obtain:

$$\pi^e = \frac{d(1-b)}{c} y^* \tag{15}$$

Remark 5 *Note, that the government is* ***not*** *able to create surprise inflation in equilibrium because the private sector will anticipate the government's response. The government would be unequivocally better off if it could commit to zero inflation ex ante - in both cases it would create output by* but with commitment it could improve its utility by having lower inflation.*

Remark 6 *The government can decrease inflation by increasing c. At first, this might sound funny: isn't c a parameter of the political process which cannot be easily manipulated? True - but government can still (i) hire a tough central banker who* ***only*** *cares about inflation lowering and (ii) make the central bank independent such that the central banker rather than the government determines inflation. These two policies effectively increase c: the government can commit to lower inflation by pursing central bank independence plus hiring a tough banker.*

Remark 7 *The last insight is extremely important: a decision maker who has a commitment problem can often benefit from (credibly) delegating decision making to an agent who has different preferences than he has. Note, that no decision maker would ever do this if he were faced with a simple decision problem: the agent would make different choices which would decrease the decision maker's utility. However, in a game this negative direct effect might be outweighed by the commitment effect.*

4.2 Repeated Barro-Gordon Model

A different solution to the government's commitment problem is a repeated interaction setting. Assume that the private sector and the government play the Barro-Gordon game in many periods (call them 'years').

Then look at the following grim-trigger type strategies:

- The private sector in period 1 and all subsequent periods in which no deviation has occurred previously sets $\pi^e = 0$.
- The government in period 1 and all subsequent periods in which no deviation has occurred previously sets $\pi = 0$.
- If there is any deviation from this pattern both the government and the central bank revert to the equilibrium in the stage game.

We can use the SPDP to check that this is an equilibrium. The only condition which is critical is the government's willingness to stick to the equilibrium and not inflate the economy. If it sticks to the equilibrium it gets $\frac{1}{1-\delta}W(0,0)$. If it deviates the government gets at best $W(\pi^*(0),0) + \frac{\delta}{1-\delta}W(\pi^S,\pi^S)$ (assuming that the government chooses its best deviation). So we want:

$$\frac{1}{1-\delta}W(0,0) \geq W(\pi^*(0),0) + \frac{\delta}{1-\delta}W(\pi^S,\pi^S) \tag{16}$$

Remark 8 *In order to use repeated game effects to solve the government's commitment problem the government has to give itself a sufficiently long time horizon (beyond the 4-5 year term) which is equivalent to making sure that it is sufficiently patient (δ large). One way is to give central bankers very long terms such that their discount factor is sufficiently large.*

5 "Gifts"

We slightly modify the standard repeated Prisoner's Dilemma model. This model gives a rationale for why people give gifts which is a common phenomenon in almost every culture. It explains a few other things as well:

- The model also explains why gifts should be objects rather than money.

- The model predicts that the best gifts are those which are bought in over-priced gift-shops at Harvard Square and that they should be (and look) expensive but have little utility to the recipient.

- Finally, the model predicts that burning time in a pointless and boring lunch or dinner with a business acquaintance is actually time well spent (or burnt).

5.1 Prisoner's Dilemma with Rematching

We consider the following environment:

- There are infinitely many agents.

- Time is discrete and agents discount time at rate δ. At the beginning of each period unaligned agents can match with a random partner (another non-aligned agent).

- Agents in a partnership play a repeated Prisoner's dilemma.

- At the end of each period a relationship breaks up with exogenous probability s (separation rate). This ensures that there is always a supply of agents willing to match.

- Agents can also break off a relationship unilaterally at the end of a period.

- Agents can only observe the history of interactions with the current partner - in particular they cannot observe the past behavior of an agent with whom they just matched. It does not really matter if they keep remembering the history of agents with whom they previously matched - the idea is that once a relationship breaks up there is essentially zero chance of meeting up with the same partner again. Therefore, it does not matter if agents remember the history of interactions with past partners.

This model is a model of finding and breaking up with 'friends'. There are many potential friends out there and the goal is to play a repeated game of cooperation with a partner.

Assume the following simple payoff structure for the stage game:

	C	D
C	3,3	−1,4
D	4,−1	0,0

Now look at one relationship in isolation. Cooperation between two friends is sustainable through a grim-trigger strategy as long as:

$$4 \leq \frac{3}{1 - \delta(1 - s)} \tag{17}$$

This implies that $\delta > \frac{1}{4(1-s)}$.

However, there is a problem: an agent could defect in the first period, break immediately off his relationship and then rematch the next period to play another game of cooperation with the next partner. The ability to rematch makes cooperation an unsustainable equilibrium. Is there a way out?

5.2 Gifts

There is a simple solution to the sequential cheating problem. Both players could start a relationship by 'burning' some utility $g > 0$ - call this a gift. Gifts have to be costly to the sender but useless to the recipient (or at least less useful than the cost to the sender). Intuitively, gifts establish that an agent is 'serious' about the relationship. They also introduce a friction into the rematching process which makes it costly for an agent to break off a relationship. If the following two conditions hold gifts can ensure the existence of a cooperative equilibrium amongst friends:

1. The gift g has to fulfill $g \geq 1$. By breaking off and rematching the agent gets a stream of utility $4 - g$ in each period versus 3 by staying in the relationship with a partner until it breaks up exogenously.

2. Continuing a relationship has to be more profitable than breaking it up and staying unaligned:

$$4 \leq \frac{3}{1 - \delta(1 - s)} \tag{18}$$

This gives the same inequality $\delta \geq \frac{1}{4(1-s)}$.

3. Finally, starting a relationship in the first place has to be more profitable than not starting a relationship:

$$\frac{3}{1 - \delta(1 - s)} - g \geq 0 \tag{19}$$

This will be the case unless agents are too impatient or the separation probability s is too large.

Remark 9 *Good presents are expensive but useless - hence money is a bad present.*

Remark 10 *An alternative to gifts is to delay the onset of a new relationship. For example, new relationships might require a probationary period of several periods during which partners cannot rematch and cannot cooperate. This is equivalent to both of them 'burning' utility. Hence this models of gifts might also be interpreted as an explanation for rituals and social norms in which partners can only start to do business with each after a certain period of time. While such a strategy would be inefficient in a world in which partners cannot break up it makes sense if there is the risk of breakup and partners realigning with others. Probation periods make such exploitative strategies costly and hence less attractive.*

Lecture XV: Games with Incomplete Information

Markus M. Möbius

April 28, 2004

- Gibbons, chapter 3
- Osborne, chapter 9

1 Introduction

Informally, a game with incomplete information is a game where the game being played is not common knowledge. This idea is tremendously important in practice where its almost always a good idea to assume that something about the game is unknown to some players. What could be unknown?

1. **Payoffs:** In a price or quantity competition model you may know that your rival maximizes profits but now what his costs are (and hence his profits).

2. **Identity of other players:** R&D race between drug companies - who else will come up with the same drug?

3. **What moves are possible:** What levels of quality can rivals in a quality competition choose?

4. **How does the outcome depend on action:** Workers work/shirk don't know probability of getting caught because product fails.

2 Examples

2.1 Example I: Crazy Incumbent

Think of a standard entry game where the incumbent is 'crazy' with probability $1-p$ and rational with probability p. The normal incumbent faces the standard entry game from lecture 11:

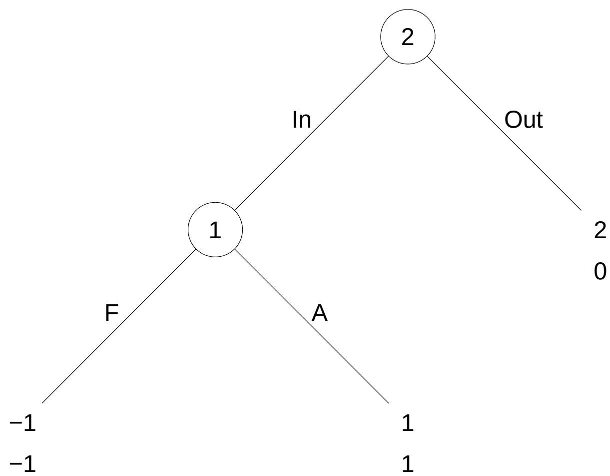

If the incumbent is crazy he will always want to fight because is is facing a different subgame:

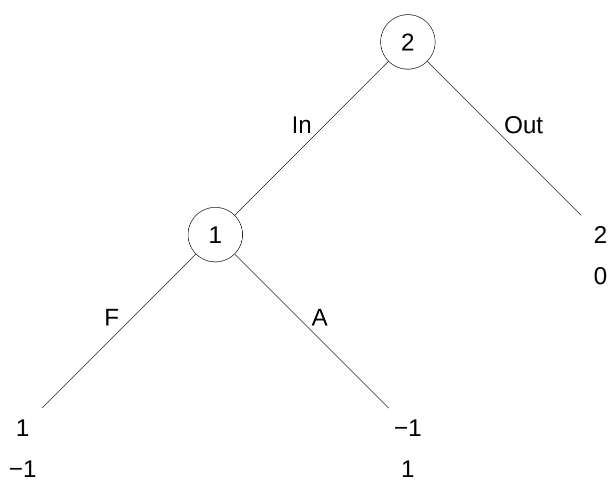

2.2 Example II: Auction

Two bidders are trying to purchase the same item at a sealed bid auction. The players simultaneously choose b_1 and b_2 and the good is sold to the highest bidder at his bid price (assume coin flip if $b_1 = b_2$). Suppose that the players' utilities are

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

The crucial incomplete information is that while each player knows his own valuation, he does not know his rival's. Assume, each has a prior that his rival's valuation is uniform on $[0, 1]$ and that this is common knowledge. Is hisTRUE?

2.3 Example III: Public Good

Two advisors of a graduate student each want the student to get a job at school X. Each can ensure this by calling someone on the phone and lying about how good the student is. Suppose the payoffs are as shown because each advisor gets utility 1 from the student being employed but has a cost of making the phone call.

	Call	Dont
Call	$1-c_1, 1-c_2$	$1-c_1, 1$
Dont	$1, 1-c_2$	0,0

Assume that the actions are chosen simultaneously and that players know only their own costs. They have prior that $c_{-i} \in U[\underline{c}, \overline{c}]$.

Alternatively, we could have player 1's cost known to all ($c_1 = \frac{1}{2}$) but $c_2 \in \{\underline{c}, \overline{c}\}$ known only to player 2.

Or, player 1 is a senior faculty member who knows from experience the cost of such phone calls ($c_1 = \frac{1}{2}$,$c_2 = \frac{2}{3}$). Player 2 is new assistant professor who has priors $c_1, c_2 \in U[0, 2]$.

3 Definitions

Definition 1 *A game with incomplete information* $G = (\Phi, S, P, u)$ *consists of*

1. *A set* $\Phi = \Phi_1 \times ... \times \Phi_I$ *where* Φ_i *is the (finite) set of possible types for player* i.[1]
2. *A set* $S = S_1 \times ... \times S_I$ *giving possible strategies for each player.*
3. *A joint probability distribution* $p(\phi_1, .., \phi_I)$ *over the types. For finite type space assume* $p(\phi_i) > 0$ *for all* $\phi_i \in \Phi_i$.
4. *A payoff function* $u_i : S \times \Phi \rightarrow \Re$.

It's useful to discuss the types in each of our examples.

- **Example I:** $\Phi_1 = \text{normal}$, $\Phi_2 \in \{\pi \text{ maximizer}, \text{crazy}\}$
- **Example II:**$\Phi_1 = \Phi_2 = (0, 1)$
- **Example III:**$\Phi_1 = \Phi_2 = [\underline{c}, \overline{c}]$

Note that payoffs can depend not only on your type but also on your rival's type as well.

Players know their own types but not the other players' types.

To analyze games of incomplete information we rely on the following observation (Harsanyi):

Observation: *Games of incomplete information can be thought of as games of complete but imperfect information where nature makes the first move and not everyone is informed about nature's move, i.e. nature chooses* Φ *but only reveals* ϕ_i *to player* i.

[1]Note, that the player i knows his type.

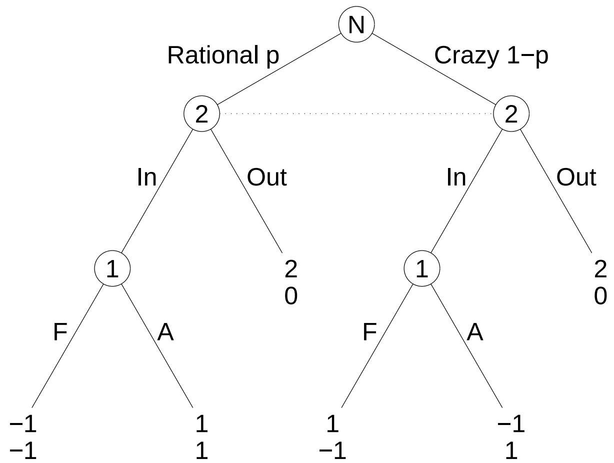

Think of nature simply as another player who rather than maximizing chooses a fixed mixed strategy.

This observation should make all of the following definitions seem completely obvious. They just say that to analyze these games we may look at NE of the game with Nature as another player.

Definition 2 *A Bayesian strategy for player i in G is a function $f_i : \Phi_i \to \Sigma_i$. Write S^{Φ_i} for the set of Bayesian strategies.*[2]

Definition 3 *A Bayesian strategy profile $(f_1^*, .., f_I^*)$ is a Bayesian Nash equilibrium if*

$$f_i^* \in \arg\max_{f_i \in S_i^{\Phi_i}} \sum_{\phi_i, \phi_{-i}} u_i\left(f_i(\phi_i), f_{-i}^*(\phi_{-i}); \phi_i, \phi_{-i}\right) p(\phi_i, \phi_{-i})$$

for all i or equivalently if for all i,ϕ_i,s_i

$$\sum_{\phi_{-i}} u_i\left(f_i^*(\phi_i), f_{-i}^*(\phi_{-i}); \phi_i, \phi_{-i}\right) p(\phi_i, \phi_{-i}) \geq \sum_{\phi_{-i}} u_i\left(s_i, f_{-i}^*(\phi_{-i}); \phi_i, \phi_{-i}\right) p(\phi_i, \phi_{-i})$$

This just says that your maximize expected payoff, and given that you know your type (that all have positive probability) this is equivalent to saying you maximize conditional on each possible type.

Remark 1 *A Bayesian Nash equilibrium is simply a Nash equilibrium of the game where Nature moves first, chooses $\phi \in \Phi$ from a distribution with probability $p(\phi)$ and reveals ϕ_i to player i.*

[2]Note, that a Bayesian strategy is simply an extensive form strategy where each type is treated as a distinct information set.

4 Solved examples

4.1 Public good I

Suppose player 1 is known to have cost $c_1 < \frac{1}{2}$. Player 2 has cost $\underline{c}$ with probability p, $\overline{c}$ with probability $1-p$. Assume that $0 < \underline{c} < 1 < \overline{c}$ and $p < \frac{1}{2}$.

	Call	Dont
Call	$1-c_1, 1-c_2$	$1-c_1, 1$
Dont	$1, 1-c_2$	0,0

Then the unique BNE is:

- $f_1^* = \text{Call}$
- $f_2^*(c) = \text{Don't Call}$ for all c.

Proof: In a BNE each type of player must play a BR so for the type $\overline{c}$ of player 2 calling is strictly dominated:

$$u_2(s_1, \text{call}; \overline{c}) < u_2(s_1, \text{Dont}; \overline{c})$$

for all s_1.

So $f_2^*(\overline{c}) = \text{Dont}$.

For player 1:

$$\begin{aligned} u_1(\text{Call}, f_2^*; c_1) &= 1 - c_1 \\ u_1(\text{Dont}, f_2^*; c_1) &= p u_1(\text{Dont}, f_2^*(\underline{c}); c_1) + (1-p) u_1(\text{Dont}, \text{Dont}; c_1) \\ &\leq p + (1-p)0 = p \end{aligned}$$

Because $1 - c_1 > p$ we know $f_1^*(c_1) = \text{Call}$.

For the type $\underline{c}$ of player 2 we have:

$$\begin{aligned} u_2(f_1^*, \text{Call}; \underline{c}) &= 1 - \underline{c} \\ u_2(f_1^*, \text{Dont}; \underline{c}) &= 1 \end{aligned}$$

because $f_1^* = Call$. So $f_2^*(\underline{c}) = \text{Dont}$.

This indicates this is the only possible BNE and our calculations have verified that it does actually work.

Process I've used is like iterated dominance.

4.2 Public Goods II

In the public good game suppose that c_1 and c_2 are drawn independently from a uniform distribution on $[0, 2]$. Then the (essentially) unique BNE is

$$f_i^*(c_i) = \begin{cases} \text{Call} & \text{if } c_i \leq \frac{2}{3} \\ \text{Don't} & \text{if } c_i > \frac{2}{3} \end{cases}$$

Proof: **Existence** is easy - just check that each is using a BR given that his rival calls with probability $\frac{1}{3}$ and doesn't call with probability $\frac{2}{3}$.

I'll show **uniqueness** to illustrate how to find the equilibrium.

Observation: *If $f_i^*(c_i) = Call$ then $f_i^*(c_i') = Call$ for all $c_i' < c_i$.*

To see this write z_{-i} for $Prob\left\{f_{-i}^*(c_{-i}) = \text{Call}\right\}$. If $f_i^*(c_i) = \text{call}$ then

$$\begin{aligned} E_{c_{-i}} u_i\left(\text{Call}, f_{-i}^*(c_{-i}); c_i\right) &\geq E_{c_{-i}} u_i\left(\text{Dont}, f_{-i}^*(c_{-i}); c_i\right) \\ 1 - c_i &\geq z_{-i} \end{aligned}$$

This clearly implies that for $c_i' < c_i$

$$E_{c_{-i}} u_i\left(\text{Call}, f_{-i}^*(c_{-i}); c_i'\right) = 1 - c_i' > E_{c_{-i}} u_i\left(\text{Dont}, f_{-i}^*(c_{-i}); c_i'\right) = z_{-i}$$

The intuition is that calling is more attractive if the cost is lower.

In light of observation a BNE must be of the form:[3]

$$f_1^*(c_1) = \begin{cases} \text{Call} & \text{if } c_1 \le c_1^* \\ \text{Don't} & \text{if } c_1 > c_1^* \end{cases}$$
$$f_2^*(c_2) = \begin{cases} \text{Call} & \text{if } c_2 \le c_2^* \\ \text{Don't} & \text{if } c_2 > c_2^* \end{cases}$$

For these strategies to be a BNE we need:

$$\begin{aligned} 1 - c_i &\ge z_{-i} \quad \text{for all } c_i < c_i^* \\ 1 - c_i &\le z_{-i} \quad \text{for all } c_i > c_i^* \end{aligned}$$

Hence $1 - c_i^* = z_{-i}$.

Because c_{-i} is uniform on $[0, 2]$ we get $z_{-i} = Prob\left\{c_{-i} < c_{-i}^*\right\} = \frac{c_{-i}^*}{2}$.

We have:

$$\begin{aligned} 1 - c_1^* &= \frac{c_2^*}{2} \\ 1 - c_2^* &= \frac{c_1^*}{2} \end{aligned}$$

It is easy to see that the unique solution to this system of equations is $c_1^* = c_2^* = \frac{2}{3}$.

Remark 2 *The result here is common in public goods situations. We get inefficient underinvestment because of the free rider problem. Each wants the other to call.*

[3] The cutoff values themselves are indeterminate - agents might or might not call. In this sense the equilibrium won't be unique. However, the cutoff values are probability zero events and hence the strategy at these points won't matter.

Lecture XVI: Auctions

Markus M. Möbius

May 6, 2004

- Gibbons, chapter 3
- Osborne, chapter 9
- Paul Klemperer's website at http://www.paulklemperer.org/ has fantastic online material on auctions and related topics.

1 Introduction

We already introduced a private-value auction with two bidders last time as an example for a game of imperfect information. In this lecture we expand this definition a little bit.

> In all our auctions there are n participants and each participant has a valuation v_i and submits a bid b_i (his action).

$$\Big\Downarrow$$

> The rules of the auction determine the probability $q_i(b_1, .., b_n)$ that agent i wins the auction and the expected price $p_i(b_1, .., b_n)$ which he pays. His utility is simple $u_i = q_i v_i - p_i$.[a]
>
> [a] The agent is risk-neutral - new issues arise if the bidders are risk-averse.

There are two important dimensions to classify auctions which are based on this template:

1. How are values v_i drawn? In the *private value* environment each v_i is drawn independently from some distribution F_i and support $[\underline{v}, \overline{v}]$. For our purposes we assume that all bidders are symmetric such that the v_i are i.i.d. draws from a common distribution F.[1] In the *correlated private value environment* the v_i are not independent - for example if I have a large v_i the other players are likely to have a large value as well. In the *pure common value* environment all bidders have the same valuation $v_i = v_j$.[2]

2. What are the rules? In the *first price auction* the highest bid wins and the bidder pays his bid. In the case of a tie a fair coin is flipped to determine the winner.[3] In the *second price auction* the highest bidder wins but pays the second-highest bid. In the *all-pay* auction all bidders pay and the highest bidder wins. All these three auctions are static games. The most famous dynamic auction is the *English auction* where the price starts at zero and starts to rise. Bidders drop out until the last remaining bidder gets the auction.[4] The Dutch auction is the opposite of the English - the price starts high and decreases until the first bidder jumps in to buy the object. The Dutch auction is strategically equivalent to the first-price auction. Note, that in May 2004 Google decided to use a Dutch auction for its IPO.

We will usually assume symmetric private value environments.

---

[1]Typically we assume that the distribution is continuous and has no atoms.

[2]There are much more general environments. A very general formulation which encompasses both private and common value auctions is due to Wilson (1977). Each bidder gets a signal t_i about her valuation which is drawn from some distribution $g_i(t_i, s)$ where s is a common noise term for all players. The players value is then a function $v(t_i, s)$. If $v_i = t_i$ we get back the private value environment. Similarly, if $v_i = v(s)$ we have the pure common value environment.

[3]In most equilibria ties are no problem because they occur with zero probability.

[4]The second price auction with private values is very similar to the English auction where the price starts at zero and increases continuously until the last bidder drops out (he pays essentially the second highest bid)

2 Solving Common Private Value Auctions

2.1 First-Price Auction

We focus on monotonic equilibria $b_i = f_i(v_i)$ such that f_i is strictly increasing (one can show this always holds). It will be simplest to work just with two bidders but the method can easily be extended to many bidders. We also assume for simplicity that values are drawn i.i.d. from the uniform distribution over $[0,1]$ (this can also be generalized - the differential equation becomes more complicated then).

The probability of player i winning the auction with bidding b is

$$Prob(f_j(v_j) \leq b) = Prob(v_j \leq f_j^{-1}(b)) = F(f_j^{-1}(b)) = f_j^{-1}(b) \tag{1}$$

The last equation follows because F is the uniform distribution.

The expected utility from bidding b is therefore:

$$Prob(f_j(v_j) \leq b)(v_i - b) = f_j^{-1}(b)(v_i - b) \tag{2}$$

The agent will choose b to maximize this utility. We differentiate with respect to b and use the first-order condition:

$$\frac{1}{f_j'(f_j^{-1}(b))}(v_i - b) - f_j^{-1}(b) = 0 \tag{3}$$

From now on we focus on symmetric equilibria such that $f_i = f_j$. We know that in equilibrium $b = f_i(v_i)$ such that:

$$\frac{1}{f'(v_i)}(v_i - f(v_i)) - v_i = 0 \tag{4}$$

This is a differential equation and can be rewritten as follows:

$$v_i = v_i f'(v_i) + f(v_i) \tag{5}$$

We can integrate both sides and get:

$$\frac{1}{2}v_i^2 + K = v_i f(v_i) \tag{6}$$

where K is a constant. This gives us finally:

$$f(v_i) = \frac{1}{2}v_i + \frac{K}{v_i} \tag{7}$$

You can check that $f(0) = 0$ - the player with a zero valuation should never bid positive amounts. Hence $K = 0$ is the only possible solution.

If you solve this exercise more generally for n bidders you get the following bidding function (uniform distribution):

$$f(v_i) = \frac{n-1}{n} v_i \tag{8}$$

This makes sense - as you increase the number of bidders there is more competition for the good and players have to make higher bids. Also, note that all bidders shade down their bid - otherwise they would not make profits.

2.2 All-Pay Auction

The all-pay auction is simple to analyze after the work we did on the first-price auction. We assume again the same setup. The corresponding utility function becomes:

$$Prob(f_j(v_j) \leq b)v_i - b = f_j^{-1}(b)v_i - b \tag{9}$$

The corresponding differential equation is:

$$\frac{1}{f'(v_i)} v_i - 1 = 0 \tag{10}$$

That means that the only solution is:

$$f(v_i) = \frac{1}{2} v_i^2 \tag{11}$$

The general solution for n bidders is:

$$f(v_i) = \frac{n-1}{n} v_i^n \tag{12}$$

2.3 Second-Price Auction

This auction is different because it has a much more robust solution:

Theorem 1 *In the second-price auction with private values (both independent and correlated) bidding one's own valuation is a weakly dominant strategy.*

This means that *no matter what the other players do* you can never do worse by bidding your own valuation $b_i = v_i$.

Proof: Denote the highest bid of all the other players except i by $\hat{b}$. Can i gain by deviating from bidding $b_i = v_i$? Assume that i bids higher such that $b_i > v_i$. This will only make a difference to the outcome of the auction for i if $v_i < \hat{b} < b_i$ in which case i will win the object now with the higher bid. But the utility from doing so is $v_i - \hat{b} < 0$ because i has to pay $\hat{b}$. Hence this is not profitable. Similarly, bidding below v_i is also non-profitable. QED

The intuition is that in the second-price auction my bid does not determine the price I pay. In the other two auction my bid equals my price. This makes me want to shade down my price - but if my bid does not affect the price but only my probability of winning then there is no reason to shade it down.

3 The Revenue-Equivalence-Theorem

THIS SECTION IS NOT EXAM-RELEVANT!

How much revenue does the auctioneer make from the auction (the total sum of payments he receives from all players)?

The expected revenue is equal to n times the expected payment from each player. Hence to compare the revenue of different auction formats we simply have to calculate the expected payment from each bidder with valuation v_i.

Let's look at the first-price, second-price and all-bid auctions with two bidders and uniform distribution.

- In the first price auction the expected payment from a player with valuation v_i is his bid $\frac{1}{2}v_i$ times the probability that he will win the auction which is v_i. Hence his expected payment is $\frac{1}{2}v_i^2$.

- In the second price auction i pays the second highest bid if he wins. Since the other player bids his valuation we know that the second highest bid is uniformly distributed over $[0, v_i]$ (conditional on i winning). Hence the expected payment from i conditional of winning is $\frac{1}{2}v_i$. The expected payment is this conditional payment times the probability of winning and is $\frac{1}{2}v_i^2$.

- In the all-pay auction player i always pays $\frac{1}{2}v_i^2$ which is equal to her expected payment.

Surprisingly, the revenue from all three auction formats is identical!

This is a special case of the revenue equivalence theorem.

Theorem 2 ***Revenue Equivalence Theorem.*** *In the symmetric independent private value case all auctions which allocate the good to the player with the highest value for it and which give zero utility to a player with valuation $\underline{v}$ are revenue equivalent.*

This theorem even applies to strange and unusual auctions such as a run-off following:

- There is a first-round auction where the five highest bidders are selected in a Dutch auction.

- Those bidders face each other in a run-off all bid auction.

Even this two-part auction will give the same revenue to the auctioneer in all equilibria where the good eventually goes to the player with the highest value.

Remark 1 *All monotonic and symmetric equilibria will satisfy the property that the highest value bidder gets the object.*

Proof of RET: Any auction mechanism which allocates the good will give player i with valuation v_i some surplus $S(v_i)$:

$$S(v_i) = q_i(v_i)v_i - p_i(v_i) \tag{13}$$

where $q_i(v_i)$ is the probability of winning with valuation v_i and $p_i(v_i)$ is the expected price. We know by assumption that $S(\underline{v}) = 0$.

Now note, that player i could pretend to be a different type $\tilde{v}$ and imitate $\tilde{v}$'s strategy (since only i knows his type this would be possible). This deviation would give surplus:

$$\hat{S}(\tilde{v}, v_i) = q_i(\tilde{v})v_i - p_i(\tilde{v}) \tag{14}$$

Now it has to be the case that i would not want to imitate type $\tilde{v}$. Hence we get:

$$\left.\frac{\partial S(\tilde{v}, v_i)}{\partial \tilde{v}}\right|_{\tilde{v}=v_i} = 0 \tag{15}$$

Now we can calculate the derivative of the surplus function:

$$\frac{dS(v_i)}{dv_i} = \frac{d\hat{S}(v_i, v_i)}{dv_i} = \underbrace{\left.\frac{\partial S(\tilde{v}, v_i)}{\partial \tilde{v}}\right|_{\tilde{v}=v_i}}_{=0} + \left.\frac{\partial S(\tilde{v}, v_i)}{\partial v_i}\right|_{\tilde{v}=v_i} = q_i(v_i) = F(v_i)^{n-1} \tag{16}$$

Finally, this gives:

$$S(v_i) = S(\underline{v}) + \int_{\underline{v}}^{v_i} F(t)^{n-1} dt = \int_{\underline{v}}^{v_i} F(t)^{n-1} dt \tag{17}$$

Hence the expected surplus for each player is identical across all auctions. But that also implies that the expected payments from each player are identical across auctions. QED

4 A Fun Common-Value Auction

"Auctioning off a Dollar" is a nice way to make some money from innocent friends who have not taken game theory. It's a simple all-pay auction game with a single common value.

Dollar Game: The auctioneer has one Dollar. There are n players. Each of them can bid for the Dollar and the highest bidder wins the dollar but each bidder has to pay hid bid.

How can we solve for the equilibrium in this game? Note, that this is not an incomplete information game because the value of the good is known to both players in advance.

- It's easy to see that there is no pure-strategy NE.
- Mixed equilibria have to have identical support over $[0, 1]$ (check!!). You can also show that the distribution F has to be continuous (no atoms).
- Since players get zero utility from bidding 0 it has to be the case that all bids give zero utility (otherwise players would not mix across them).
- One can easily show that all equilibria are symmetric. But assuming this for now we get the probability of winning the dollar when bidding b is $F(b)^{n-1}$. Hence we have:

$$F(b)^{n-1} - b = 0 \tag{18}$$

From this we get that $F(b) = b^{\frac{1}{n-1}}$. With two players this equals the uniform distribution. With more players this distribution is more concentrated around 0 as one would expect because competition for the Dollar heats up and in an all-pay auction bidders would shade down their bids more strongly.

5 Winner's Curse

The winner's curse arises in first-price auctions with a common value environment such as the following:

- All players get a signal t_i. This signal is drawn from a distribution $g(t_i, v)$ where v is the common value of the object.
- A simple way to generate such an environment is $t_i = v + \epsilon_i$ - every player observes the value with some noise. This is a good model where oil companies bid for an oil tract. Since the price of oil is set at the world market all companies have roughly the same valuation for the oil well. But they cannot observe the oil directly - they can only get seismic evaluations which gives an unbiased signal of the true value of oil in the ground. If one averages across all of these evaluations one gets a pretty good estimate of v.

A common mistake people make in playing this game is to bid too close to their signal t_i.

- Assume you bid exactly your signal $b_i = t_i = v + \epsilon_i$.
- In this case the most optimistic bid wins (very high ϵ) and the winner makes a loss.
- Engineers observed this phenomenon in the 1950s/60s in Texas - winning bidders would make negative profits.
- The phenomenon is called winner's curse because winning is 'bad' news. Winning means that one has the highest ϵ_i. To counteract the winner's curse one has to shade down one's bid appropriately.

Lecture XVII: Dynamic Games with Incomplete Information

Markus M. Möbius

May 6, 2004

- Gibbons, sections 4.1 and 4.2
- Osborne, chapter 10

1 Introduction

In the last two lectures I introduced the idea of incomplete information. We analyzed some important simultaneous move games such as sealed bid auctions and public goods.

In practice, almost all of the interesting models with incomplete information are dynamic games also. Before we talk about these games we'll need a new solution concept called Perfect Bayesian Equilibrium.

Intuitively, PBE is to extensive form games with incomplete games what SPE is to extensive form games with complete information. The concept we did last time, BNE is a simply the familiar Nash equilibrium under the Harsanyi representation of incomplete information. In principle, we could use the Harsanyi representation and SPE in dynamic games of incomplete information. However, dynamic games with incomplete information typically don't have enough subgames to do SPE. Therefore, many 'non-credible' threats are possible again and we get too many unreasonable SPE's. PBE allows subgame reasoning at information sets which are not single nodes whereas SPE only applies at single node information sets of players (because only those can be part of a proper subgame).

The following example illustrates some problems with SPE.

1.1 Example I - SPE

Our first example has no incomplete information at all.

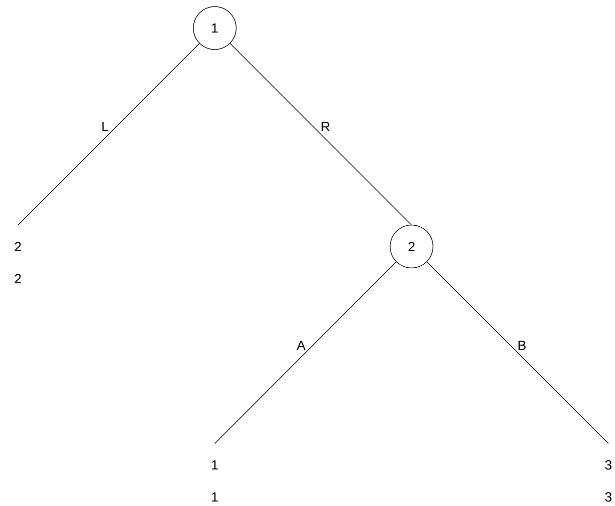

Its unique SPE is (R,B).

The next game looks formally the same - however, SPE is the same as NE because the game has no proper subgames.

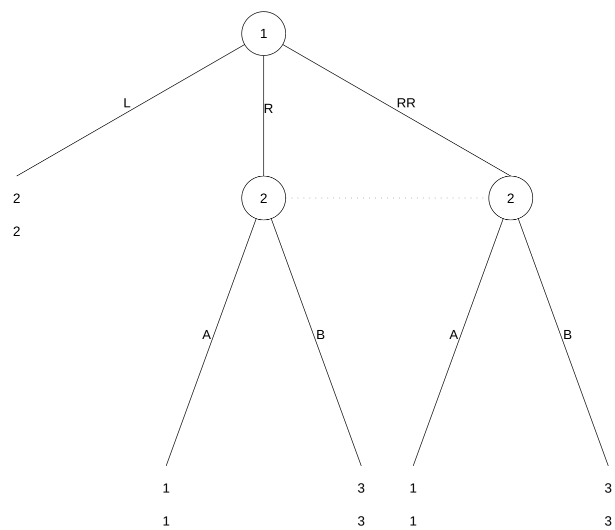

The old SPE survives - all $(pR + (1-p)\,RR, B)$ for all p is SPE. But there are suddenly strange SPE such as $(L, qA + (1-q)\,B)$ for $q \geq \frac{1}{2}$. Player 2's strategy looks like an non-credible threat again - but out notion of SPE can't rule it out!

Remember: *SPE can fail to rule out actions which are not optimal given any 'beliefs' about uncertainty.*

Remark 1 *This problem becomes severe with incomplete information: moves of Nature are not observed by one or both players. Hence the resulting extensive form game will have no or few subgames. This and the above example illustrate the need to replace the concept of a 'subgame' with the concept of a 'continuation game'.*

1.2 Example II: Spence's Job-Market Signalling

The most famous example of dynamic game with incomplete information is Spence's signalling game. There are two players - a firm and a worker. The worker has some private information about his ability and has the option of acquiring some education. Education is always costly, but less so for more able workers. **However, education does not improve the worker's productivity at all!** In Spence's model education merely serves as a signal to firms. His model allows equilibria where able workers will acquire education and less able workers won't. Hence firms will pay high wages only to those who acquired education - however, they do this because education has revealed the type of the player rather than improved his productivity.

Clearly this is an extreme assumption - in reality education has presumably dual roles: there is some signalling and some productivity enhancement. But it is an intriguing insight that education might be nothing more than a costly signal which allows more able workers to differentiate themselves from less able ones.

Let's look at the formal set-up of the game:

- **Stage 0:** Nature chooses the ability θ of a worker. Suppose $\Theta = \{2, 3\}$ and that $Prob\,(\theta = 2) = p$ and $Prob\,(\theta = 3) = 1 - p$.

- **Stage I:** Player 1 (worker) observes his type and chooses eduction level $e \in \{0, 1\}$. Education has cost $\frac{ce}{\theta}$. Note, that higher ability workers have lower cost and that getting no education is costless.

- **Stage II:** Player 2 (the competitive labor market) chooses the wage rate $w\,(e)$ of workers after observing the education level.

Suppose that $u_1(e, w, ; \theta) = w - \frac{ce}{\theta}$ and that $u_2(e, w; \theta) = -(w - \theta)^2$. Note, that education does not enter the firm's utility function. Also note, that the BR of the firm is to set wages equal to the expected ability of the worker under this utility function. This is exactly what the competitive labor market would do if θ is equal to the productivity of a worker (the dollar amount of output he produces). If a firm pays above the expected productivity it will run a loss, and if it pays below some other firm would come in and offer more to the worker. So the market should offer exactly the expected productivity. The particular (rather strange-looking) utility function we have chosen implements the market outcome with a single firm - it's a simple shortcut.

Spence's game is a *signalling game*. Each signalling game has the same three-part structure: nature chooses types, the sender (worker) observes his type and takes and action, the receiver (firm) sees that action but not the worker's type. Hence the firm tries to deduce the worker's type using his action. His action therefore serves as a signal. Spence's game is extreme because the signal (education) has no value to the firm except for its signalling function. This is not the case for all signalling models: think of a car manufacturer who can be of low or high quality and wants to send a signal to the consumer that he is a high-quality producer. He can offer a short or long warranty for the car. The extended warranty will not only signal his type but also benefit the consumer.

The (Harsanyi) extensive form representation of Spence's game (and any other signalling game) is given below.

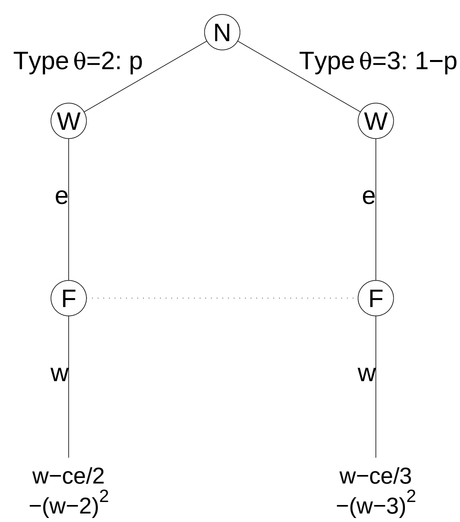

1.3 Why does SPE concept together with Harsanyi representation not work?

We could find the set of SPE in the Harsanyi representation of the game. The problem is that the game has no proper subgame in the second round when the firm makes its decision. Therefore, the firm can make unreasonable threats such as the following: both workers buy education and the firms pays the educated worker $w = 3 - p$ (his expected productivity), and the uneducated worker gets $w = -235.11$ (or something else). Clearly, every worker would get education, and the firm plays a BR to a worker getting education (check for yourself using the Harsanyi representation).

However, the threat of paying a negative wage is unreasonable. Once the firm sees a worker who has no education it should realize that the worker has a least ability level 2 and should therefore at least get a wage of $w = 2$.

2 Perfect Bayesian Equilibrium

Let G be a multistage game with incomplete information and observed actions in the Harsanyi representation. Write Θ_i for the set of possible types for player i and H_i to be the set of possible information sets of player i. For each information set $h_i \in H_i$ denote the set of nodes in the information set with $X_i(h_i)$ and $X \bigcup_{H_i} X_i(h_i)$.

A strategy in G is a function $s_i : H_i \to \Delta(A_i)$. Beliefs are a function $\mu_i : H_i \to \Delta(X_i)$ such that the support of belief $\mu_i(h_i)$ is within $X_i(h_i)$.

Definition 1 *A PBE is a strategy profile s^** ***together with*** *a belief system μ such that*

1. *At every information set strategies are optimal given beliefs and opponents' strategies (**sequential rationality**).*

$$\sigma_i^*(h) \quad \text{maximizes} \quad E_{\mu_i(x|h_i)} u_i\left(\sigma_i, \sigma_{-i}^* | h, \theta_i, \theta_{-i}\right)$$

2. *Beliefs are always updated according to Bayes rule **when applicable**.*

The first requirement replaces subgame perfection. The second requirement makes sure that beliefs are derived in a rational manner - assuming that you know the other players' strategies you try to derive as many beliefs as possible. Branches of the game tree which are reached with zero probability cannot be derived using Bayes rule: here you can choose arbitrary beliefs. However, the precise specification will typically matter for deciding whether an equilibrium is PBE or not.

Remark 2 *In the case of complete information and observed actions PBE reduces to SPE because beliefs are trivial: each information set is a singleton and the belief you attach to being there (given that you are in the corresponding information set) is simply 1.*

2.1 What's Bayes Rule?

There is a close connection between agent's actions and their beliefs. Think of job signalling game. We have to specify the beliefs of the firm in the second stage when it does not know for sure the current node, but only the information set.

Let's go through various strategies of the worker:

- *The high ability worker gets education and the low ability worker does not:* $e(\theta = 2) = 0$ and $e(\theta = 3) = 1$. In this case my beliefs at the information set $e = 1$ should be $Prob(\text{High}|e = 1) = 1$ and similarly, $Prob(\text{High}|e = 0) = 0$.

- *Both workers get education.* In this case, we should have:

$$Prob(\text{High}|e = 1) = 1 - p \tag{1}$$

  The beliefs after observing $e = 0$ cannot be determined by Bayes rule because it's a probability zero event - we should never see it if players follow their actions. This means that we can choose beliefs freely at this information set.

- *The high ability worker gets education and the low ability worker gets education with probability q.* This case is less trivial. What's the probability of seeing worker get education - it's $1 - p + pq$. What's the probability of a worker being high ability and getting education? It's $1 - p$. Hence the probability that the worker is high ability after we have observed him getting education is $\frac{1-p}{1-p+pq}$. This is the non-trivial part of Bayes rule.

Formally, we can derive the beliefs at some information set h_i of player i as follows. There is a probability $p(\theta_j)$ that the other player is of type θ_j. These probabilities are determined by nature. Player j (i.e. the worker) has taken some action a_j such that the information set h_i was reached. Each type of player j takes action a_j with some probability $\sigma_j^*(a_j|\theta_j)$ according to his equilibrium strategy. Applying Bayes rule we can then derive the belief of player i that player j has type θ_j at information set h_i:

$$\mu_i(\theta_j|a_j) = \frac{p(\theta_j)\,\sigma_j^*(a_j|\theta_j)}{\sum_{\tilde{\theta}_j \in \Theta_j} p\left(\tilde{\theta}_j\right)\sigma_j^*\left(a_j|\tilde{\theta}_j\right)} \tag{2}$$

1. In the job signalling game with separating beliefs Bayes rule gives us exactly what we expect - we belief that a worker who gets education is high type.

2. In the pooling case Bayes rule gives us $Prob\,(\text{High}|e=1) = \frac{1-p}{p\times 1+(1-p)\times 1} = 1-p$. **Note, that Bayes rule does NOT apply for finding the beliefs after observing $e=0$ because the denominator is zero.**

3. In the semi-pooling case we get $Prob\,(\text{High}|e=1) = \frac{(1-p)\times 1}{p\times q+(1-p)\times 1}$. Similarly, $Prob\,(\text{High}|e=0) = \frac{(1-p)\times 0}{p\times(1-q)+(1-p)\times 0} = 0$.

3 Signalling Games and PBE

It turns out that signalling games are a very important class of dynamic games with incomplete information in applications. Because the PBE concept is much easier to state for the signalling game environment we define it once again in this section for signalling games. You should convince yourself that the more general definition from the previous section reduces to the definition below in the case of signalling games.

3.1 Definitions and Examples

Every signalling game has a sender, a receiver and two periods. The sender has private information about his type and can take an action in the first action. The receiver observes the action (signal) but not the type of the sender, and takes his action in return.

- **Stage 0:** Nature chooses the type $\theta_1 \in \Theta_1$of player 1 from probability distribution p.
- **Stage 1:** Player 1 observes θ_1 and chooses $a_1 \in A_1$.
- **Stage 2:** Player 2 observes a_1 and chooses $a_2 \in A_2$.

The players utilities are:

$$u_1 = u_1(a_1, a_2; \theta_1) \tag{3}$$
$$u_2 = u_2(a_1, a_2; \theta_1) \tag{4}$$

3.1.1 Example 1: Spence's Job Signalling Game

- worker is sender; firm is receiver
- θ is the ability of the worker (private information to him)
- $A_1 = \{\text{educ}, \text{no educ}\}$
- $A_2 = \{\text{wage rate}\}$

3.1.2 Example 2: Initial Public Offering

- player 1 - owner of private firm
- player 2 - potential investor
- Θ - future profitability
- A_1 - fraction of company retained
- A_2 - price paid by investor for stock

3.1.3 Example 3: Monetary Policy

- player 1 = FED
- player 2 - firms
- Θ - Fed's preference for inflation/ unemployment
- A_1 - first period inflation
- A_2 - expectation of second period inflation

3.1.4 Example 4: Pretrial Negotiation

- player 1 - defendant
- player 2 - plaintiff
- Θ - extent of defendant's negligence
- A_1 - settlement offer
- A_2 - accept/reject

3.2 PBE in Signalling Games

A PBE in the signalling game is a strategy profile $(s_1^*(\theta_1), s_2^*(a_1))$ together with beliefs $\mu_2(\theta_1|a_1)$ for player 2 such that

1. Players strategies are optimal given their beliefs and the opponents' strategies, i.e.

$$s_1^*(\theta_1) \quad \text{maximizes} \quad u_1(a_1, s_2^*(a_1); \theta_1) \text{ for all } \theta_1 \in \Theta_1 \tag{5}$$

$$s_2^*(a_1) \quad \text{maximizes} \quad \sum_{\theta_1 \in \Theta_1} u_2(a_1, a_2; \theta_1) \mu_2(\theta_1|a_1) \quad \text{for all } a_1 \in A_1 \tag{6}$$

2. Player 2's beliefs are compatible with Bayes' rule. If any type of player 1 plays a_1 with positive probability then

$$\mu_2(\theta_1|a_1) = \frac{p(\theta_1) \, Prob(s_1^*(\theta_1) = a_1)}{\sum_{\theta_1' \in \Theta_1} p(\theta_1') \, Prob(s_1^*(\theta_1') = a_1)} \quad \text{for all } \theta_1 \in \Theta_1$$

3.3 Types of PBE in Signalling Games

To help solve for PBE's it helps to think of all PBE's as taking one of the following three forms"

1. **Separating** - different types take different actions and player 2 learns type from observing the action
2. **Pooling** - all types of player 1 take same action; no info revealed
3. **Semi-Separating** - one or more types mixes; partial learning (often only type of equilibrium

Remark 3 *In the second stage of the education game the "market" must have an expectation that player 1 is type $\theta = 2$ and attach probability $\mu(2|a_1)$ to the player being type 2. The wage in the second period* ***must be*** *between 2 and 3. This rules out the unreasonable threat of the NE I gave you in the education game (with negative wages).*[1]

Remark 4 *In the education game suppose the equilibrium strategies are $s_1^*(\theta = 2) = 0$ and $s_1^*(\theta = 3) = 1$, i.e. only high types get education. Then for any prior $(p, 1-p)$ at the start of the game beliefs must be:*

$$\begin{aligned}
\mu_2(\theta = 2|e = 0) &= 1 \\
\mu_2(\theta = 3|e = 0) &= 0 \\
\mu_2(\theta = 2|e = 1) &= 0 \\
\mu_2(\theta = 3|e = 1) &= 1
\end{aligned}$$

If player 1's strategy is $s_1^*(\theta = 2) = \frac{1}{2} \times 0 + \frac{1}{2} \times 1$ and $s_1^*(\theta = 3) = 1$:

$$\begin{aligned}
\mu_2(\theta = 2|e = 0) &= 1 \\
\mu_2(\theta = 3|e = 0) &= 0 \\
\mu_2(\theta = 2|e = 1) &= \frac{\frac{p}{2}}{\frac{p}{2} + 1 - p} = \frac{p}{2-p} \\
\mu_2(\theta = 3|e = 1) &= \frac{2-2p}{2-p}
\end{aligned}$$

Also note, that Bayes rule does NOT apply after an actions which should not occur in equilibrium. Suppose $s_1^*(\theta = 2) = s_1^*(\theta = 3) = 1$ then it's OK to assume

$$\begin{aligned}
\mu_2(\theta = 2|e = 0) &= \frac{57}{64} \\
\mu_2(\theta = 3|e = 0) &= \frac{7}{64} \\
\mu_2(\theta = 2|e = 1) &= p \\
\mu_2(\theta = 3|e = 1) &= 1-p
\end{aligned}$$

The first pair is arbitrary.

---

[1] It also rules out unreasonable SPE in the example SPE I which I have initially. Under any beliefs player 2 should strictly prefer B.

Remark 5 *Examples SPE II and SPE III from the introduction now make sense - if players update according to Bayes rule we get the 'reasonable' beliefs of players of being with equal probability in one of the two nodes.*

4 Solving the Job Signalling Game

Finally, after 11 tough pages we can solve our signalling game. The solution depends mainly on the cost parameter c.

4.1 Intermediate Costs $2 \leq c \leq 3$

A separating equilibrium of the model is when only the able worker buys education and the firm pays wage 2 to the worker without education and wage 3 to the worker with education. The firm beliefs that the worker is able iff he gets educated.

- The beliefs are consistent with the equilibrium strategy profile.
- Now look at optimality. Player 2 sets the wage to the expected wage so he is maximizing.
- Player 1 of type $\theta = 2$ gets $3 - \frac{c}{2} \leq 2$ for $a_1 = 1$ and 2 for $a_1 = 0$. Hence he should not buy education.
- Player 1 of type $\theta = 3$ gets $3 - \frac{c}{3} \geq 2$ for $a_1 = 1$ and 2 for $a_1 = 0$. Hence he should get educated.

1. Note that for too small or too big costs there is no separating PBE.
2. There is no separating PBE where the $\theta = 2$ type gets an education and the $\theta = 3$ type does not.

4.2 Small Costs $c \leq 1$

A pooling equilibrium of the model is that both workers buy education and that the firm pays wage $w = 3 - p$ if it observes education, and wage 2 otherwise. The firm believes that the worker is able with probability $1 - p$ if it observes education, and that the worker is of low ability if it observes no education.

- The beliefs are consistent with Bayes rule for $e = 1$. If $e = 0$ has been observed Bayes rule does not apply because $e = 0$ should never occur - hence any belief is fine. The belief that the worker is low type if he does not get education makes sure that the worker gets punished for not getting educated.

- The firm pays expected wage - hence it's optimal response. The low ability guy won't deviate as long $2.5 - \frac{c}{2} \geq 2$ and the high ability type won't deviate as long as $2.5 - \frac{c}{3} \geq 2$. For $c \leq 1$ both conditions are true.

1. While this pooling equilibrium only works for small c there is always another pooling equilibrium where no worker gets education and the firms thinks that any worker who gets education is of the low type.

4.3 $1 < c < 2$

Assume that $p = \frac{1}{2}$ for this section. In the parameter range $1 < c < 2$ there is a semiseparating PBE of the model. The high ability worker buys education and the low ability worker buys education with positive probability q. The wage is $w = 2$ if the firm observes no education and set to $w = 2 + \frac{1}{1+q}$ if it observes education. The beliefs that the worker is high type is zero if he gets no education and $\frac{1}{1+q}$ if he does.

- Beliefs are consistent (check!).

- Firm plays BR.

- Player 1 of low type won't deviate as long as $2 + \frac{1}{1+q} - \frac{c}{2} \leq 2$.

- Player 1 of high type won't deviate as long as $2 + \frac{1}{1+q} - \frac{c}{3} \geq 2$.

Set $1 + q = \frac{2}{c}$. It's easy to check that the first condition is binding and the second condition is strictly true. So we are done if we choose $q = \frac{2}{c} - 1$. Note, that as $c \to 2$ we get back the separating equilibrium and as $c \to 1$ we get the pooling one.

Lecture XVIII: Games with Incomplete Information II - More Examples

Markus M. Möbius

May 6, 2004

- Gibbons, section 4.2
- Osborne, chapter 10

1 Introduction

This lecture gives more examples of games of incomplete information, in particular signalling games.

2 The Lobbying Game

We consider the following model of lobbying.

- Nature chooses whether the lobbyist's industry is headed for Good or Bad times and reveals the state of the world $\{G, B\}$ to the lobbyist.
- The a priori probability of Good times is p.
- The Lobbyist can then send a message to Congress. Following this message, Congress chooses whether or not to enact a subsidy. Let $\{S, N\}$ denote the actions available to Congress. At the end of the period, the state of the world is revealed to Congress.
- A subsidy costs Congress k. It generates a return $r > k$ for Congress if and only if times are Bad. We assume that $(1 - p)r < k$.

- The Lobbyist gets a payoff of zero if the subsidy is not passed, a payoff of 1 if the subsidy passes and times are Bad, and a subsidy of 1/2 if the subsidy passes and times are good.

1. **Is there a PBE in which the subsidy passes? Answer: NO**

2. **Now suppose lobbying Congress is costly. In particular, the lobbyist must incur a cost c to be heard. Show that if $c > 1/2$, there is a PBE in which the subsidy passes whenever the state of the world is bad.**

3 Legal Settlements

- There are two players, a plaintiff and a defendant in a civil suit. The plaintiff knows whether or not he will win the case if he goes to trial, but the defendant does not have this information.

- The defendant knows that the plaintiff knows who would win, and the defendant has prior beliefs that there is probability $\frac{1}{3}$ that the plaintiff will win; these prior beliefs are common knowledge.

- If the plaintiff wins, his payoff is 3 and the defendant's payoff is -4; if the plaintiff loses, his payoff is -1 and the defendant's is 0. (This corresponds to the defendant paying cash damages of 3 if the plaintiff wins, and the loser of the case paying court costs of 1.)

- The plaintiff has two possible actions: He can ask for either a low settlement of $m = 1$ or a high settlement of $m = 2$. If the defendant accepts a settlement offer of m, the plaintiff's payoff is m and the defendant's is $-m$. If the defendant rejects the settlement offer, the case goes to court.

1. List all the pure-strategy PBE strategy profiles. For each such profile, specify the beliefs of the defendant as a function of m, and verify that the combination of these beliefs and the profile is in fact a PBE.

2. Explain why the other profiles are not PBE.

4 Corporate Investment

This is a variant of the IPO game. An entrepreneur needs financing to realize a new project and can offer an outside investor an equity stake in his company. The stake gives the investor a share of the future (next period's) cashflow of the company: the profits from the existing business plus the profits from the project. The profitability of the nw project is known to both investor and entrepreneur. However, only the entrepreneur knows the profitability of the existing company. The investor therefore runs the risk of investing in an unprofitable business. The new project requires investment I and gives payoff $R > (1 + r) I$ in the next period (where r is the interest rate).

- Nature chooses the type of the entrepreneur's firm which can be highly profitable ($\pi = H$) or less profitable ($\pi = L$). The business is not so profitable with probability p.
- The entrepreneur observes π and then offers equity stake s such that $0 \leq s \leq 1$.
- The investor can accept or reject.

5 Monetary Policy

5.1 The Barro-Gordon model (1983)

There are two periods. In the first period firms form expectations about inflation π_e. Their payoff is $-(\pi - \pi_e)^2$. In the second period the Fed sets inflation π. The Fed has objective function:

$$-c\pi^2 - (y - y^*)^2 \tag{1}$$

Actual output y is:

$$y = by^* + d(\pi - \pi^*) \tag{2}$$

where $b < 1$.

5.2 Barro-Gordon with Signalling

Now assume that the Fed can either be weak ($c = W$) or strong ($c = S$) such that $S > W > 0$. The firm and the Fed play the Barro-Gordon game now for two periods: the first period can now be used to signal the resolve of the Fed.